



(51) International Patent Classification:

G06F 11/30 (2006.01) H04L 9/32 (2006.01)  
G06F 15/16 (2006.01) H04L 12/24 (2006.01)  
G06F 21/00 (2013.01) H04L 12/26 (2006.01)

(21) International Application Number:

PCT/US2016/040094

(22) International Filing Date:

29 June 2016 (29.06.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/136,687 22 April 2016 (22.04.2016) US  
15/136,762 22 April 2016 (22.04.2016) US

(71) Applicant: SOPHOS LIMITED [GB/GB]; The Pentagon, Abingdon Science Park, Abingdon OX14 3YP (GB).

(72) Inventors; and

(71) Applicants (for US only): ACKERMAN, Karl [US/US]; Sophos, Inc., 3 Van De Graaf Drive, Burlington, MA 01803 (US). HARRIS, Mark, David [GB/GB]; Sophos Limited, The Pentagon, Abingdon Science Park, Abingdon OX14 3YP (GB). REED, Simon, Neil [GB/GB]; Sophos Limited, The Pentagon, Abingdon Science Park, Abingdon OX14

3YP (GB). THOMAS, Andrew, J. [GB/GB]; Sophos Limited, The Pentagon, Abingdon Science Park, Abingdon OX14 3YP (GB). RAY, Kenneth, D. [US/US]; Sophos, Inc., 3 Van De Graaf Drive, Burlington, MA 01803 (US). STUTZ, Daniel [DE/DE]; Sophos Limited, The Pentagon, Abingdon Science Park, Abingdon OX14 3YP (GB). HOWARD, Fraser [GB/GB]; Sophos Limited, The Pentagon, Abingdon Science Park, Abingdon OX14 3YP (GB). SAMOSSEIKO, Dmitri [CA/CA]; Sophos Limited, The Pentagon, Abingdon Science Park, Abingdon OX14 3YP (GB).

(74) Agent: MAZZARESE, Robert, A. et al.; Strategic Patents, P.C., P.O. Box 920629, Needham, MA 02492 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(54) Title: DETECTING TRIGGERING EVENTS FOR DISTRIBUTED DENIAL OF SERVICE ATTACKS

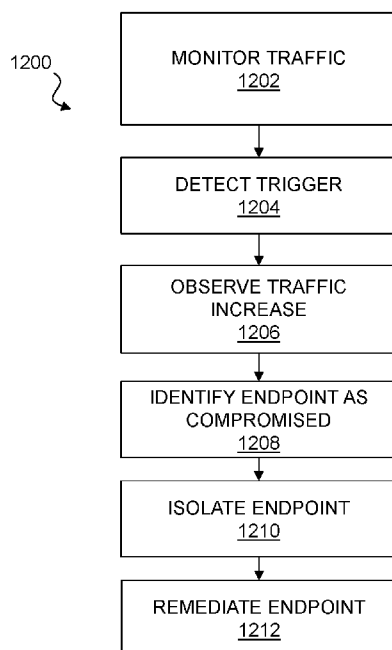


FIG. 12

(57) Abstract: An endpoint in an enterprise network is monitored, and when a potential trigger for a distributed denial of service (DDoS) attack is followed by an increase in network traffic from the endpoint to a high reputation network address, the endpoint is treated as a DDoS service bot and isolated from the network until remediation can be performed.



**(84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Published:**

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

## DETECTING TRIGGERING EVENTS FOR DISTRIBUTED DENIAL OF SERVICE ATTACKS

### TECHNICAL FIELD

**[0001]** This application relates to techniques for improving security of communications with endpoints in an enterprise network.

### BACKGROUND

**[0002]** Enterprise networks can contain valuable information that forms an increasingly attractive target for malicious actors. Useful techniques for securing endpoints in a network against malicious activity are described by way of example in commonly-owned U.S. Pat. App. No. 14/263,955, filed on April 28, 2014, U.S. Pat. App. No. 14/485,759, filed on September 14, 2014, and U.S. Pat. App. No. 15/042,862 filed on February 12, 2016, each of which is hereby incorporated by reference in its entirety.

**[0003]** There remains a need for improved security of communications to and from endpoints in an enterprise network.

### SUMMARY

**[0004]** An endpoint in an enterprise network is monitored, and when a potential trigger for a distributed denial of service (DDoS) attack is followed by an increase in network traffic from the endpoint to a high reputation network address, the endpoint is treated as a DDoS service bot and isolated from the network until remediation can be performed.

**[0005]** A computer program product for protecting against distributed denial of service attacks from an enterprise network may include computer executable code embodied in a non-transitory computer readable medium that, when executing on one or more computing devices, performs the steps of monitoring outbound traffic from an endpoint in the enterprise network, detecting a potential trigger event on the endpoint for a distributed denial of service attack, observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event, in response to the increase in network traffic, identifying the endpoint as having a compromised state in which the endpoint serves as a distributed denial of service bot for the distributed denial of service attack, and preventing network traffic from the endpoint until the compromised state is remediated.

**[0006]** Implementations may include one or more of the following features. The potential trigger event may include a detection of a process on the endpoint that monitors a clock. The potential trigger event may include a detection of a process on the endpoint that monitors a file on the endpoint. The potential trigger event may include a detection of network traffic between the endpoint and a low reputation network address. The potential trigger event may include a detection of network traffic from a command and control site. The potential trigger event may include a detection of an increase in network traffic to the high reputation network address from one or more other endpoints in the enterprise network. Monitoring the outbound traffic may include monitoring the outbound traffic at a gateway for the enterprise network. Monitoring the outbound traffic may include monitoring the outbound traffic at the endpoint.

**[0007]** A method may include monitoring outbound traffic from an endpoint in an enterprise network, detecting a potential trigger event for a distributed denial of service attack, observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event, in response to the increase in network traffic, identifying the endpoint as having a compromised state wherein the endpoint serves as a distributed denial of service bot, and preventing network traffic from the endpoint until the compromised state is remediated.

**[0008]** Implementations may include one or more of the following features. The potential trigger event may include a detection of a process on the endpoint that monitors a clock. The potential trigger event may include a detection of a process on the endpoint that monitors a file on the endpoint. The potential trigger event may include a receipt of a message on the endpoint from a low reputation network address. The potential trigger event may include a receipt of a message on the endpoint from a command and control site. The potential trigger event may include a detection of an increase in network traffic to the high reputation network address from one or more other endpoints in the enterprise network. Monitoring the outbound traffic may include monitoring the outbound traffic at a gateway for the enterprise network. Monitoring the outbound traffic may include monitoring the outbound traffic at the endpoint. Observing an increase in network traffic may include observing an increase in network traffic beyond a predetermined threshold. The method may further include remediating the endpoint.

**[0009]** A threat management facility for managing an enterprise network may include a network interface for coupling the threat management facility to the enterprise network, a memory, and a processor. The processor may be configured by computer executable code stored

in the memory to protect against a distributed denial of service attack originating from within the enterprise network by performing the steps of monitoring outbound traffic from an endpoint in the enterprise network, detecting a potential trigger event on the network for a distributed denial of service attack, observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event, in response to the increase in network traffic, identifying the endpoint as having a compromised state wherein the endpoint serves as a distributed denial of service bot, and preventing network traffic from the endpoint until the compromised state is remediated. The processor may be further configured to initiate a remediation of the endpoint.

**[0010]** In the context of network activity by an endpoint in an enterprise network, malware detection is improved by using a combination of reputation information for a network address that is accessed by the endpoint with reputation information for an application on the endpoint that is accessing the network address. This information, when combined with a network usage history for the application, provides improved differentiation between malicious network activity and legitimate, user-initiated network activity.

**[0011]** A computer program product for protecting against malicious network activity in an enterprise network may include computer executable code embodied in a non-transitory computer readable medium that, when executing on one or more computing devices, performs the steps of detecting a connection from an application executing on an endpoint in the enterprise network to a network address outside the enterprise network, determining a reputation of the application using a first reputation lookup at a threat management facility for the enterprise network, determining a reputation of the network address using a second reputation lookup at the threat management facility, determining a network usage history for the application using a log of network activity maintained on the endpoint, evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application, and initiating remediation action on the endpoint when the compromised condition is detected.

**[0012]** Implementations may include one or more of the following features. The reputation of the application may include at least one of known malicious, suspect, unknown, or known good. The reputation of the network address may include at least one of known malicious, suspect, unknown, or known good. The reputation of the network address may be based on crowd-sourced information about the network address. The computer program product

may further include code that performs the step of monitoring network activity by the application and storing the network usage history for the application in the log on the endpoint. The computer program product may further include code that performs the step of associating a network communication to the network address by a service executing on the endpoint with the application when the application controls the network communication, and adding the network communication to the network usage history for the application. The one or more computing devices may include the endpoint. The one or more computing devices may include the threat management facility.

**[0013]** A method may include detecting a connection from an endpoint in an enterprise network to a network address, determining a reputation of an application using the connection, determining a reputation of the network address, determining a network usage history for the application, evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application, and initiating remediation action on the endpoint when the compromised condition is detected.

**[0014]** Implementations may include one or more of the following features. The reputation of the application may be requested from a threat management facility for the enterprise network. The reputation of the application may include at least one of known malicious, suspect, unknown, or known good. The reputation of the network address may be requested from a threat management facility for the enterprise network. The reputation of the network address may include at least one of known malicious, suspect, unknown, or known good. The network usage history for the application may be retrieved from a log on the endpoint. The reputation of the network address may be based on crowd-sourced information about the network address. The method may further include monitoring network activity by the application and storing the network usage history for the application in a log on the endpoint. The method may further include associating a network communication to the network address by a service executing on the endpoint with the application when the application controls the network communication, and adding the network communication to the network usage history for the application.

**[0015]** A threat management facility for managing an enterprise network may include a network interface for coupling the threat management facility to the enterprise network, a memory, and a processor configured by computer executable code stored in the memory to

protect against malicious network activity in the enterprise network by performing the steps of detecting a connection from an endpoint in an enterprise network to a network address, determining a reputation of an application using the connection, determining a reputation of the network address, determining a network usage history for the application, evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application, and initiating remediation action on the endpoint when the compromised condition is detected. The reputation of the application may include at least one of known malicious, suspect, unknown, or known good, and the reputation of the network address may include at least one of known malicious, suspect, unknown, or known good. The network usage history for the application may be retrieved from a log on the endpoint.

**[0016]** Protocol suites such as hypertext transfer protocol (HTTP) using secure socket layer (SSL) can facilitate secure network communications. When using this type of secure communication, network addresses are typically expressed as numeric internet protocol addresses rather than the human-readable uniform resource locators (URLs) that are entered into a browser address bar by a human user. This property can be exploited to differentiate between secure and insecure communications, and to detect certain instances where a malicious proxy has been deployed to intercept network traffic with an endpoint.

**[0017]** A computer program product for detecting a malicious proxy on an endpoint in an enterprise network may include computer executable code embodied in a non-transitory computer readable medium that, when executing on one or more computing devices, performs the steps of monitoring outbound traffic from the endpoint with remote network addresses outside the enterprise network, detecting use of a secure communication protocol with a request from the endpoint to one of the remote network addresses, identifying a plaintext network address within the request, and, in response to identifying a plaintext network address in the request, initiating remediation of a potentially maliciously local proxy on the endpoint.

**[0018]** Implementations may include one or more of the following features. The secure communication protocol may include a hypertext transfer protocol using secure socket layer or transport layer security. The plaintext network address may include an alphanumeric address other than an internet protocol address. The plaintext network address may include a uniform resource locator. Monitoring outbound traffic may include looking up a reputation for destinations of outbound communications. Monitoring outbound traffic may include monitoring

outbound traffic at a gateway for the enterprise network. Initiating remediation of the potentially malicious local proxy may include quarantining the endpoint until the potentially malicious local proxy can be removed. Initiating remediation of the potentially malicious local proxy may include reversing a malware identification for the potentially malicious proxy by identifying a non-malicious source of the local proxy. Initiating remediation of the potentially malicious local proxy may include verifying a malware identification for the potentially malicious local proxy by performing the steps of identifying a process that initiated the request, determining a reputation of the process, if the reputation of the process is low or unknown, confirming the malware identification, and, if the reputation of the process is good, confirming the malware identification only when a calling process for the process that has a low or unknown reputation.

**[0019]** A method may include monitoring outbound traffic from an endpoint in an enterprise network, detecting use of a secure communication protocol with a request from the endpoint, identifying a plaintext network address within the request, and in response to identifying a plaintext network address in the request, remediating a potentially maliciously local proxy on the endpoint.

**[0020]** Implementations may include one or more of the following features. The secure communication protocol may include a hypertext transfer protocol using secure socket layer or transport layer security. The plaintext network address may include an alphanumeric address other than an internet protocol address. The plaintext network address may include a uniform resource locator. Monitoring outbound traffic may include looking up a reputation for destinations of outbound communications. Monitoring outbound traffic may include monitoring outbound traffic at a gateway for the enterprise network. Initiating remediation of the potentially malicious local proxy may include quarantining the endpoint until the potentially malicious local proxy can be removed. Initiating remediation of the potentially malicious local proxy may include reversing a malware identification for the potentially malicious proxy by identifying a non-malicious source of the local proxy. Initiating remediation of the potentially malicious local proxy may include verifying a malware identification for the potentially malicious local proxy by performing the steps of identifying a process that initiated the request, determining a reputation of the process, if the reputation of the process is low or unknown, confirming the malware identification, and, if the reputation of the process is good, confirming the malware identification only when a calling process for the process that has a low or unknown reputation.

[0021] A device may include a network interface, a memory, and a processor. The processor may be configured by computer executable code stored in the memory to perform the steps of monitoring outbound traffic from an endpoint in an enterprise network, detecting use of a secure communication protocol with a request from the endpoint, identifying a plaintext network address within the request, and in response to identifying a plaintext network address in the request, remediating a potentially maliciously local proxy on the endpoint. The secure communication protocol may be a hypertext transfer protocol using secure socket layer or transport layer security.

#### BRIEF DESCRIPTION OF THE FIGURES

[0022] The foregoing and other objects, features and advantages of the devices, systems, and methods described herein will be apparent from the following description of particular embodiments thereof, as illustrated in the accompanying drawings. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the devices, systems, and methods described herein.

[0023] Fig. 1 illustrates an environment for threat management.

[0024] Fig. 2 illustrates a computer system.

[0025] Fig. 3 illustrates a threat management system.

[0026] Fig. 4 illustrates a system for behavioral tracking, coloring, and generation of indications of compromise (IOCs).

[0027] Fig. 5 illustrates a system for encryption management.

[0028] Fig. 6 illustrates a threat management system using heartbeats.

[0029] Fig. 7 shows an architecture for endpoint protection in an enterprise network security system.

[0030] Fig. 8 shows a method for labeling network flows.

[0031] Fig. 9 illustrates an Internet Protocol packet.

[0032] Fig. 10 illustrates a method for secure labeling of network flows.

[0033] Fig. 11 shows a process for verifying user presence on an endpoint.

[0034] Fig. 12 shows a method for mitigating distributed denial of service attacks from an enterprise network.

[0035] Fig. 13 shows a method for detecting endpoint compromise based on network usage history.

[0036] Fig. 14 shows useful malware detection insights based on the monitoring described with reference to Fig. 13.

[0037] Fig. 15 shows a method for local proxy detection.

#### DETAILED DESCRIPTION

[0038] Embodiments will now be described with reference to the accompanying figures, in which preferred embodiments are shown. The foregoing may, however, be embodied in many different forms and should not be construed as limited to the illustrated embodiments set forth herein.

[0039] All documents mentioned herein are hereby incorporated by reference in their entirety. References to items in the singular should be understood to include items in the plural, and vice versa, unless explicitly stated otherwise or clear from the context. Grammatical conjunctions are intended to express any and all disjunctive and conjunctive combinations of conjoined clauses, sentences, words, and the like, unless otherwise stated or clear from the context. Thus, the term “or” should generally be understood to mean “and/or” and so forth.

[0040] Recitation of ranges of values herein are not intended to be limiting, referring instead individually to any and all values falling within the range, unless otherwise indicated herein, and each separate value within such a range is incorporated into the specification as if it were individually recited herein. The words “about,” “approximately,” or the like, when accompanying a numerical value, are to be construed as indicating a deviation as would be appreciated by one of ordinary skill in the art to operate satisfactorily for an intended purpose. Ranges of values and/or numeric values are provided herein as examples only, and do not constitute a limitation on the scope of the described embodiments. The use of any and all examples, or exemplary language (“e.g.,” “such as,” or the like) provided herein, is intended merely to better illuminate the embodiments and does not pose a limitation on the scope of the embodiments or the claims. No language in the specification should be construed as indicating any unclaimed element as essential to the practice of the embodiments.

[0041] In the following description, it is understood that terms such as “first,” “second,” “third,” “above,” “below,” and the like, are words of convenience and are not to be construed as implying a chronological order or otherwise limiting any corresponding element unless expressly state otherwise.

**[0042]** Fig. 1 illustrates an environment for threat management. Specifically, Fig. 1 depicts a block diagram of a threat management system providing protection to an enterprise against a plurality of threats – a context in which the following techniques may usefully be deployed. One aspect relates to corporate policy management and implementation through a unified threat management facility 100. As will be explained in more detail below, a threat management facility 100 may be used to protect computer assets from many threats, both computer-generated threats and user-generated threats. The threat management facility 100 may be multi-dimensional in that it may be designed to protect corporate assets from a variety of threats and it may be adapted to learn about threats in one dimension (e.g. worm detection) and apply the knowledge in another dimension (e.g. spam detection). Policy management is one of the dimensions for which the threat management facility can provide a control capability. A corporation or other entity may institute a policy that prevents certain people (e.g. employees, groups of employees, types of employees, guest of the corporation, etc.) from accessing certain types of computer programs. For example, the corporation may elect to prevent its accounting department from using a particular version of an instant messaging service or all such services. In this example, the policy management facility 112 may be used to update the policies of all corporate computing assets with a proper policy control facility or it may update a select few. By using the threat management facility 100 to facilitate the setting, updating and control of such policies the corporation only needs to be concerned with keeping the threat management facility 100 up to date on such policies. The threat management facility 100 can take care of updating all of the other corporate computing assets.

**[0043]** It should be understood that the threat management facility 100 may provide multiple services, and policy management may be offered as one of the services. We will now turn to a description of certain capabilities and components of the threat management system 100.

**[0044]** Over recent years, malware has become a major problem across the Internet 154. From both a technical perspective and a user perspective, the categorization of a specific threat type, whether as virus, worm, spam, phishing exploration, spyware, adware, or the like, is becoming reduced in significance. The threat, no matter how it is categorized, may need to be stopped at various points of a networked computing environment, such as one of an enterprise facility 102, including at one or more laptops, desktops, servers, gateways, communication ports, handheld or mobile devices, firewalls, and the like. Similarly, there may be less and less benefit

to the user in having different solutions for known and unknown threats. As such, a consolidated threat management facility 100 may need to apply a similar set of technologies and capabilities for all threats. In certain embodiments, the threat management facility 100 may provide a single agent on the desktop, and a single scan of any suspect file. This approach may eliminate the inevitable overlaps and gaps in protection caused by treating viruses and spyware as separate problems, while simultaneously simplifying administration and minimizing desktop load. As the number and range of types of threats has increased, so may have the level of connectivity available to all IT users. This may have led to a rapid increase in the speed at which threats may move. Today, an unprotected PC connected to the Internet 154 may be infected quickly (perhaps within 10 minutes) which may require acceleration for the delivery of threat protection. Where once monthly updates may have been sufficient, the threat management facility 100 may automatically and seamlessly update its product set against spam and virus threats quickly, for instance, every five minutes, every minute, continuously, or the like. Analysis and testing may be increasingly automated, and also may be performed more frequently; for instance, it may be completed in 15 minutes, and may do so without compromising quality. The threat management facility 100 may also extend techniques that may have been developed for virus and malware protection, and provide them to enterprise facility 102 network administrators to better control their environments. In addition to stopping malicious code, the threat management facility 100 may provide policy management that may be able to control legitimate applications, such as VoIP, instant messaging, peer-to-peer file-sharing, and the like, that may undermine productivity and network performance within the enterprise facility 102.

**[0045]** The threat management facility 100 may provide an enterprise facility 102 protection from computer-based malware, including viruses, spyware, adware, Trojans, intrusion, spam, policy abuse, uncontrolled access, and the like, where the enterprise facility 102 may be any entity with a networked computer-based infrastructure. In an embodiment, Fig. 1 may depict a block diagram of the threat management facility 100 providing protection to an enterprise against a plurality of threats. The enterprise facility 102 may be corporate, commercial, educational, governmental, or the like, and the enterprise facility's 102 computer network may be distributed amongst a plurality of facilities, and in a plurality of geographical locations, and may include administration 134, a firewall 138A, an appliance 140A, server 142A, network devices 148A–B, clients 144A–D, such as protected by computer security facilities 152, and the like. It will be understood that any reference herein to client facilities may include the

clients 144A–D shown in Fig. 1 and vice-versa. The threat management facility 100 may include a plurality of functions, such as security management facility 122, policy management facility 112, update facility 120, definitions facility 114, network access rules facility 124, remedial action facility 128, detection techniques facility 130, testing facility 118, threat research facility 132, and the like. In embodiments, the threat protection provided by the threat management facility 100 may extend beyond the network boundaries of the enterprise facility 102 to include clients 144D (or client facilities) that have moved into network connectivity not directly associated or controlled by the enterprise facility 102. Threats to client facilities may come from a plurality of sources, such as from network threats 104, physical proximity threats 110, secondary location threats 108, and the like. Clients 144A–D may be protected from threats even when the client 144A–D is not located in association with the enterprise 102, such as when a client 144E–F moves in and out of the enterprise facility 102, for example when interfacing with an unprotected server 142C through the Internet 154, when a client 144F is moving into a secondary location threat 108 such as interfacing with components 140B, 142B, 148C, 148D that are not protected, and the like. In embodiments, the threat management facility 100 may provide an enterprise facility 102 protection from a plurality of threats to multiplatform computer resources in a plurality of locations and network configurations, with an integrated system approach. It should be understood that an enterprise model is applicable to organizations and users of any size or type. For example, an enterprise may be or may include a group or association of endpoints, networks, users, and the like within or outside of one or more protected locations. It should be understood that an enterprise may include one or more offices or business locations, or one or more homes, where each location, or portions of each location, or a collection of locations may be treated as a client facility.

**[0046]** In embodiments, the threat management facility 100 may be provided as a stand-alone solution. In other embodiments, the threat management facility 100 may be integrated into a third-party product. An application programming interface (e.g. a source code interface) may be provided such that the threat management facility 100 may be integrated. For instance, the threat management facility 100 may be stand-alone in that it provides direct threat protection to an enterprise or computer resource, where protection is subscribed to directly 100. Alternatively, the threat management facility 100 may offer protection indirectly, through a third-party product, where an enterprise may subscribe to services through the third-party product, and threat

protection to the enterprise may be provided by the threat management facility 100 through the third-party product.

**[0047]** The security management facility 122 may include a plurality of elements that provide protection from malware to enterprise facility 102 computer resources, including endpoint security and control, email security and control, web security and control, reputation-based filtering, control of unauthorized users, control of guest and non-compliant computers, and the like. The security management facility 122 may be a software application that may provide malicious code and malicious application protection to a client facility computing resource. The security management facility 122 may have the ability to scan the client facility files for malicious code, remove or quarantine certain applications and files, prevent certain actions, perform remedial actions and perform other security measures. In embodiments, scanning the client facility may include scanning some or all of the files stored to the client facility on a periodic basis, scanning an application when the application is executed, scanning files as the files are transmitted to or from the client facility, or the like. The scanning of the applications and files may be performed to detect known malicious code or known unwanted applications. In an embodiment, new malicious code and unwanted applications may be continually developed and distributed, and updates to the known code database may be provided on a periodic basis, on a demand basis, on an alert basis, or the like.

**[0048]** The security management facility 122 may provide email security and control, where security management may help to eliminate spam, viruses, spyware and phishing, control of email content, and the like. The security management facility's 122 email security and control may protect against inbound and outbound threats, protect email infrastructure, prevent data leakage, provide spam filtering, and the like. In an embodiment, security management facility 122 may provide for web security and control, where security management may help to detect or block viruses, spyware, malware, unwanted applications, help control web browsing, and the like, which may provide comprehensive web access control enabling safe, productive web browsing. Web security and control may provide Internet use policies, reporting on suspect devices, security and content filtering, active monitoring of network traffic, URI filtering, and the like. In an embodiment, the security management facility 122 may provide for network access control, which may provide control over network connections. Network control may stop unauthorized, guest, or non-compliant systems from accessing networks, and may control network traffic that may not be bypassed from the client level. In addition, network access

control may control access to virtual private networks (VPN), where VPNs may be a communications network tunneled through another network, establishing a logical connection acting as a virtual network. In embodiments, a VPN may be treated in the same manner as a physical network.

**[0049]** The security management facility 122 may provide host intrusion prevention through behavioral based protection, which may guard against unknown threats by analyzing behavior before software code executes. Behavioral based protection may monitor code when it runs and intervene if the code is deemed to be suspicious or malicious. Advantages of behavioral based protection over runtime protection may include code being prevented from running. Whereas runtime protection may only interrupt code that has already partly executed, behavioral protection can identify malicious code at the gateway or on the file servers and delete the code before it can reach endpoint computers and the like.

**[0050]** The security management facility 122 may provide reputation filtering, which may target or identify sources of known malware. For instance, reputation filtering may include lists of URIs of known sources of malware or known suspicious IP addresses, or domains, say for spam, that when detected may invoke an action by the threat management facility 100, such as dropping them immediately. By dropping the source before any interaction can initiate, potential threat sources may be thwarted before any exchange of data can be made.

**[0051]** In embodiments, information may be sent from the enterprise back to a third party, a vendor, or the like, which may lead to improved performance of the threat management facility 100. For example, the types, times, and number of virus interactions that a client experiences may provide useful information for the preventions of future virus threats. This type of feedback may be useful for any aspect of threat detection. Feedback of information may also be associated with behaviors of individuals within the enterprise, such as being associated with most common violations of policy, network access, unauthorized application loading, unauthorized external device use, and the like. In embodiments, this type of information feedback may enable the evaluation or profiling of client actions that are violations of policy that may provide a predictive model for the improvement of enterprise policies.

**[0052]** The security management facility 122 may support overall security of the enterprise facility 102 network or set of enterprise facility 102 networks, e.g., by providing updates of malicious code information to the enterprise facility 102 network and associated client facilities. The updates may include a planned update, an update in reaction to a threat notice, an

update in reaction to a request for an update, an update based on a search of known malicious code information, or the like. The administration facility 134 may provide control over the security management facility 122 when updates are performed. The updates may be automatically transmitted without an administration facility's 134 direct control, manually transmitted by the administration facility 134, or otherwise distributed. The security management facility 122 may manage the receipt of malicious code descriptions from a provider, distribution of the malicious code descriptions to enterprise facility 102 networks, distribution of the malicious code descriptions to client facilities, and so forth.

**[0053]** The threat management facility 100 may provide a policy management facility 112 that may be able to block non-malicious applications, such as VoIP, instant messaging, peer-to-peer file-sharing, and the like, that may undermine productivity and network performance within the enterprise facility 102. The policy management facility 112 may be a set of rules or policies that may indicate enterprise facility 102 access permissions for the client facility, such as access permissions associated with the network, applications, external computer devices, and the like. The policy management facility 112 may include a database, a text file, a combination of databases and text files, or the like. In an embodiment, a policy database may be a block list, a black list, an allowed list, a white list, or the like that may provide a list of enterprise facility 102 external network locations/applications that may or may not be accessed by the client facility. The policy management facility 112 may include rules that may be interpreted with respect to an enterprise facility 102 network access request to determine if the request should be allowed. The rules may provide a generic rule for the type of access that may be granted. The rules may be related to the policies of an enterprise facility 102 for access rights for the enterprise facility's 102 client facility. For example, there may be a rule that does not permit access to sporting websites. When a website is requested by the client facility, a security facility may access the rules within a policy facility to determine if the requested access is related to a sporting website. In an embodiment, the security facility may analyze the requested website to determine if the website matches with any of the policy facility rules.

**[0054]** The policy management facility 112 may be similar to the security management facility 122 but with the addition of enterprise facility 102 wide access rules and policies that may be distributed to maintain control of client facility access to enterprise facility 102 network resources. The policies may be defined for application type, subset of application capabilities, organization hierarchy, computer facility type, user type, network location, time of day,

connection type, or the like. Policies may be maintained by the administration facility 134, through the threat management facility 100, in association with a third party, or the like. For example, a policy may restrict IM activity to only support personnel for communicating with customers. This may allow communication for departments requiring access, but may maintain the network bandwidth for other activities by restricting the use of IM to only the personnel that need access to instant messaging (IM) in support of the enterprise facility 102. In an embodiment, the policy management facility 112 may be a stand-alone application, may be part of the network server facility 142, may be part of the enterprise facility 102 network, may be part of the client facility, or the like.

**[0055]** The threat management facility 100 may provide configuration management, which may be similar to policy management, but may specifically examine the configuration set of applications, operating systems, hardware, and the like, and manage changes to their configurations. Assessment of a configuration may be made against a standard configuration policy, detection of configuration changes, remediation of improper configuration, application of new configurations, and the like. An enterprise may keep a set of standard configuration rules and policies which may represent the desired state of the device. For example, a client firewall may be running and installed, but in the disabled state, where remediation may be to enable the firewall. In another example, the enterprise may set a rule that disallows the use of USB disks, and sends a configuration change to all clients, which turns off USB drive access via a registry.

**[0056]** The threat management facility 100 may also provide for the removal of applications that potentially interfere with the operation of the threat management facility 100, such as competitor products that may also be attempting similar threat management functions. The removal of such products may be initiated automatically whenever such products are detected. In the case where such applications are services are provided indirectly through a third-party product, the application may be suspended until action is taken to remove or disable the third-party product's protection facility.

**[0057]** Threat management against a quickly evolving malware environment may require timely updates, and thus an update management facility 120 may be provided by the threat management facility 100. In addition, a policy management facility 112 may also require update management (e.g., as provided by the update facility 120 herein described). The update management for the security facility 122 and policy management facility 112 may be provided directly by the threat management facility 100, such as by a hosted system or in conjunction with

the administration facility 134. In embodiments, the threat management facility 100 may provide for patch management, where a patch may be an update to an operating system, an application, a system tool, or the like, where one of the reasons for the patch is to reduce vulnerability to threats.

**[0058]** The security facility 122 and policy management facility 112 may push information to the enterprise facility 102 network and/or client facility. The enterprise facility 102 network and/or client facility may also or instead pull information from the security facility 122 and policy management facility 112 network server facilities 142, or there may be a combination of pushing and pulling of information between the security facility 122 and the policy management facility 112 network servers 142, enterprise facility 102 network, and client facilities, or the like. For example, the enterprise facility 102 network and/or client facility may pull information from the security facility 122 and policy management facility 112 network server facility 142 may request the information using the security facility 122 and policy management facility 112 update module; the request may be based on a certain time period, by a certain time, by a date, on demand, or the like. In another example, the security facility 122 and policy management facility 112 network servers 142 may push the information to the enterprise facility's 102 network and/or client facility by providing notification that there are updates available for download and then transmitting the information. The combination of the security management 122 network server facility 142 and security update module may function substantially the same as the policy management facility 112 network server and policy update module by providing information to the enterprise facility 102 network and the client facility in a push or pull method. In an embodiment, the policy management facility 112 and the security facility 122 management update modules may work in concert to provide information to the enterprise facility's 102 network and/or client facility for control of application execution. In an embodiment, the policy update module and security update module may be combined into a single update module.

**[0059]** As threats are identified and characterized, the threat management facility 100 may create definition updates that may be used to allow the threat management facility 100 to detect and remediate the latest malicious software, unwanted applications, configuration and policy changes, and the like. The threat definition facility 114 may contain threat identification updates, also referred to as definition files. A definition file may be a virus identity file that may include definitions of known or potential malicious code. The virus identity (IDE) definition files

may provide information that may identify malicious code within files, applications, or the like. The definition files may be accessed by security management facility 122 when scanning files or applications within the client facility for the determination of malicious code that may be within the file or application. The definition files may contain a number of commands, definitions, or instructions, to be parsed and acted upon, or the like. In embodiments, the client facility may be updated with new definition files periodically to provide the client facility with the most recent malicious code definitions; the updating may be performed on a set time period, may be updated on demand from the client facility, may be updated on demand from the network, may be updated on a received malicious code alert, or the like. In an embodiment, the client facility may request an update to the definition files from an update facility 120 within the network, may request updated definition files from a computing facility external to the network, updated definition files may be provided to the client facility 114 from within the network, definition files may be provided to the client facility from an external computing facility from an external network, or the like.

**[0060]** A definition management facility 114 may provide timely updates of definition files information to the network, client facilities, and the like. New and altered malicious code and malicious applications may be continually created and distributed to networks worldwide. The definition files that maintain the definitions of the malicious code and malicious application information for the protection of the networks and client facilities may need continual updating to provide continual defense of the network and client facility from the malicious code and malicious applications. The definition files management may provide for automatic and manual methods of updating the definition files. In embodiments, the network may receive definition files and distribute the definition files to the network client facilities, the client facilities may receive the definition files directly, or the network and client facilities may both receive the definition files, or the like. In an embodiment, the definition files may be updated on a fixed periodic basis, on demand by the network and/or the client facility, as a result of an alert of a new malicious code or malicious application, or the like. In an embodiment, the definition files may be released as a supplemental file to an existing definition files to provide for rapid updating of the definition files.

**[0061]** In a similar manner, the security management facility 122 may be used to scan an outgoing file and verify that the outgoing file is permitted to be transmitted per the enterprise facility 102 rules and policies. By checking outgoing files, the security management facility 122

may be able discover malicious code infected files that were not detected as incoming files as a result of the client facility having been updated with either new definition files or policy management facility 112 information. The definition files may discover the malicious code infected file by having received updates of developing malicious code from the administration facility 134, updates from a definition files provider, or the like. The policy management facility 112 may discover the malicious code infected file by having received new updates from the administration facility 134, from a rules provider, or the like.

**[0062]** The threat management facility 100 may provide controlled access to the enterprise facility 102 networks. For instance, a manager of the enterprise facility 102 may want to restrict access to certain applications, networks, files, printers, servers, databases, or the like. In addition, the manager of the enterprise facility 102 may want to restrict user access based on certain criteria, such as the user's location, usage history, need to know, job position, connection type, time of day, method of authentication, client-system configuration, or the like. Network access rules may be developed for the enterprise facility 102, or pre-packaged by a supplier, and managed by the threat management facility 100 in conjunction with the administration facility 134.

**[0063]** A network access rules facility 124 may be responsible for determining if a client facility application should be granted access to a requested network location. The network location may be on the same network as the facility or may be on another network. In an embodiment, the network access rules facility 124 may verify access rights for client facilities from within the network or may verify access rights of computer facilities from external networks. When network access for a client facility is denied, the network access rules facility 124 may send an information file to the client facility containing. For example, the information sent by the network access rules facility 124 may be a data file. The data file may contain a number of commands, definitions, instructions, or the like to be parsed and acted upon through the remedial action facility 128, or the like. The information sent by the network access facility rules facility 124 may be a command or command file that the remedial action facility 128 may access and take action upon.

**[0064]** The network access rules facility 124 may include databases such as a block list, a black list, an allowed list, a white list, an unacceptable network site database, an acceptable network site database, a network site reputation database, or the like of network access locations that may or may not be accessed by the client facility. Additionally, the network access rules

facility 124 may incorporate rule evaluation; the rule evaluation may parse network access requests and apply the parsed information to network access rules. The network access rule facility 124 may have a generic set of rules that may be in support of an enterprise facility's 102 network access policies, such as denying access to certain types of websites, controlling instant messenger accesses, or the like. Rule evaluation may include regular expression rule evaluation, or other rule evaluation method for interpreting the network access request and comparing the interpretation to the established rules for network access. In an embodiment, the network access rules facility 124 may receive a rules evaluation request from the network access control and may return the rules evaluation to the network access control.

**[0065]** Similar to the threat definitions facility 114, the network access rule facility 124 may provide updated rules and policies to the enterprise facility 102. The network access rules facility 124 may be maintained by the network administration facility 134, using network access rules facility 124 management. In an embodiment, the network administration facility 134 may be able to maintain a set of access rules manually by adding rules, changing rules, deleting rules, or the like. Additionally, the administration facility 134 may retrieve predefined rule sets from a remote provider of a set of rules to be applied to an entire enterprise facility 102. The network administration facility 134 may be able to modify the predefined rules as needed for a particular enterprise facility 102 using the network access rules management facility 124.

**[0066]** When a threat or policy violation is detected by the threat management facility 100, the threat management facility 100 may perform or initiate a remedial action facility 128. Remedial action may take a plurality of forms, such as terminating or modifying an ongoing process or interaction, sending a warning to a client or administration facility 134 of an ongoing process or interaction, executing a program or application to remediate against a threat or violation, record interactions for subsequent evaluation, or the like. Remedial action may be associated with an application that responds to information that a client facility network access request has been denied. In an embodiment, when the data file is received, remedial action may parse the data file, interpret the various aspects of the data file, and act on the parsed data file information to determine actions to be taken on an application requesting access to a denied network location. In an embodiment, when the data file is received, remedial action may access the threat definitions to parse the data file and determine an action to be taken on an application requesting access to a denied network location. In an embodiment, the information received from the facility may be a command or a command file. The remedial action facility may carry out any

commands that are received or parsed from a data file from the facility without performing any interpretation of the commands. In an embodiment, the remedial action facility may interact with the received information and may perform various actions on a client requesting access to a denied network location. The action may be one or more of continuing to block all requests to a denied network location, a malicious code scan on the application, a malicious code scan on the client facility, quarantine of the application, terminating the application, isolation of the application, isolation of the client facility to a location within the network that restricts network access, blocking a network access port from a client facility, reporting the application to an administration facility 134, or the like.

**[0067]** Remedial action may be provided as a result of a detection of a threat or violation. The detection techniques facility 130 may include monitoring the enterprise facility 102 network or endpoint devices, such as by monitoring streaming data through the gateway, across the network, through routers and hubs, and the like. The detection techniques facility 130 may include monitoring activity and stored files on computing facilities, such as on server facilities 142, desktop computers, laptop computers, other mobile computing devices, and the like. Detection techniques, such as scanning a computer's stored files, may provide the capability of checking files for stored threats, either in the active or passive state. Detection techniques, such as streaming file management, may provide the capability of checking files received at the network, gateway facility, client facility, and the like. This may provide the capability of not allowing a streaming file or portions of the streaming file containing malicious code from entering the client facility, gateway facility, or network. In an embodiment, the streaming file may be broken into blocks of information, and a plurality of virus identities may be used to check each of the blocks of information for malicious code. In an embodiment, any blocks that are not determined to be clear of malicious code may not be delivered to the client facility, gateway facility, or network.

**[0068]** Verifying that the threat management facility 100 is detecting threats and violations to established policy, may require the ability to test the system, either at the system level or for a particular computing component. The testing facility 118 may allow the administration facility 134 to coordinate the testing of the security configurations of client facility computing facilities on a network. The administration facility 134 may be able to send test files to a set of client facility computing facilities to test the ability of the client facility to determine acceptability of the test file. After the test file has been transmitted, a recording

facility may record the actions taken by the client facility in reaction to the test file. The recording facility may aggregate the testing information from the client facility and report the testing information to the administration facility 134. The administration facility 134 may be able to determine the level of preparedness of the client facility computing facilities by the reported information. Remedial action may be taken for any of the client facility computing facilities as determined by the administration facility 134; remedial action may be taken by the administration facility 134 or by the user of the client facility.

**[0069]** The threat research facility 132 may provide a continuously ongoing effort to maintain the threat protection capabilities of the threat management facility 100 in light of continuous generation of new or evolved forms of malware. Threat research may include researchers and analysts working on known and emerging malware, such as viruses, rootkits a spyware, as well as other computer threats such as phishing, spam, scams, and the like. In embodiments, through threat research, the threat management facility 100 may be able to provide swift, global responses to the latest threats.

**[0070]** The threat management facility 100 may provide threat protection to the enterprise facility 102, where the enterprise facility 102 may include a plurality of networked components, such as client facility, server facility 142, administration facility 134, firewall 138, gateway, hubs and routers 148, threat management appliance 140, desktop users, mobile users, and the like. In embodiments, it may be the endpoint computer security facility 152, located on a computer's desktop, which may provide threat protection to a user, and associated enterprise facility 102. In embodiments, the term endpoint may refer to a computer system that may source data, receive data, evaluate data, buffer data, or the like (such as a user's desktop computer as an endpoint computer), a firewall as a data evaluation endpoint computer system, a laptop as a mobile endpoint computer, a personal digital assistant or tablet as a hand-held endpoint computer, a mobile phone as an endpoint computer, or the like. In embodiments, endpoint may refer to a source or destination for data, including such components where the destination is characterized by an evaluation point for data, and where the data may be sent to a subsequent destination after evaluation. The endpoint computer security facility 152 may be an application loaded onto the computer platform or computer support component, where the application may accommodate the plurality of computer platforms and/or functional requirements of the component. For instance, a client facility computer may be one of a plurality of computer platforms, such as Windows, Macintosh, Linux, and the like, where the endpoint computer

security facility 152 may be adapted to the specific platform, while maintaining a uniform product and product services across platforms. Additionally, components may have different functions to serve within the enterprise facility's 102 networked computer-based infrastructure. For instance, computer support components provided as hubs and routers 148, server facility 142, firewalls 138, and the like, may require unique security application software to protect their portion of the system infrastructure, while providing an element in an integrated threat management system that extends out beyond the threat management facility 100 to incorporate all computer resources under its protection.

[0071] The enterprise facility 102 may include a plurality of client facility computing platforms on which the endpoint computer security facility 152 is adapted. A client facility computing platform may be a computer system that is able to access a service on another computer, such as a server facility 142, via a network. This client facility server facility 142 model may apply to a plurality of networked applications, such as a client facility connecting to an enterprise facility 102 application server facility 142, a web browser client facility connecting to a web server facility 142, an e-mail client facility retrieving e-mail from an Internet 154 service provider's mail storage servers 142, and the like. In embodiments, traditional large client facility applications may be switched to websites, which may increase the browser's role as a client facility. Clients 144 may be classified as a function of the extent to which they perform their own processing. For instance, client facilities are sometimes classified as a fat client facility or thin client facility. The fat client facility, also known as a thick client facility or rich client facility, may be a client facility that performs the bulk of data processing operations itself, and does not necessarily rely on the server facility 142. The fat client facility may be most common in the form of a personal computer, where the personal computer may operate independent of any server facility 142. Programming environments for fat clients 144 may include CURI, Delphi, Droplets, Java, win32, X11, and the like. Thin clients 144 may offer minimal processing capabilities, for instance, the thin client facility may primarily provide a graphical user interface provided by an application server facility 142, which may perform the bulk of any required data processing. Programming environments for thin clients 144 may include JavaScript/AJAX, ASP, JSP, Ruby on Rails, Python's Django, PHP, and the like. The client facility may also be a mix of the two, such as processing data locally, but relying on a server facility 142 for data storage. As a result, this hybrid client facility may provide benefits from both the fat client facility type, such as multimedia support and high performance, and the thin client facility type, such as high

manageability and flexibility. In embodiments, the threat management facility 100, and associated endpoint computer security facility 152, may provide seamless threat protection to the plurality of clients 144, and client facility types, across the enterprise facility 102.

**[0072]** The enterprise facility 102 may include a plurality of server facilities 142, such as application servers, communications servers, file servers, database servers, proxy servers, mail servers, fax servers, game servers, web servers, and the like. A server facility 142, which may also be referred to as a server facility 142 application, server facility 142 operating system, server facility 142 computer, or the like, may be an application program or operating system that accepts client facility connections in order to service requests from clients 144. The server facility 142 application may run on the same computer as the client facility using it, or the server facility 142 and the client facility may be running on different computers and communicating across the network. Server facility 142 applications may be divided among server facility 142 computers, with the dividing depending upon the workload. For instance, under light load conditions all server facility 142 applications may run on a single computer and under heavy load conditions a single server facility 142 application may run on multiple computers. In embodiments, the threat management facility 100 may provide threat protection to server facilities 142 within the enterprise facility 102 as load conditions and application changes are made.

**[0073]** A server facility 142 may also be an appliance facility 140, where the appliance facility 140 provides specific services onto the network. Though the appliance facility 140 is a server facility 142 computer, that may be loaded with a server facility 142 operating system and server facility 142 application, the enterprise facility 102 user may not need to configure it, as the configuration may have been performed by a third party. In an embodiment, an enterprise facility 102 appliance may be a server facility 142 appliance that has been configured and adapted for use with the threat management facility 100, and located within the facilities of the enterprise facility 102. The enterprise facility's 102 threat management appliance may enable the enterprise facility 102 to administer an on-site local managed threat protection configuration, where the administration facility 134 may access the threat resources through an interface, such as a web portal. In an alternate embodiment, the enterprise facility 102 may be managed remotely from a third party, vendor, or the like, without an appliance facility 140 located within the enterprise facility 102. In this instance, the appliance functionality may be a shared hardware product between pluralities of enterprises 102. In embodiments, the appliance facility 140 may be located

at the enterprise facility 102, where the enterprise facility 102 maintains a degree of control. In embodiments, a hosted service may be provided, where the appliance 140 may still be an on-site black box to the enterprise facility 102, physically placed there because of infrastructure requirements, but managed by a third party, vendor, or the like.

**[0074]** Simple server facility 142 appliances may also be utilized across the enterprise facility's 102 network infrastructure, such as switches, routers, wireless routers, hubs and routers, gateways, print servers, net modems, and the like. These simple server facility appliances may not require configuration by the enterprise facility 102, but may require protection from threats via an endpoint computer security facility 152. These appliances may provide interconnection services within the enterprise facility 102 network, and therefore may advance the spread of a threat if not properly protected.

**[0075]** A client facility may be protected from threats from within the enterprise facility 102 network using a personal firewall, which may be a hardware firewall, software firewall, or combination of these, that controls network traffic to and from a client. The personal firewall may permit or deny communications based on a security policy. Personal firewalls may be designed for use by end-users, which may result in protection for only the computer on which it's installed. Personal firewalls may be able to control network traffic by providing prompts each time a connection is attempted and adapting security policy accordingly. Personal firewalls may also provide some level of intrusion detection, which may allow the software to terminate or block connectivity where it suspects an intrusion is being attempted. Other features that may be provided by a personal firewall may include alerts about outgoing connection attempts, control of program access to networks, hiding the client from port scans by not responding to unsolicited network traffic, monitoring of applications that may be listening for incoming connections, monitoring and regulation of incoming and outgoing network traffic, prevention of unwanted network traffic from installed applications, reporting applications that make connection attempts, reporting destination servers with which applications may be attempting communications, and the like. In embodiments, the personal firewall may be provided by the threat management facility 100.

**[0076]** Another important component that may be protected by an endpoint computer security facility 152 is a network firewall facility 138, which may be a hardware or software device that may be configured to permit, deny, or proxy data through a computer network that has different levels of trust in its source of data. For instance, an internal enterprise facility 102

network may have a high level of trust, because the source of all data has been sourced from within the enterprise facility 102. An example of a low level of trust is the Internet 154, because the source of data may be unknown. A zone with an intermediate trust level, situated between the Internet 154 and a trusted internal network, may be referred to as a “perimeter network.” Since firewall facilities 138 represent boundaries between threat levels, the endpoint computer security facility 152 associated with the firewall facility 138 may provide resources that may control the flow of threats at this enterprise facility 102 network entry point. Firewall facilities 138, and associated endpoint computer security facility 152, may also be associated with a network node that may be equipped for interfacing between networks that use different protocols. In embodiments, the endpoint computer security facility 152 may provide threat protection in a plurality of network infrastructure locations, such as at the enterprise facility 102 network entry point, i.e. the firewall facility 138 or gateway; at the server facility 142; at distribution points within the network, i.e. the hubs and routers 148; at the desktop of client facility computers; and the like. In embodiments, the most effective location for threat detection may be at the user’s computer desktop endpoint computer security facility 152.

[0077] The interface between the threat management facility 100 and the enterprise facility 102, and through the appliance facility 140 to embedded endpoint computer security facilities, may include a set of tools that may be the same for all enterprise implementations, but allow each enterprise to implement different controls. In embodiments, these controls may include both automatic actions and managed actions. Automatic actions may include downloads of the endpoint computer security facility 152 to components of the enterprise facility 102, downloads of updates to existing endpoint computer security facilities of the enterprise facility 102, uploaded network interaction requests from enterprise facility 102 components to the threat management facility 100, and the like. In embodiments, automatic interactions between the enterprise facility 102 and the threat management facility 100 may be configured by the threat management facility 100 and an administration facility 134 in the enterprise facility 102. The administration facility 134 may configure policy rules that determine interactions, such as developing rules for accessing applications, as in who is authorized and when applications may be used; establishing rules for ethical behavior and activities; rules governing the use of entertainment software such as games, or personal use software such as IM and VoIP; rules for determining access to enterprise facility 102 computing resources, including authentication, levels of access, risk assessment, and usage history tracking; rules for when an action is not

allowed, such as whether an action is completely designed or just modified in its execution; and the like. The administration facility 134 may also establish license management, which in turn may further determine interactions associated with a licensed application. In embodiments, interactions between the threat management facility 100 and the enterprise facility 102 may provide threat protection to the enterprise facility 102 by managing the flow of network data into and out of the enterprise facility 102 through automatic actions that may be configured by the threat management facility 100 or the administration facility 134.

**[0078]** Client facilities within the enterprise facility 102 may be connected to the enterprise facility 102 network by way of wired network facilities 148A or wireless network facilities 148B. Client facilities connected to the enterprise facility 102 network via a wired facility 148A or wireless facility 148B may receive similar protection, as both connection types are ultimately connected to the same enterprise facility 102 network, with the same endpoint computer security facility 152, and the same threat protected enterprise facility 102 environment. Mobile wireless facility clients 144B–F, because of their ability to connect to any wireless 148B,D network access point, may connect to the Internet 154 outside the enterprise facility 102, and therefore outside the threat-protected environment of the enterprise facility 102. In this instance the mobile client facility (e.g., the clients 144 B–F), if not for the presence of the endpoint computer security facility 152 may experience a malware attack or perform actions counter to enterprise facility 102 established policies. In addition, there may be a plurality of ways for the threat management facility 100 to protect the out-of-enterprise facility 102 mobile client facility (e.g., the clients 144 D–F) that has an embedded endpoint computer security facility 152, such as by providing URI filtering in personal routers, using a web appliance as a DNS proxy, or the like. Mobile client facilities that are components of the enterprise facility 102 but temporarily outside connectivity with the enterprise facility 102 network may be provided with the same threat protection and policy control as client facilities inside the enterprise facility 102. In addition, mobile the client facilities may receive the same interactions to and from the threat management facility 100 as client facilities inside the enterprise facility 102, where the mobile client facilities may be considered a virtual extension of the enterprise facility 102, receiving all the same services via their embedded endpoint computer security facility 152.

**[0079]** Interactions between the threat management facility 100 and the components of the enterprise facility 102, including mobile client facility extensions of the enterprise facility 102, may ultimately be connected through the Internet 154. Threat management facility 100

downloads and upgrades to the enterprise facility 102 may be passed from the firewalled networks of the threat management facility 100 through to the endpoint computer security facility 152 equipped components of the enterprise facility 102. In turn the endpoint computer security facility 152 components of the enterprise facility 102 may upload policy and access requests back across the Internet 154 and through to the threat management facility 100. The Internet 154 however, is also the path through which threats may be transmitted from their source. These network threats 104 may include threats from a plurality of sources, including without limitation, websites, e-mail, IM, VoIP, application software, and the like. These threats may attempt to attack a mobile enterprise client facility (e.g., the clients 144B–F) equipped with an endpoint computer security facility 152, but in embodiments, as long as the mobile client facility is embedded with an endpoint computer security facility 152, as described above, threats may have no better success than if the mobile client facility were inside the enterprise facility 102.

**[0080]** However, if the mobile client facility were to attempt to connect into an unprotected connection point, such as at a secondary location 108 that is not a part of the enterprise facility 102, the mobile client facility may be required to request network interactions through the threat management facility 100, where contacting the threat management facility 100 may be performed prior to any other network action. In embodiments, the client facility's 144 endpoint computer security facility 152 may manage actions in unprotected network environments such as when the client facility (e.g., client 144F) is in a secondary location 108 or connecting wirelessly to a non-enterprise facility 102 wireless Internet connection, where the endpoint computer security facility 152 may dictate what actions are allowed, blocked, modified, or the like. For instance, if the client facility's 144 endpoint computer security facility 152 is unable to establish a secured connection to the threat management facility 100, the endpoint computer security facility 152 may inform the user of such, and recommend that the connection not be made. In the instance when the user chooses to connect despite the recommendation, the endpoint computer security facility 152 may perform specific actions during or after the unprotected connection is made, including running scans during the connection period, running scans after the connection is terminated, storing interactions for subsequent threat and policy evaluation, contacting the threat management facility 100 upon first instance of a secured connection for further actions and or scanning, restricting access to network and local resources, or the like. In embodiments, the endpoint computer security facility 152 may perform specific

actions to remediate possible threat incursions or policy violations during or after the unprotected connection.

**[0081]** The secondary location 108 may have no endpoint computer security facilities 152 as a part of its computer components, such as its firewalls 138B, servers 142B, clients 144G, hubs and routers 148C–D, and the like. As a result, the computer components of the secondary location 108 may be open to threat attacks, and become potential sources of threats, as well as any mobile enterprise facility clients 144B–F that may be connected to the secondary location's 108 network. In this instance, these computer components may now unknowingly spread a threat to other components connected to the network.

**[0082]** Some threats may not come directly from the Internet 154, such as from non-enterprise facility controlled mobile devices that are physically brought into the enterprise facility 102 and connected to the enterprise facility 102 client facilities. The connection may be made from direct connection with the enterprise facility's 102 client facility, such as through a USB port, or in physical proximity with the enterprise facility's 102 client facility such that a wireless facility connection can be established, such as through a Bluetooth connection. These physical proximity threats 110 may be another mobile computing device, a portable memory storage device, a mobile communications device, or the like, such as CDs and DVDs, memory sticks, flash drives, external hard drives, cell phones, PDAs, MP3 players, digital cameras, point-to-point devices, digital picture frames, digital pens, navigation devices, tablets, appliances, and the like. A physical proximity threat 110 may have been previously infiltrated by network threats while connected to an unprotected network connection outside the enterprise facility 102, and when connected to the enterprise facility 102 client facility, pose a threat. Because of their mobile nature, physical proximity threats 110 may infiltrate computing resources in any location, such as being physically brought into the enterprise facility 102 site, connected to an enterprise facility 102 client facility while that client facility is mobile, plugged into an unprotected client facility at a secondary location 108, and the like. A mobile device, once connected to an unprotected computer resource, may become a physical proximity threat 110. In embodiments, the endpoint computer security facility 152 may provide enterprise facility 102 computing resources with threat protection against physical proximity threats 110, for instance, through scanning the device prior to allowing data transfers, through security validation certificates, through establishing a safe zone within the enterprise facility 102 computing resource to transfer data into for evaluation, and the like.

**[0083]** Having provided an overall context for threat detection, the description now turns to a brief discussion of an example of a computer system that may be used for any of the entities and facilities described above.

**[0084]** Fig. 2 illustrates a computer system. In general, the computer system 200 may include a computing device 210 connected to a network 202, e.g., through an external device 204. The computing device 210 may be or include any type of network endpoint or endpoints as described herein, e.g., with reference to Fig. 1 above. For example, the computing device 210 may include a desktop computer workstation. The computing device 210 may also or instead be any suitable device that has processes and communicates over a network 202, including without limitation a laptop computer, a desktop computer, a personal digital assistant, a tablet, a mobile phone, a television, a set top box, a wearable computer (e.g., watch, jewelry, or clothing), a home device (e.g., a thermostat or a home appliance controller), just as some examples. The computing device 210 may also or instead include a server, or it may be disposed on a server.

**[0085]** The computing device 210 may be used for any of the entities described in the threat management environment described above with reference to Fig. 1. For example, the computing device 210 may be a server, a client an enterprise facility, a threat management facility, or any of the other facilities or computing devices described therein. In certain aspects, the computing device 210 may be implemented using hardware (e.g., a desktop computer), software (e.g., in a virtual machine or the like), or a combination of software and hardware (e.g., with programs executing on the desktop computer), and the computing device 210 may be a standalone device, a device integrated into another entity or device, a platform distributed across multiple entities, or a virtualized device executing in a virtualization environment.

**[0086]** The network 202 may include any network described above, e.g., data network(s) or internetwork(s) suitable for communicating data and control information among participants in the computer system 200. This may include public networks such as the Internet, private networks, and telecommunications networks such as the Public Switched Telephone Network or cellular networks using third generation cellular technology (e.g., 3G or IMT-2000), fourth generation cellular technology (e.g., 4G, LTE, MT-Advanced, E-UTRA, etc.) or WiMax-Advanced (IEEE 802.16m)) and/or other technologies, as well as any of a variety of corporate area, metropolitan area, campus or other local area networks or enterprise networks, along with any switches, routers, hubs, gateways, and the like that might be used to carry data among

participants in the computer system 200. The network 202 may also include a combination of data networks, and need not be limited to a strictly public or private network.

**[0087]** The external device 204 may be any computer or other remote resource that connects to the computing device 210 through the network 202. This may include threat management resources such as any of those contemplated above, gateways or other network devices, remote servers or the like containing content requested by the computing device 210, a network storage device or resource, a device hosting malicious content, or any other resource or device that might connect to the computing device 210 through the network 202.

**[0088]** The computing device 210 may include a processor 212, a memory 214, a network interface 216, a data store 218, and one or more input/output devices 220. The computing device 210 may further include or be in communication with peripherals 222 and other external input/output devices 224.

**[0089]** The processor 212 may be any as described herein, and in general be capable of processing instructions for execution within the computing device 210 or computer system 200. The processor 212 may include a single-threaded processor or a multi-threaded processor. The processor 212 may be capable of processing instructions stored in the memory 214 or on the data store 218.

**[0090]** The memory 214 may store information within the computing device 210 or computer system 200. The memory 214 may include any volatile or non-volatile memory or other computer-readable medium, including without limitation a Random Access Memory (RAM), a flash memory, a Read Only Memory (ROM), a Programmable Read-only Memory (PROM), an Erasable PROM (EPROM), registers, and so forth. The memory 214 may store program instructions, program data, executables, and other software and data useful for controlling operation of the computing device 200 and configuring the computing device 200 to perform functions for a user. The memory 214 may include a number of different stages and types for different aspects of operation of the computing device 210. For example, a processor may include on-board memory and/or cache for faster access to certain data or instructions, and a separate, main memory or the like may be included to expand memory capacity as desired.

**[0091]** The memory 214 may, in general, include a non-volatile computer readable medium containing computer code that, when executed by the computing device 200 creates an execution environment for a computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a

combination of the foregoing, and/or code that performs some or all of the steps set forth in the various flow charts and other algorithmic descriptions set forth herein. While a single memory 214 is depicted, it will be understood that any number of memories may be usefully incorporated into the computing device 210. For example, a first memory may provide non-volatile storage such as a disk drive for permanent or long-term storage of files and code even when the computing device 210 is powered down. A second memory such as a random access memory may provide volatile (but higher speed) memory for storing instructions and data for executing processes. A third memory may be used to improve performance by providing even higher speed memory physically adjacent to the processor 212 for registers, caching and so forth.

**[0092]** The network interface 216 may include any hardware and/or software for connecting the computing device 210 in a communicating relationship with other resources through the network 202. This may include remote resources accessible through the Internet, as well as local resources available using short range communications protocols using, e.g., physical connections (e.g., Ethernet), radio frequency communications (e.g., WiFi), optical communications, (e.g., fiber optics, infrared, or the like), ultrasonic communications, or any combination of these or other media that might be used to carry data between the computing device 210 and other devices. The network interface 216 may, for example, include a router, a modem, a network card, an infrared transceiver, a radio frequency (RF) transceiver, a near field communications interface, a radio-frequency identification (RFID) tag reader, or any other data reading or writing resource or the like.

**[0093]** More generally, the network interface 216 may include any combination of hardware and software suitable for coupling the components of the computing device 210 to other computing or communications resources. By way of example and not limitation, this may include electronics for a wired or wireless Ethernet connection operating according to the IEEE 802.11 standard (or any variation thereof), or any other short or long range wireless networking components or the like. This may include hardware for short range data communications such as Bluetooth or an infrared transceiver, which may be used to couple to other local devices, or to connect to a local area network or the like that is in turn coupled to a data network 202 such as the Internet. This may also or instead include hardware/software for a WiMax connection or a cellular network connection (using, e.g., CDMA, GSM, LTE, or any other suitable protocol or combination of protocols). The network interface 216 may be included as part of the input/output devices 220 or vice-versa.

**[0094]** The data store 218 may be any internal memory store providing a computer-readable medium such as a disk drive, an optical drive, a magnetic drive, a flash drive, or other device capable of providing mass storage for the computing device 210. The data store 218 may store computer readable instructions, data structures, program modules, and other data for the computing device 210 or computer system 200 in a non-volatile form for subsequent retrieval and use. For example, the data store 218 may store without limitation one or more of the operating system, application programs, program data, databases, files, and other program modules or other software objects and the like.

**[0095]** The input/output interface 220 may support input from and output to other devices that might couple to the computing device 210. This may, for example, include serial ports (e.g., RS-232 ports), universal serial bus (USB) ports, optical ports, Ethernet ports, telephone ports, audio jacks, component audio/video inputs, HDMI ports, and so forth, any of which might be used to form wired connections to other local devices. This may also or instead include an infrared interface, RF interface, magnetic card reader, or other input/output system for coupling in a communicating relationship with other local devices. It will be understood that, while the network interface 216 for network communications is described separately from the input/output interface 220 for local device communications, these two interfaces may be the same, or may share functionality, such as where a USB port is used to attach to a WiFi accessory, or where an Ethernet connection is used to couple to a local network attached storage.

**[0096]** A peripheral 222 may include any device used to provide information to or receive information from the computing device 200. This may include human input/output (I/O) devices such as a keyboard, a mouse, a mouse pad, a track ball, a joystick, a microphone, a foot pedal, a camera, a touch screen, a scanner, or other device that might be employed by the user 230 to provide input to the computing device 210. This may also or instead include a display, a speaker, a printer, a projector, a headset or any other audiovisual device for presenting information to a user. The peripheral 222 may also or instead include a digital signal processing device, an actuator, or other device to support control or communication to other devices or components. Other I/O devices suitable for use as a peripheral 222 include haptic devices, three-dimensional rendering systems, augmented-reality displays, magnetic card readers, and so forth. In one aspect, the peripheral 222 may serve as the network interface 216, such as with a USB device configured to provide communications via short range (e.g., BlueTooth, WiFi, Infrared, RF, or the like) or long range (e.g., cellular data or WiMax) communications protocols. In

another aspect, the peripheral 222 may provide a device to augment operation of the computing device 210, such as a global positioning system (GPS) device, a security dongle, or the like. In another aspect, the peripheral may be a storage device such as a flash card, USB drive, or other solid state device, or an optical drive, a magnetic drive, a disk drive, or other device or combination of devices suitable for bulk storage. More generally, any device or combination of devices suitable for use with the computing device 200 may be used as a peripheral 222 as contemplated herein.

**[0097]** Other hardware 226 may be incorporated into the computing device 200 such as a co-processor, a digital signal processing system, a math co-processor, a graphics engine, a video driver, and so forth. The other hardware 226 may also or instead include expanded input/output ports, extra memory, additional drives (e.g., a DVD drive or other accessory), and so forth.

**[0098]** A bus 232 or combination of busses may serve as an electromechanical platform for interconnecting components of the computing device 200 such as the processor 212, memory 214, network interface 216, other hardware 226, data store 218, and input/output interface. As shown in the figure, each of the components of the computing device 210 may be interconnected using a system bus 232 or other communication mechanism for communicating information.

**[0099]** Methods and systems described herein can be realized using the processor 212 of the computer system 200 to execute one or more sequences of instructions contained in the memory 214 to perform predetermined tasks. In embodiments, the computing device 200 may be deployed as a number of parallel processors synchronized to execute code together for improved performance, or the computing device 200 may be realized in a virtualized environment where software on a hypervisor or other virtualization management facility emulates components of the computing device 200 as appropriate to reproduce some or all of the functions of a hardware instantiation of the computing device 200.

**[00100]** Fig. 3 illustrates a threat management system as contemplated herein. In general, the system may include an endpoint 302, a firewall 304, a server 306 and a threat management facility 308 coupled to one another directly or indirectly through a data network 305, all as generally described above. Each of the entities depicted in Fig. 3 may, for example, be implemented on one or more computing devices such as the computing device described above with reference to Fig.2. A number of systems may be distributed across these various components to support threat detection, such as a coloring system 310, a key management system 312 and a heartbeat system 314, each of which may include software components

executing on any of the foregoing system components, and each of which may communicate with the threat management facility 308 and an endpoint threat detection agent 320 executing on the endpoint 302 to support improved threat detection and remediation.

**[00101]** The coloring system 310 may be used to label or ‘color’ software objects for improved tracking and detection of potentially harmful activity. The coloring system 310 may, for example, label files, executables, processes, network communications, data sources and so forth with any suitable. A variety of techniques may be used to select static and/or dynamic labels for any of these various software objects, and to manage the mechanics of applying and propagating coloring information as appropriate. For example, a process may inherit a color from an application that launches the process. Similarly a file may inherit a color from a process when it is created or opened by a process, and/or a process may inherit a color from a file that the process has opened. More generally, any type of labeling, as well as rules for propagating, inheriting, changing, or otherwise manipulating such labels, may be used by the coloring system 310 as contemplated herein. A suitable coloring system is described in greater detail below with reference to Fig. 4.

**[00102]** The key management system 312 may support management of keys for the endpoint 302 in order to selectively permit or prevent access to content on the endpoint 302 on a file-specific basis, a process-specific basis, an application-specific basis, a user-specific basis, or any other suitable basis in order to prevent data leakage, and in order to support more fine-grained and immediate control over access to content on the endpoint 302 when a security compromise is detected. Thus for example, if a particular process executing on the endpoint is compromised, or potentially compromised or otherwise under suspicion, keys to that process may be revoked in order to prevent, e.g., data leakage or other malicious activity. A suitable key management system useful in this context is described in greater detail below with reference to Fig. 5.

**[00103]** The heartbeat system 314 may be used to provide periodic or aperiodic information from the endpoint 302 or other system components about system health, security, status, and so forth. A heartbeat may be encrypted or plaintext, or some combination of these, and may be communicated unidirectionally (e.g., from the endpoint 308 to the threat management facility 308) or bidirectionally (e.g., between the endpoint 302 and the server 306, or any other pair of system components) on any useful schedule. A suitable heartbeat system is described in greater detail below with reference to Fig. 6.

**[00104]** In general, these various monitoring and management systems may cooperate to provide improved threat detection and response. For example, the coloring system 310 may be used to evaluate when a particular process is potentially opening inappropriate files, and a potential threat may be confirmed based on an interrupted heartbeat from the heartbeat system 314. The key management system 312 may then be deployed to revoke keys to the process so that no further files can be opened, deleted or otherwise modified. More generally, the cooperation of these systems enables a wide variety of reactive measures that can improve detection and remediation of potential threats to an endpoint.

**[00105]** Fig. 4 illustrates a system for behavioral tracking, coloring, and generation of indications of compromise (IOCs). In general, the system 400 may include a number of entities participating in a threat management process such as any of the entities and threat management processes described herein. The threat management process may for example employ techniques such as behavioral tracking, encryption, endpoint recording, reputation-based threat detection, behavioral-based threat detection, signature-based threat detection, and combinations of the foregoing, or any other suitable techniques for detecting threats to endpoints in a network.

**[00106]** In general, the system 400 may include a number of endpoints 402, 412 and a threat management facility 404 in an enterprise 410, such as any of the enterprises described above. An external analysis facility 406 may analyze threat data and provide rules and the like for use by the threat management facility 404 and endpoints 402, 412 in managing threats to the enterprise 410. The threat management facility 404 may reside in a local appliance (e.g., embedded within, or locally coupled to the endpoint 402), a virtual appliance (e.g., which could be run by a protected set of systems on their own network system(s)), a private cloud, a public cloud, and so forth. The analysis facility 406 may store locally-derived threat information. The analysis facility 406 may also or instead receive threat information from a third party source 416 such as MITRE Corporation or any other public, private, educational or other organization that gathers information on network threats and provides analysis and threat detection information for use by others. Each of these components may be configured with suitable programming to participate in the various threat detection and management techniques contemplated herein. The threat management facility 404 may monitor any stream of data from an endpoint 402 exclusively, or use the full context of intelligence from the stream of all protected endpoints 402, 412 or some combination of these.

**[00107]** The endpoint 402 may be any of the endpoints described herein, or any other device or network asset that might join or participate in the enterprise 410 or otherwise operate on an enterprise network. This may, for example, include a server, a client such as a desktop computer or a mobile computing device (e.g., a laptop computer, a wearable device, a tablet, and the like), a cellular phone, a smart phone, or other computing device suitable for participating in the enterprise 410.

**[00108]** In general, the endpoint 402 may include any number of computing objects such as an object 418 labeled with a descriptor 420. While the term object has a number of specific meanings in the art, and in particular in object-oriented programming, it will be understood that the term 'object' as used herein is intended to be significantly broader, and may include any data, process, file or combination of these including without limitation any process, application, executable, script, dynamic linked library, file, data, database, data source, data structure, function, resource locator (e.g., uniform resource locator (URL) or other uniform resource identifier (URI)), or the like that might be manipulated by one of the computing devices described herein.

**[00109]** An object 418 may also or instead include a remote resource, such as a resource identified in a URL. That is, while the objects 418 in Fig. 4 are depicted as residing on the endpoint 402, an object 418 may also reside elsewhere in the system 400, while still being labeled with a descriptor 420 and tracked by the monitor 421 of the endpoint 402. The object 418 may be an item that is performing an action or causing an event, or the object 418 may be an item that is receiving the action or result of an event (i.e., the item in the system 400 being acted upon).

**[00110]** Where the object 418 is data or includes data, the object 418 may be encrypted or otherwise protected, or the object 418 may be unencrypted or otherwise unprotected. The object 418 may be a process or other computing object that performs an action, which may include a single event or a collection or sequence of events taken by a process. The object 418 may also or instead include an item such as a file or lines of code that are executable to perform such actions. The object 418 may also or instead include a computing component upon which an action is taken, e.g., a system setting (e.g., a registry key or the like), a data file, a URL, or the like. The object 418 may exhibit a behavior such as an interaction with another object or component of the system 400.

**[00111]** In one aspect, objects 418 may be described in terms of persistence. The object 418 may, for example, be a part of a process, and remain persistent as long as that process is alive. The object 418 may instead be persistent across an endpoint 402 and remain persistent as long as an endpoint 402 is active or alive. The object 418 may instead be a global object having persistence outside of an endpoint 418, such as a URL or a data store. In other words, the object 418 may be a persistent object with persistence outside of the endpoint.

**[00112]** Although many if not most objects 418 will typically be benign objects that may be found on a normal, operating endpoint, an object 418 may contain software associated with an advanced persistent threat (APT) or other malware that resides partially or entirely on the endpoint 402. The associated software may have reached the endpoint 402 in a variety of ways, and may have been placed manually or automatically on the endpoint 402 by a malicious source. It will be understood that the associated software may take any number of forms and have any number of components. For example, the associated software may include an executable file that can execute independently, or the associated software may be a macro, plug-in, or the like that executes within another application. Similarly, the associated software may manifest as one or more processes or threads executing on the endpoint 402. Further, the associated software may install from a file on the endpoint 402 (or a file remote from the endpoint 402), and the associated software may create one or more files such as data files or the like while executing. Associated software should be understood to generally include all such files and processes except where a specific file or process is more specifically noted.

**[00113]** A threat such as an APT may also take the form of an attack where no altered or additional software is directly added or modified on the endpoint 402. Instead, an adversary may reuse existing software on the system 400 to perform the attacks. It is for this reason that simply scanning for associated software may be insufficient for the detection of APTs and it may be preferable to detect APTs based on the behavior of the software and associated objects 418 that are used by, for, and with that software.

**[00114]** An object coloring system 414 may apply descriptors 420 to objects 418 on the endpoint 402. This may be performed continuously by a background process on the endpoint 402, or it may occur whenever an object 418 is involved in an action, such as when a process makes a call to an application programming interface (API) or takes some other action, or when a URL is used to initiate a network request, or when a read or a write is performed on data in a file. This may also or instead include a combination of these approaches as well as other

approaches, such as by labeling a file or application when it is moved to the endpoint 402, or when the endpoint 402 is started up or instantiated. In general, the object coloring system 414 may add, remove or change a color at any location and at any moment that can be practicably instrumented on a computer system.

**[00115]** As noted above, the term ‘object’ as used herein is intended to include a wide range of computing objects and as such, the manner in which particular objects 418 are labeled or ‘colored’ with descriptors 420 may vary significantly. Any object 418 that is performing an action may be colored at the time of and/or with a label corresponding to the action, or likewise any object 418 that is the target of the action may be colored at the time that it is used and/or with a label corresponding to a process or the like using the object 418. Furthermore, the operating system runtime representation of the object 418 may be colored, or the persistent object outside of the operating system may be colored (as is the case for a File Handle or File Object within the operating system or the actual file as stored in a file system), such as within an encryption header or other header applied to the file, or as part of a directory attribute or any other persistent location within the file or file system. A former coloring may be ephemerally tracked while the operating system maintains the representation and the latter may persist long after any reboots of the same operating system and likewise have meaning when read or used by other endpoints 402. For processes, each file handle may be supplemented with a pointer or other mechanism for locating a descriptor 420 for a particular object 420 that is a process. More specifically, each object 418 may be colored in any manner suitable for appending information to that object 418 so that the corresponding descriptor 420 can be retrieved and, where appropriate, updated.

**[00116]** The coloring system 414 may apply any suitable rules for adding and changing descriptors 420 for objects 418. For example, when a process with a certain descriptor accesses data with a different descriptor, the descriptor for the process may be updated to correspond to the data, or the descriptor for the data may be updated to correspond to the process, or some combination of these. Any action by or upon an object 418 may trigger a coloring rule so that descriptors 420 can be revised at any relevant time(s) during processing.

**[00117]** In one aspect, colors will not explicitly indicate a compromised security state or other good/bad types of distinctions (although they may be adapted to this use). Instead, colors may record some known information or understanding about an object 418, such as a source, a purpose, and so forth. In this context, colors will not be used to label actual or potential security

compromises, but to identify inconsistencies among interacting objects 418, and to restrict or control access and use accordingly. For example, where an endpoint uses file-system-based encryption as described herein, a process that is colored as exposed to external resources (e.g., the Internet) may be prohibited from accessing cleartext data for protected files. Colors can also be used in other contexts such as intrusion prevention, routing rules, and detection of odd or questionable behavior.

**[00118]** In one aspect, colors may be implemented as flags associated with objects 418 that provide a short hand cache of potentially relevant information. While this information could also be obtained for an object 418 through a careful inspection of related activity logs or other data recording activities, the use of a cache of flags for coloring information makes the coloring information directly available and immediately actionable, as distinguished from post hoc forensic activities that are otherwise supported by data logging.

**[00119]** In one aspect, colors as contemplated herein may fall into two different categories: static colors and dynamic colors. Static colors may be explicitly applied based on, e.g., a controlling application. For example, a static color may specify a status of an application or data, or an associated type of application (e.g., productivity, mail client, messaging, browser, word processing, financial, spreadsheet, etc.). In this context, a process will generally inherit static colors from a source executable, and will permit inferences for appropriate behavior and related processes. Dynamic colors may be assigned based on direct observation of executing processes, and may not be inherited or transferred among processes (although the presence of a dynamic color may be used to draw another coloring inference upon interaction with another process). Thus, the inheritance of colors may depend in part upon the type of color that is applied, or upon explicit inheritance rules provided for a particular color.

**[00120]** A descriptor 420 may take a variety of forms, and may in general include any information selected for relevance to threat detection. This may, for example, be a simple categorization of data or processes such as trusted or untrusted. For example, in one embodiment described herein, data and processes are labeled as either 'IN' (e.g., trusted) or 'OUT' (e.g., untrusted). The specific content of the label is unimportant, and this may be a binary flag, text string, encrypted data or other human-readable and/or machine-readable identifier, provided that the descriptor 420 can facilitate discrimination among labeled files – in this example, between trusted objects 418 and untrusted objects 418 so that, e.g., trusted data can be selectively decrypted or encrypted for use with trusted processes. Similarly, data may be labeled as

corporate data or private data, with similar type-dependent processing provided. For example, private data may be encrypted with a key exclusively controlled by the data owner, while corporate data may be encrypted using a remotely managed key ring for an enterprise operated by the corporation.

**[00121]** In another aspect, the descriptor 420 may provide a multi-tiered or hierarchical description of the object 418 including any information useful for characterizing the object 418 in a threat management context. For example, in one useful configuration the descriptor 420 may include a type or category, static threat detection attributes, and an explicit identification. The type or category for the object 418 may be any category or the like that characterizes a general nature or use of the object 418 as inferred from behavior and other characteristics. This may, for example, include categories such as ‘game,’ ‘financial,’ ‘application,’ ‘electronic mail,’ ‘image,’ ‘video,’ ‘browser,’ ‘antivirus,’ and so forth. The category may be more granular, or may include hierarchical categories such as ‘application:spreadsheet,’ ‘application:word\_processing,’ and so forth. Such colors may be directly inferred from a single action, a sequence of actions, or a combination of actions and other colors, including, e.g., colors of processes and files related to a particular action, or other objects 418 that provide context for a particular action or group of actions. One or more colors may also or instead be explicitly provided by a user or a process, or otherwise automatically or manually attributed to computer objects as contemplated herein.

**[00122]** The static threat detection attributes may be any readily ascertainable characteristics of the object 418 useful in threat detection. This may, for example, include an antivirus signature, a hash, a file size, file privileges, a process user, a path or director, and so forth. Static threat detection attributes may also include attributes that are derived by or supplied from other sources. For example, static threat detection attributes may include a reputation for an object 418, which may be expressed in any suitable or useful level of granularity such as with discrete categories (trusted/untrusted/unknown) or with a numerical score or other quantitative indicator. The explicit identification may, in general, be what an object 418 calls itself, e.g., a file name or process name.

**[00123]** Some actions may transfer colors from the subject of the action to the target of the action. For example, when a process creates sub-processes, the sub-processes may inherit the colors of its parent(s). By way of another example, when a process is initially loaded from an executable, it may inherit the color(s) stored in the file system for or with the executable.

**[00124]** In general, the descriptor 420 may be provided in any suitable format. The descriptor 420 may for example be formed as a vector of binary flags or other attributes that form the ‘color’ or description of an object 418. The descriptor 420 may also, where appropriate, include scalar quantities for certain properties. For example, it may be relevant how many times a system file was accessed, how many file handles a process has open, how many times a remote resource was requested or how long a remote resource is connected, and this information may be suitably encoded in the descriptor 420 for use in coloring objects with the coloring system 414 and applying rules for IOC detection by the IOC monitor 421.

**[00125]** An indication of compromise (IOC) monitor 421 may be provided to instrument the endpoint 402 so that any observable actions by or involving various objects 418 can be detected. As with the coloring system 414, it will be understood that the types of observable actions will vary significantly, and the manner in which the endpoint 402 is instrumented to detect such actions will depend on the particular type of object 418. For example, for files or the like, an API for a file system may be used to detect reads, writes, and other access (e.g., open, read, write, move, copy, delete, etc.), and may be configured to report to or otherwise initiate monitoring of the action taken with the file through the file system. As another example, kernel objects may be instrumented at the corresponding object handle or in some other manner. As a further example, a kernel driver may be used for intercepting a process startup. While a wide variety of objects are contemplated herein, one of ordinary skill in the art may create suitable instrumentation for any computing object so that it may be monitored by the IOC monitor 421.

**[00126]** It will be noted that suitable instrumentation may be created for a variety of functions and circumstances. For example, instrumentation may usefully track requests for network access or other actions back to a particular application or process, or data payloads back to a particular file or data location. One of ordinary skill in the art can readily implement suitable traces and/or logging for any such information that might be useful in a particular IOC monitoring operation.

**[00127]** In general, the IOC monitor 421 applies rules to determine when there is an IOC 422 suitable for reporting to a threat management facility 404. It will be understood that an endpoint 402 may, in suitable circumstances and with appropriate information, take immediate local action to remediate a threat. However, the monitor 421 may advantageously accumulate a sequence of actions, and still more advantageously may identify inconsistencies or unexpected behavior within a group of actions with improved sensitivity by comparing descriptors 420 for

various objects 418 involved in relevant actions and events. In this manner, rules may be applied based upon the descriptors 420 that better discriminate malicious activity while reducing the quantity and frequency of information that must be communicated to a remote threat management facility 404. At the same time, all of the relevant information provided by the descriptors 420 can be sent in an IOC 422 when communicating a potential issue to the threat management facility 404. For example, during the course of execution, a specific process (as evidenced by its observed actions) may be assigned color descriptors indicating that it is a browser process. Further, the specific process may be assigned an attribute indicating that it has exposed itself to external URLs or other external data. Subsequently, the same process may be observed to be taking an action suitable for an internal or system process, such as opening up shared memory to another process that has coloring descriptions indicating that it is a system process. When this last action is observed, an inconsistency in the various color descriptors between the subject of the action—the externally exposed browser process—and the target of the action may result in a well-defined IOC, which may be directly processed with immediate local action taken. The IOC may also or instead be reported externally as appropriate.

**[00128]** Thus, an endpoint 402 in an enterprise 410 may be instrumented with a coloring system 414 and monitor 421 to better detect potentially malicious activity using descriptors 420 that have been selected for relevance to threat detection along with a corresponding set of rules developed for the particular descriptors 420 that are being used to label or color various objects 418. By way of example, the object 418 may be a web browser that starts off being colored as a ‘browser’ and an ‘internet facing’ application. Based on this descriptor 420, a range of behaviors or actions may be considered normal, such as accessing remote network resources. However, if an object 418 colored with this descriptor 420 attempted to elevate privileges for a process, or to access a registry or system files, then this inconsistency in action may trigger a rule violation and result in an IOC 422.

**[00129]** In general, any action or series of actions that cumulatively invoke a particular reporting or action rule may be combined into an IOC 422 and communicated to the threat management facility 404. For example, an IOC 422 may include a malicious or strange behavior, or an indication of a malicious or strange behavior. The IOC 422 may be a normalized IOC that expresses one or more actions in a platform independent manner. That is, the IOC 422 may express a malicious behavior or suspected malicious behavior without reference to platform-specific information such as details of an operating system (e.g., iOS, MacOS, Windows,

Android, Linux, and so forth), hardware, applications, naming conventions, and so forth. Thus, a normalized IOC may be suitable for identifying a particular threat across multiple platforms, and may include platform independent processes, actions, or behaviors, or may express such process, actions, or behaviors in a platform independent manner. The normalized IOC may be generated from the IOC 422, e.g., it may be a converted version of the IOC 422 suitable for use with multiple platforms, or it may simply be any IOC 422 that has been created in a platform independent form. Process colorization (i.e., using the coloring system 414) as described herein may be used to create a normalized IOC.

**[00130]** In general, a threat management facility 404 for the enterprise 410 may include an IOC collector 426 that receives the IOC 422 from the endpoint 402 and determines an appropriate action. This may include any suitable remedial action, or where one or more IOCs 422 are inconclusive, continued monitoring or increased monitoring as appropriate.

**[00131]** The threat management facility 404 may provide a variety of threat management or monitoring tools 424, any of which may be deployed in response to IOCs 422 collected by the IOC collector 426. These tools 424 may include without limitation a scanning engine, whitelisting/blacklisting, reputation analysis, web filtering, an emulator, protection architecture, live protection, runtime detection, APT detection, network antivirus products, IOC detection, access logs, a heartbeat, a sandbox or quarantine system, and so forth.

**[00132]** The analysis facility 406 may provide a remote processing resource for analyzing malicious activities and creating rules 434 suitable for detecting IOCs 422 based on objects 420 and descriptors 420. It is generally contemplated that suitable attributes of certain descriptors 418 and one or more rules 434 may be developed together so that objects 418 can be appropriately labeled with descriptors 420 that permit invocation of rules 434 and creation of IOCs 422 at appropriate times. The analysis facility 406 may include a variety of analysis tools 428 including, without limitation, tools for regular expression, whitelisting/blacklisting, crowd sourcing, identifiers, and so forth. The analysis tools 428 may also or instead include information and tools such as URL look-ups, genotypes, identities, file look-up, reputations, and so forth. The analysis facility 406 may also provide numerous related functions such as an interface for receiving information on new, unknown files or processes, and for testing of such code or content in a sandbox on the analysis facility 406.

**[00133]** The analysis facility 406 may also or instead include a compromise detector 430, where the compromise detector 430 is configured to receive new threat information for

analysis and creation of new rules and descriptors as appropriate, as well as corresponding remedial actions. The compromise detector 430 may include any tools described herein or otherwise known in the art for detecting compromises or evaluating new threats in an enterprise 410.

**[00134]** In general, a rule 434 may be manually created with corresponding human-readable semantics, e.g., where a process is labeled as a browser process or other category or type that can be interpreted by a human. It should, however, be appreciated that the compromise detector 430 may also be configured to automatically generate descriptors 420 and rules 434 suitable for distribution to a threat management facility 404 and an endpoint 402. In this latter mode, the meaning of a particular descriptor 420 may not have a readily expressible human-readable meaning. Thus, it will be understood that attributes selected for relevance to threat detection may include conventional attributes, as well as attributes without conventional labels or meaning except in the context of a particular, computer-generated rule for threat detection.

**[00135]** In general, the analysis facility 406 may be within an enterprise 410, or the analysis facility 406 may be external to the enterprise 410 and administered, for example, by a trusted third party. Further, a third-party source 416 may provide additional threat data 438 or analyses for use by the analysis facility 406 and the threat management facility 404. The third-party resource 416 may be a data resource that provides threat data 438 and analyses, where the threat data 438 is any data that is useful in detecting, monitoring, or analyzing threats. For example, the threat data 438 may include a database of threats, signatures, and the like. By way of example, the third-party resource 416 may be a resource provided by The MITRE Corporation.

**[00136]** The system 400 may include a reputation engine 440 storing a plurality of reputations 442. The reputation engine 440 may include a reputation management system for the generation, analysis, identification, editing, storing, etc., of reputations 442. The reputation engine 440 may include reputation-based filtering, which may be similar to the reputation filtering discussed above with reference to Fig. 1. The reputation engine 440 may be located on the threat management facility 404 or the endpoint 402 as shown in Fig. 4, or the reputation engine 440 may be located elsewhere in the system 400. The reputation engine 440 may receive an IOC 422 or a stream of IOCs 422, and may generate or utilize reputations 442 for the IOCs 422. The reputation engine 440 may also or instead receive actions, behaviors, events, interactions, and so forth, and may generate or utilize reputations 442 for any of the foregoing.

The reputation engine 440 may generate or revise a reputation 442 based on behaviors, actions, events, interactions, IOCs 422, other reputations 442, a history of events, data, rules, state of encryption, colors, and so forth. The reputation engine 440 may utilize a third-party resource, e.g., for the third-party resource's reputation data.

**[00137]** The reputations 442 may include reputations for any of the objects 418 as described herein. In general, the reputations 442 may relate to the trustworthiness of the objects 418 or an attribute thereof (e.g., the source of the object 418, a behavior of the object 418, another object interacting with the object 418, and so forth). The reputations 442 may include lists of known sources of malware or known suspicious objects 418. The reputations 442 may also or instead include lists of known safe or trusted resources or objects 418. The reputations 442 may be stored in a reputations database included on the reputation engine 440 or located elsewhere in the system 400. The reputations 442 may be expressed in any suitable or useful level of granularity such as with discrete categories (e.g., trusted, untrusted, unknown, malicious, safe, etc.) or with a numerical score or other quantitative indicator. The reputations 442 may also be scaled.

**[00138]** In general, in the system 400 of Fig. 4, a malicious activity on the endpoint 402 may be detected by the IOC monitor 421, and a corresponding IOC 422 may be transmitted to the threat management facility 404 for remedial action as appropriate. The threat management facility 404 may further communicate one or more IOCs 422 to the analysis facility 406 for additional analyses and/or resolution of inconclusive results. Other details and variations are provided below. While the use of coloring and IOCs as contemplated herein can improve threat detection and remediation in a number of ways, the system 400 can be further improved with granular control over access to endpoint data using an encryption system. A system for key-based management of processes and files on an endpoint is now discussed in greater detail.

**[00139]** Fig. 5 illustrates a system for encryption management. Generally, the system 500 may include endpoints 502, an administration host 504, and a threat management facility 506, which may include policy manager 508 and key manager 510. The system 500 may provide for the management of users 512, policies 514, keys 516 (e.g., disposed on key rings 518), and endpoints 502 (e.g., from the administration host 504). The system 500 may utilize various storage and processing resources, which may be disposed in a cloud or the like.

**[00140]** The endpoints 502 may be any of the endpoints as described herein, e.g., with reference to the other figures. The endpoints 502 may also or instead include other end user

devices and other devices to be managed. The endpoints 502 may include a web browser for use by the users 512, with supporting cryptographic functions implemented using cryptographic libraries in the web browser. The endpoints 502 may communicate with the other components of the system 500 using any suitable communication interface, which may include Secure Socket Layer (SSL) encryption, Hypertext Transfer Protocol Secure (HTTPS), and so forth for additional security.

**[00141]** The endpoints 502 may include objects as described herein. For example, the endpoints 502 may include processes 520 and files 522. The processes 520 may be labeled (e.g., by a coloring system using descriptors as described above) in such a manner that the process is ‘IN,’ where the process 520 is in compliance with policies 514 administered for the endpoint 502 from a remote threat management facility 506, or the process is ‘OUT,’ where the process 520 is out of compliance with a policy (or a number of policies) in the policies 514 for an enterprise. This may provide IN processes 520A and OUT processes 520B as shown in Fig. 5. The files 522 may be similarly labeled by a coloring system with descriptors that identify each file 522 as IN, where the file 522 complies with the policies 514 and is accordingly encrypted using, e.g., a remotely managed key ring 518, or the file is OUT, where the file 522 does not conform to the policies 514 and is accordingly not encrypted using the remotely managed key ring 518. This may provide IN files 522A and OUT files 522B as shown in Fig. 5. One skilled in the art will recognize that other objects of the endpoint 502 or other components of the system 500 may be labeled in a similar manner where they are either IN or OUT. By coloring objects in this manner and basing key access on the corresponding color, the “IN” software objects may operate in a protected environment that objectively appears to be in compliance with the policies 514. Other files and processes may still be used on the endpoint 502, but they will operate in an “OUT” or unprotected environment that cannot obtain access to any of the “IN” content or functionality.

**[00142]** In an implementation, the system 500 may include determining whether an endpoint 502 is IN or OUT or whether a component of the endpoint 502 is IN or OUT, which may be based upon a set of rules (e.g., the rules outlined herein) or policies such as the policies 514 described herein. In some aspects, if the entire endpoint 502 is OUT – that is, out of compliance with one or more policies 514, the endpoint 502 will not have key access or access to any protected content. Conversely, if the endpoint 502 is IN, the endpoint 502 may have access to protected content. Thus in one aspect, the notion of IN/OUT may be applied at an endpoint level, and data protection may be a consequence of endpoint protection. Endpoint protection may

also or instead be applied at a more granular level, e.g., by determining whether executables, processes 520, files 522, etc., on the endpoint 502 are IN or OUT, which may be based upon rules or policies 514 as described herein.

**[00143]** The administration host 504 may include a web browser, which may include a cryptography library 524 and a web user interface (e.g., HTML, JavaScript, etc.). An administrator may utilize the web user interface to administer a key management system and perform administrative functions such as creating and distributing keys 516, establishing security policies, creating key hierarchies and rules, and so forth. The endpoint 502 may also include a cryptographic library 524 implementing cryptographic protocols for using key material in the key ring 518 to encrypt and decrypt data as needed.

**[00144]** The threat management facility 506 may include any of the threat management facilities or similar systems described herein. In general, the threat management facility 506 may include a policy manager 508 and key manager 510. Alternatively, one or more of the policy manager 508 and key manager 510 may be located elsewhere on a network.

**[00145]** The policy manager 508 may implement one or more policies 514, and maintain, distribute, and monitor the policies for devices in an enterprise. The policies 514 may include any policies 514 relating to secure operation of endpoints 502 in an enterprise. This may, for example, include hardware configuration policies, software configuration policies, communication policies, update policies, or any other policies relating to, e.g., the configuration of an endpoint 502, communications by an endpoint 502, software executing on an endpoint 502 and so forth. Policies 514 may include usage criteria based on, e.g., signatures, indications of compromise, reputation, user identity, and so forth. With respect to the key management system contemplated herein, the policies 514 may include a cryptographic protocol design, key servers, user procedures, and other relevant protocols, or these cryptographic protocols may be provided elsewhere for use by the policy manager 508. The policies 514 may also include any rules for compliance including those mentioned above or any other suitable rules or algorithms that can be applied to determine whether objects and components are 'IN' or 'OUT' as contemplated herein.

**[00146]** The key manager 510 may be part of the threat management facility 506, or it may be remotely managed elsewhere, e.g., in a remote cloud resource or the like. The key manager 510 may also or instead be disposed on the administration host 504 and one or more endpoints 502 in a manner independent of the threat management facility 506. In this manner, all cryptographic operations may be isolated from the threat management facility 506 and instead

may be performed by a web browser or the like executing on the administration host 504 or an endpoint 502. The key manager 510 may manage the keys 516, including managing the generation, exchange, storage, use, and replacement of keys 516. The key manager 510 may include a key ring 518, where the keys 516 are disposed on the key ring 518 using one root key 526. The key manager 510 may also or instead include a variety of key management and other secure processes, including without limitation, administrator registration, establishing trust to endpoints 502, key distribution to endpoints 502, policy deployment, endpoint status reporting, and local key backup.

**[00147]** The users 512 may have full access to encrypted data. Alternatively, the users 512 may have limited access to encrypted data, or no access to encrypted data. Access may be limited to users 512 using endpoints 502 that are deemed ‘IN’ by the system, as well as to processes 520 that are IN, as further described herein.

**[00148]** The keys 210 may include cryptographic keys in a cryptosystem, i.e., decryption keys. In one aspect, the keys 210 may be disposed on one key ring 218 using one root key 220. In general, the keys 210 may be created and managed using, e.g., symmetric key technology, asymmetric key technology, or any other key technology or combination of key technologies suitable for securing data in an enterprise including, for example the Data Encryption Standard (DES), Triple DES, Advanced Encryption Standard (AES), and so forth. The cryptosystem may also or instead include any suitable public key infrastructure or the like supporting the distribution and use of keys for encryption, digital signatures, and so forth.

**[00149]** The key ring 518 may facilitate simplified management of the system 500. For example, by reducing the data protection system down to a single key ring 518, the system can eliminate or reduce the overhead for management of keys 516. In one aspect, all of the data on a key ring 518 is protected by one root key 526. By reducing the data protection system down to a single key ring 518 protected by one root key 526, all privileged users 512 on uncompromised platforms can have access to all protected data. In this embodiment, data is either ‘IN’ (i.e., encrypted), or it’s ‘OUT’ (i.e., not encrypted). In one aspect, the default system does not include any additional shade of access control.

**[00150]** The cryptography library 524 may be disposed on the administration host 504 as shown in Fig. 5. The cryptography library 524 may also be disposed on the endpoint 502, e.g., in a web browser, or it may be disposed on another component of the system 500, or any combination of these. The cryptographic library 524 may be installed by an administrator. In

general, key material 530 from the key ring 518 may be stored in a cache 532 on the endpoint 502 within any suitable memory on the endpoint 502 for use in encryption and decryption as contemplated herein. As noted above, an enterprise that systematically uses coloring and indications of compromise can be improved through the use of a key management system as contemplated herein. This system may be still further improved with the addition of a heartbeat system that communicates heartbeats from an endpoint containing health and status information about the endpoint. A suitable heartbeat system is now described in greater detail.

**[00151]** Fig. 6 illustrates a threat management system using heartbeats. In general, a system 600 may include an endpoint 602, a gateway 604, a threat management system 606, and an enterprise management system 608 that manages an enterprise including the endpoint 602, the gateway 604, and one or more additional endpoints 610. Each of these components may be configured with suitable programming to participate in the detection and remediation of an advanced persistent threat (APT) or other malware threat as contemplated herein. Although the term “gateway” is used for the device between an endpoint and an external network, it will be appreciated that this device may also or instead include a switch, router, firewall, and/or other network elements, any of which may be included in the “gateway” as that term is used herein.

**[00152]** The endpoint 602 may be any of the endpoints described herein, or any other device or network asset that might join or participate in an enterprise network. The endpoint 602 may contain a threat 612 such as an advanced persistent threat, virus, or similar malware that resides on the endpoint 602. The threat 612 may have reached the endpoint 602 in a variety of ways, and may have been placed manually or automatically on the endpoint 602 by a malicious source. It will be understood that the threat 612 may take any number of forms and have any number of components. For example, the threat 612 may include an executable file that can execute independently, or the threat 612 may be a macro, plug-in, or the like that executes within another application. Similarly, the threat 612 may manifest as one or more processes or threads executing on the endpoint 602. The threat 612 may install from a file on the endpoint 602 or a file remote from the endpoint 602, and the threat 612 may create one or more other files such as data files or the like while executing. Advanced persistent threats can be particularly difficult to detect and remediate, and the systems and methods contemplated herein can advantageously provide improved sensitivity to such threats, as well as enabling improved remediation strategies. However, the systems and methods contemplated herein may also or instead be used to detect and remediate other types of malware threats. As such, in this context references to a particular

type of threat (e.g., an advanced persistent threat) should be understood to generally include any type of malware or other threat to an endpoint or enterprise unless a more specific threat or threat type is explicitly provided or otherwise clear from the context.

**[00153]** The threat 612 may be analyzed by one or more threat countermeasures on the endpoint 602 such as a whitelisting filter 614 that approves each item of code before executing on the endpoint 602 and prevents execution of non-whitelisted code. The endpoint 602 may also include an antivirus engine 616 or other malware detection software that uses any of a variety of techniques to identify malicious code by reputation or other characteristics. A runtime detection engine 618 may also monitor executing code to identify possible threats. More generally, any of a variety of threat detection techniques may be applied to the threat 612 before and during execution. In general, a threat 612 may evade these and other security measures and begin executing as a process 620 on the endpoint 602.

**[00154]** Network traffic 622 from the process 620 may be monitored and logged by a traffic monitor 624 on the endpoint 602. The traffic monitor 624 may, for example, log a time and a source of each network request from the endpoint 602. Where the endpoint 602 is within an enterprise network, the network traffic 622 may pass through the gateway 604 in transit to a data network such as the Internet. While the gateway 604 may be logically or physically positioned between the endpoint 602 and an external data network, it will be understood that other configurations are possible. For example, where the endpoint 602 is associated with an enterprise network but operating remotely, the endpoint 602 may form a VPN or other secure tunnel or the like to the gateway 604 for use of a threat management system 606, enterprise management system 608, and any other enterprise resources.

**[00155]** The endpoint 602 may use a heartbeat 626 to periodically and securely communicate status to the gateway 604. The heartbeat 626 may be created by a health monitor 628 within the endpoint 602, and may be transmitted to a remote health monitor 630, for example, at the gateway 604. The health monitor 628 may monitor system health in a variety of ways, such as by checking the status of individual software items executing on the endpoint 602, checking that antivirus and other security software is up to date (e.g., with current virus definition files and so forth) and running correctly, checking the integrity of cryptographic key stores, checking for compliance with enterprise security policies, and checking any other hardware or software components of the endpoint 602 as necessary or helpful for health monitoring. The health monitor 628 may thus condition the issuance of a heartbeat 626 on a

satisfactory status of the endpoint 602 according to any suitable criteria, enterprise policies, and other evaluation techniques. The remote health monitor 630 may also or instead be provided at the threat management facility 650, for example as part of the threat management system 606 or the enterprise management system 608.

**[00156]** The heartbeat 626 may be secured in any suitable manner so that the health monitor 630 can reliably confirm the source of the heartbeat 626 and the status of the endpoint 602. To this end, the heartbeat 626 may be cryptographically signed or secured using a private key so that the monitor 630 can authenticate the origin of the heartbeat 626 using a corresponding public key. In one aspect, the heartbeat 626 may include a combination of plaintext information and encrypted information, such as where the status information for the endpoint is provided in plaintext while a digital signature for authentication is cryptographically secured. In another aspect, all of the information in the heartbeat 626 may be encrypted.

**[00157]** In one aspect, a key vault 632 may be provided on the endpoint to support cryptographic functions associated with a secure heartbeat. An obfuscated key vault 632 may support numerous useful functions, including without limitation, private key decryption, asymmetric signing, and validation with a chain of trust to a specific root validation certificate. A variety of suitable key management and cryptographic systems are known in the art and may be usefully employed to support the use of a secure heartbeat as contemplated herein. The system may support a secure heartbeat in numerous ways. For example, the system may ensure that signing and decryption keys can only be used in authorized ways and inside an intended Access Control mechanism. The system may use “anti-lifting” techniques to ensure that a signing key can only be used when the endpoint is healthy. The system may ensure that attacking software cannot, without first reverse-engineering the key vault 632, extract the original key material. The system may also usefully ensure that an attacker cannot undetectably replace the public keys in a root certificate store, either directly or indirectly, such as in an attack that tries to cause the code to validate against a different set of root keys without directly replacing any keys in the root store.

**[00158]** A robust heartbeat 626 may usefully provide defensive mechanisms against reverse engineering of obfuscated content (e.g., the private key material stored in key vault 632, the code used to validate the correct running of the remainder of the systems as part of the heartbeat 626 code itself) and any anti-lifting protections to prevent malware from directly using the endpoint 602 (or the health monitor 628 on the endpoint 602) to continue to send out signed

heartbeat packets (e.g. stating that “all is well” with the endpoint) after security mechanisms have been impaired, disabled, or otherwise compromised in any way. Lifting in this manner by malicious code can be materially mitigated by providing statistical validation (e.g., with checksums of code) of call stacks, calling processes, and core processes. Likewise, statistical checks as well as checksum integrations into the cryptographic calculations may protect against code changes in the heartbeat 626 code itself.

**[00159]** A variety of useful techniques may be employed to improve security of the key vault 632 and the heartbeat 626. For example, the system may use domain shifting so that original key material is inferred based on hardware and software properties readily available to the key vault 632, and to ensure that key material uses non-standard or varying algorithms. Software properties may, for example, include readily determined system values such as hashes of nearby code. In another aspect, the keys may be domain shifted in a manner unique to the endpoint 602 so that the manner of statistical validation of call stacks and core software is unique to the endpoint 602. Further the key vault may be provisioned so that a public key stored in the key vault 632 is signed with a certificate (or into a certificate chain) that can be externally validated by a network appliance or other trusted third party or directly by the health monitor.

**[00160]** The heartbeat 626 may encode any useful status information, and may be transmitted from the endpoint 602 on any desired schedule including any periodic, aperiodic, random, deterministic, or other schedule. Configured in this manner, the heartbeat 626 can provide secure, tamper-resistant instrumentation for status of the endpoint 602, and in particular an indication that the endpoint 602 is online and uncompromised. A disappearance of the heartbeat 626 from the endpoint 602 may indicate that the endpoint 602 has been compromised; however this may also simply indicate that the endpoint 602 has been powered off or intentionally disconnected from the network. Thus, other criteria may be used in addition to the disappearance or interruption of the heartbeat 626 to more accurately detect malicious software. Some such techniques are described below, but it will be understood that this may include any supplemental information that might tend to make an attack on the endpoint 602 more or less likely. For example, if the heartbeat 626 is interrupted but the endpoint 602 is still sourcing network traffic, then an inference might suitably be made that the endpoint 602 is compromised.

**[00161]** The threat management system 606 may, in general, be any of the threat management systems described herein. The enterprise management system 608 generally provides tools and interfaces for administration of the enterprise and various endpoints 610 and

other resources or assets attached thereto. It will be understood that, the functions of the threat management system 606 and the enterprise management system 608 may vary, and general threat management and administration functions may be distributed in a variety of ways between and among these and other components. This is generally indicated in Fig. 6 as a threat management facility 650 that includes the threat management system 606 and the enterprise management system 608. It will be understood that either or both of these system may be administered by third parties on behalf of the enterprise, or managed completely within the enterprise, or some combination of these, all without departing from the scope of this disclosure. It will similarly be understood that a reference herein to a threat management facility 650 is not intended to imply any particular combination of functions or components, and shall only be understood to include such functions or components as explicitly stated in a particular context, or as necessary to provide countermeasures for advanced persistent threats as contemplated herein. It also should be understood that the heartbeat may be monitored and/or managed by the threat management system 606, the enterprise management system 608, or another component of the threat management facility 650.

**[00162]** The system 600 may include a certificate authority 660 or similar trust authority or the like (shown as a “trusted third party” in the figure). In order to provide a meaningfully secure heartbeat 626, the heartbeat 626 may be secured with reference to a trusted authority such as a certificate authority 660 that can issue cryptographic certificates allowing other entities to rely on assertions about identity (e.g., by enabling verification with a trusted third party), and to enable cryptographically secure communications. The cryptographic techniques for creating and using such certificates and relationships are well known, and are not repeated here. The certificate authority 660 may be administered by the enterprise management system 608 or some other internal resource of the enterprise, or the certificate authority 660 may be administered by a trusted third party such as any of a variety of commercially available certificate authorities or the like. Thus, the certificate authority 660, or some other similar cloud service or the like, may operate as a security broker to register, e.g., endpoints 602, 610, the gateway 604, the threat management facility 650, and so forth, and provide cryptographic material for each of the other trusting entities to securely communicate with one another.

**[00163]** Once registered with the certificate authority 660 in this fashion, the heartbeat may be used to establish trust between the endpoint 602 and other entities, and to validate the source of the heartbeat 626 when it is received. More generally, a heartbeat 626 secured in this

manner may provide an encrypted channel between network entities such as an endpoint 602 and the gateway 604 (or a firewall or the like). The nature of the communication may provide a technique for validating the source, as well as obfuscating the contents with encryption. Thus when, for example, the endpoint 602 provides information about a good/healthy state or a bad/compromised state, the recipient may rely on this state information and act accordingly.

**[00164]** Fig. 7 shows an architecture for endpoint protection in an enterprise network security system. In general, an endpoint may include a processing environment 702, a file system 706, a threat monitor 720 and a key wrapper 730.

**[00165]** The processing environment 702 may, for example, be any environment such as an operating system or the like suitable for executing one or more processes 704.

**[00166]** Each process 704 may be an instance of a computer program, portion of a computer program or other code executing within the processing environment 702. A process 704 may execute, e.g., on a processor, group of processors, or other processing circuitry or platform for executing computer-executable code. A process 704 may include executable computer code, as well as an allocation of memory, file descriptors or handles for data sources and sinks, security attributes such as an owner and any associated permissions, and a context including the content of physical memory used by the process 704. More generally, a process 704 may include any code executing on an endpoint such as any of the endpoints described herein.

**[00167]** The file system 706 is generally associated with an operating system that provides the processing environment 702, and serves as an intermediary between processes 704 executing in the processing environment 702 and one or more files 708 stored on the endpoint. The file system 706 may provide a directory structure or other construct to facilitate organization of the files 708, and the file system 706 generally supports file functions such as creating, deleting, opening, closing, reading, writing, and so forth.

**[00168]** An extension 710 may be included in the file system 706 by modifying the operating system kernel. While other programming techniques may be employed to perform the functions of an extension 710 as contemplated herein, direct modifications of or additions to the operating system permit the extension 710 to operate transparently to the processing environment 702 and the processes 704 without requiring any modifications or adaptations. The extension 710 may, for example, be implemented as a file system filter (in a MICROSOFT WINDOWS environment) or a mount point to a directory (in an APPLE iOS environment). The extension

710 to the files system as contemplated herein performs two concurrent functions. First, the extension 710 communicates with a threat monitor 720 in order to receive updates on the security status and exposure status of the processes 704 or the endpoint. Second the extension 710 communicates with a key wrapper 730 that provides key material for encrypting and decrypting data in the files 708. Finally, the extension 710 operates to conditionally provide encryption and decryption of the files 708 for the processes 704 based on a current security or exposure state, as described in greater detail below.

**[00169]** The threat monitor 720 may include any suitable threat monitoring, malware detection, antivirus program or the like suitable for monitoring and reporting on a security state of an endpoint or individual processes 704 executing thereon. This may include local threat monitoring using, e.g., behavioral analysis or static analysis. The threat monitor 720 may also or instead use reputation to evaluate the security state of processes 704 based on the processes 704 themselves, source files or executable code for the processes 704, or network activity initiated by the processes 704. For example, if a process 704 requests data from a remote URL that is known to have a bad reputation, this information may be used to infer a compromised security state of the endpoint. While a threat monitor 720 may operate locally, the threat monitor 720 may also or instead use remote resources such as a gateway carrying traffic to and from the endpoint, or a remote threat management facility that provides reputation information, malware signatures, policy information and the like for the endpoint and other devices within an enterprise such as the enterprise described above.

**[00170]** In general, the threat monitor 720 provides monitoring of a security state and an exposure state of the endpoint. The security state may, for example, be ‘compromised’, ‘secure’, or some other state or combination of states. This may be based on detections of known malware, suspicious activity, policy violations and so forth. The exposure state may be ‘exposed’ or ‘unexposed’, reflecting whether or not a particular process 704 or file 708 has been exposed to potentially unsafe content. Thus exposure does not necessarily represent a specific threat, but the potential for exposure to unsafe content. This may be tracked in a variety of ways, such as by using the coloring system described above with reference to Fig. 5.

**[00171]** The key wrapper 730 may contain a key ring 732 with one or more keys 734 for encrypting and decrypting files 708. The key ring 732 may be cryptographically protected within the key wrapper 730 in order to prevent malicious access thereto, and the key wrapper 730 may communicate with the extension 710 to provide keys 734 for accessing the files 708 at

appropriate times, depending, for example, on whether processes 704 are secure or exposed. In one aspect, the files 708 are stored in a non-volatile memory such as a disk drive, or in a random access memory that provides a cache for the disk drive, and the key wrapper 730 may be stored in a separate physical memory such as a volatile memory accessible to the operating system and the extension 710 but not to processes 704 executing in the user space of the processing environment 702.

**[00172]** In one aspect, every document or file on the endpoint may have a separate key. This may be, for example, a unique, symmetric key that can be used for encryption and decryption of the corresponding file. The key wrapper 730 may control access to the key material for encrypting and decrypting individual files, and may be used by the extension 710 to control access by individual processes 704 executing on the endpoint. As described herein, the extension 710 may generally control access to files 708 based on an exposure state, a security state, or other context such as the user of a calling process or the like. In the event of a severe compromise, or a detection of a compromise independent of particular processes, a key shredding procedure may be invoked to destroy the entire key wrapper 730 immediately and prevent any further access to the files 708. In such circumstances, the keys can only be recovered by the endpoint when a remediation is confirmed. Alternatively, the files may be accessed directly and decrypted from a secure, remote resource that can access the keys 734.

**[00173]** Fig. 8 shows a method for labeling network flows. In general, network flows may be explicitly labeled according to source applications (or source processes or the like) to permit tracking and management according to a source after a network flow leaves an endpoint.

**[00174]** As shown in step 802, the process 800 may include executing a process on an endpoint, such as any of the endpoints described herein. The process 800 may, for example, be associated with an application on the endpoint.

**[00175]** As shown in step 804, the process 800 may, at some point during its execution, create a network message. In general, a network message may include a header containing, e.g., a source address, a destination address, and other information necessary or helpful for supporting data network communications between the endpoint and an intended destination. The network message may, for example, be in response to an explicit user input such as directing a browser to a web address, or in response to an implicit user input such as requesting a software update or performing an online backup. This may also or instead be fully automated under any of a variety of normal circumstances. This may also, however, be outbound traffic from malware, and by

labeling the network traffic, malicious activity may be more easily detected after the network message has left the endpoint.

**[00176]** As shown in step 806, the process 800 may include adding a label to the network message. This may, for example, include generating a label that includes context information related to the network message, such as an identifier for an application or process. This identifier may be cryptographically signed or encrypted to protect the contents thereof, and may include additional information about, e.g., the status of the endpoint, the reputation of the application or the destination address, or any other information necessary or helpful for making improved routing or processing decisions about network traffic. Thus, the process 800 may include cryptographically signing the label, encrypting the contents of the label, or any other suitable cryptographic steps. Packet structures for network communication are generally known, and may for example include a header with control and routing information, along with a payload of data for communication to a recipient. The header may include various items of information useful for network communications such as a preamble that identifies a beginning of the packet, a source address, a destination address, a sequence number (e.g., for spreading a large file or the like across multiple packets/payload), a packet length, a packet type, a cyclic redundancy check or other checksum or the like, and any other useful information. A label as contemplated herein may be usefully incorporated into the header of such a packet in any suitable format and any suitable location. For example, internet protocol version '4' and version '6' both permit optional information to be included in a packet, and the label may be included as optional information. For example, the label may be included by encapsulating a packet within another packet that includes label information. For example, a label may be included by repurposing a portion of a packet header. For example, a label may be included by injecting additional label data into the packet, or for example, at the start of the data portion of the packet.

**[00177]** As shown in step 808, the process 800 may include transmitting the network message from the endpoint, e.g., through a network interface or the like to a data network.

**[00178]** According to the foregoing, in one aspect there is disclosed herein an endpoint configured as described above to label outbound network messages to identify a source process or application. Thus there is disclosed herein a system comprising an endpoint including a network interface configured to couple the endpoint in a communicating relationship with a data network. The endpoint may include a memory and a processor such as any of the memories, processors, or processing circuitry described herein. In general, the processor may be configured

to execute instructions stored in the memory to perform the steps of generating a label for a network message created by a process executing on the endpoint and associated with an application, where the network message includes a payload and a header and where the label includes an identifier for the application, adding the label to the header of the network message, and transmitting the network message through the network interface to a remote location on the data network. At the same time, a network device or the like may be configured to receive the network message and process it in a variety of ways, as further discussed herein.

**[00179]** As shown in step 810, the process 800 may include receiving a network message from an endpoint at a network device. The network message may be the network message described above, or any other packet or other data structure including a payload of data and header information such as a source address for the endpoint and a destination address for an intended recipient of the network message. The network message may also include a label that identifies an application that generated the network message on the endpoint. This may for example, be determined by examining the properties of a process associated with the application, or otherwise relating a process or other source of the network message with a particular application.

**[00180]** It should be understood that this receiving step may be performed at a wide range of network devices. In one aspect, this may include a router, a gateway, a firewall, a threat management facility for an enterprise, or any other network device for handling network communications to and/or from the endpoint. This may include a network device on a local network used by the endpoint, a cloud-based network device used in an enterprise network security system, a physical gateway between an enterprise network and an external network, or some combination of these. In one aspect, the firewall may be a local firewall for the endpoint that is configured to locally enforce routing rules based on information obtained in the label.

**[00181]** While a source application may be a useful item of information to be included in the label, it will be understood that other context information may also or instead be included. For example, this may include the identity of a process of an application that created the network message, a user of the process that created the network message, or other identity or source information, such as a name or other identity of the endpoint. For the endpoint name this may include a fully qualified name such as a device identifier that is unique within the domain of the endpoint. This may also or instead include a reputation of the application or process that created the network message if this is locally available to or determined by the endpoint. In this context,

the reputation may specify a particular reputation (e.g., known good, known bad, suspected bad, or the like) or the reputation may specify an absence of specific information (e.g., unknown). This may also or instead include health information about the endpoint or a process, coloring information for the process that identifies exposure of the process to other processes or data on the endpoint, reputation or exposure information for data used by the process, and so forth. The label may include status information related to the endpoint, such as an activity level of the endpoint, a status of software installed on the endpoint, a security condition of the endpoint (e.g., anti-malware software installed, a warning state of control software, and so forth), or other useful status indications in evaluating context of the network traffic and the like.

**[00182]** It should also be understood that, while the process 800 described herein may involve proactive labeling of network flows, or labeling of all network flows from an endpoint, the process 800 may also or instead be adapted to be responsive to detection and query by a network device, such as a threat management facility, a firewall, a gateway, a server, or the like. For example, a malware attack that exfiltrates data may begin with a file upload using a file transfer protocol (FTP) request or an HTTP PUT. When upload activity from the endpoint is detected at a network device, this may cause the network device to actively query the endpoint to request information about the application that requested the upload, such as an identifier of the application, a reputation of the application, or the like. For example, the network device may direct the endpoint to label all traffic from a user, from an application, from a process, or from a device. For example, the network device may direct the endpoint to label all traffic from processes that are currently active. For example, the network device may direct the endpoint to label all traffic from users that are currently active. For example, the network device may direct the endpoint to label all traffic from a user associated with a process. For example, the network device may first direct the endpoint to label all network traffic to identify a process. The network device may then identify a process of interest, and direct the endpoint to label all traffic associated with the process of interest, or associated with a user associated with the process of interest. These and other similar adaptations are intended to fall within the scope of this disclosure.

**[00183]** The process may also or instead be adapted to be responsive to detecting suspicious activity at the endpoint. Upon detecting suspicious activity, the endpoint may begin to label network traffic associated with the suspicious activity. The label may indicate that the labeled network traffic is associated with suspicious activity. The label may provide additional

context for the network traffic and associated application or process. A network device receiving labeled network traffic from an endpoint may then take additional measures, such as observing and recording the content of the network traffic, inspecting the content of the network traffic, verifying the reputation of the destination(s) of the network traffic, verifying the reputations of the application(s) generating the network traffic, determining whether there is other network traffic from other devices on the network directed to the destination, notifying an administrator, and so on.

**[00184]** The process may also or instead be adapted to be responsive to detecting suspicious activity at an endpoint that hosts one or more virtual machines. Upon detecting suspicious activity, the endpoint may label network traffic with context associated with the virtual machine. For example, the label may include an identifier for the virtual machine, e.g., if the virtual machine is not using a unique network address. The label may indicate that the labeled network traffic is associated with suspicious activity. The label may provide additional context for the network traffic and associated application or process. A network device, or a virtual machine host, or another virtual machine on the same host that is receiving labeled network traffic from the virtual machine, may then take additional measures.

**[00185]** The process may also or instead be adapted to be responsive to detecting suspicious activity at a network device. Upon detecting suspicious activity, the first network device may begin to label network traffic associated with the suspicious activity. The label may indicate that the labeled network traffic is associated with suspicious activity. The label may provide additional context for the labeled network traffic that may be available at the network device, for example, that the traffic transited the network device, or came from a particular subnet, or was recorded by the network device. The label may provide a reference to additional information collected by the network device about the suspicious network traffic. For example, the reference may be to a recording of the content of the network traffic. A network device receiving labeled network traffic from an endpoint may then take additional measures, such as observing and recording the content of the network traffic, inspecting the content of the network traffic, verifying the reputation of the destination(s) of the network traffic, verifying the endpoint(s) generating the network traffic, determining whether there is other network traffic from other devices on the network directed to the destination, notifying an administrator, and so forth.

**[00186]** As shown in step 812, the process 800 may include processing the network message on a network device that receives the network message to extract the label. This may, for example, include cryptographically verifying an authenticity of the label or a source of the label, decrypting the label to extract encoded information, or any other form of processing.

**[00187]** As shown in step 814, the process 800 may include processing the network message at a network device based on the label. While routing is generally contemplated as described below, a number of additional processing steps may usefully be performed. For example, the process 800 may include caching the label, for example associated with the application and a reputation of the application on the network device.

**[00188]** As shown in step 816, the process 800 may include routing the network message based on the application (or process or other source on the endpoint) that generated the network message. The routing may be a conditional routing based on the information in the label, or information determined from the information in the label. This may include routing rules or combinations of rules that may be based, for example, on user identity, process/application, reputation, and so forth. For example, the processor may be further configured to perform the steps of determining a reputation of the application and routing the network message to the destination address conditionally based on the reputation of the application that generated the network message. That is, if the reputation of the application is known and good, then the traffic may be routed as requested by the endpoint according to the routing information in the network traffic. However, if the reputation of the application is known and bad, then the traffic may be sequestered in any suitable manner. For example, the traffic may be dropped. Additional steps may be taken. For example, when the reputation of the application is uncertain, or other information is collectively inconclusive, a sandbox based on the endpoint may be created and used to communicate with the destination address to test for malicious activity. In another aspect, where information is inconclusive, the network traffic may be permitted, but more aggressive monitoring may be initiated until a conclusive evaluation of the source application can be obtained. For example, the network device may observe and record the network traffic, inspect the content of the network traffic, verify the reputation of the destination(s) of the network traffic, verify the reputations of the application(s) generating the network traffic, determine whether there is other network traffic from other devices on the network directed to the destination, notify an administrator, and so forth. In another aspect, an

indication of malware or other compromise on the endpoint may be used as a basis for initiating remediation of the endpoint.

**[00189]** Likewise, in the case of uncertain application reputation, some users may have additional permissions with respect to network traffic routing decisions. For example, a user designated as an accounting user, who may have access to more sensitive information may have traffic blocked, while a user in sales or customer support may be permitted with additional monitoring or safeguards.

**[00190]** In this context, it will be appreciated that an application or process identifier in the label may be used as a single item of information in a broader context for the network message. For example, the label may include a health status of the endpoint, other status information of the endpoint, a user identifier for a user of the application or a related process (e.g., automated machine activity—e.g., daemon, backup service—might be treated differently from human activity, categories of users may be treated differently than others, and the like), a name of the endpoint, reputation information for the source application, and so forth. Thus, the application identifier may be useful by itself, but the accuracy of attack detection may be improved still further by using additional context for the network message, which may include various pieces of contextual information inserted directly into the label for the network message, as well as other information available to a network device called upon to make a routing decision for a particular network message (e.g., allowing, blocking, rerouting, and so forth). The routing decision may include blocking the message. All such variations that would be apparent to one of ordinary skill in the art are intended to fall within the scope of this disclosure.

**[00191]** According to the foregoing, there is disclosed herein a system that includes a network device that may use labels provided by endpoints that identify source applications or processes and the like for network messages. In one aspect, a system disclosed herein includes a network device including a network interface configured to couple the network device in a communicating relationship with a data network that includes an endpoint, a memory on the network device, and a processor on the network device. The processor may be configured to execute instructions stored in the memory to perform the steps of receiving a network message from an endpoint through the network interface, where the network message includes a source address for the endpoint, a destination address for an intended recipient of the network message, a label that identifies an application that generated the network message on the endpoint, and a payload of data. The processor may be further configured to execute instructions stored in the

memory to perform the steps of processing the network message to extract the label, and routing the network message based on the application that generated the network message.

**[00192]** In one use of this system, the processor may be configured to determine a reputation of the application and to route the network message to the destination address conditionally based on the reputation of the application that generated the network message, for example as described above. Where information such as user information is available, the processor may be further configured to extract a user identifier from the label that identifies a user of the application on the endpoint, which may also be usefully employed for various intelligent routing processes as contemplated herein.

**[00193]** Fig. 9 illustrates an Internet Protocol (IP) packet format. In general, the packet 900 may include preamble information 902, a source address 904, a destination address 906, option information 908, and a payload of data 910. As described herein, the source address 904 may generally specify a network location that originated the packet 900, but may not provide any more specific information about a source user, a source machine, a source application, and other information as contemplated herein.

**[00194]** In one aspect, this information may be usefully added to the packet 900 by inserting relevant information into the option information 908, or in any other suitable location within the packet 900 or packet header. For example, the Internet Protocol version '4' (IPv4) packet may be configured such that additional information may be included in the packet header, which may be specified by the internet header length (IHL) data and an option type field. This field is identified in Fig. 9 as option information 908. In ordinary use, the option type field is sub-divided into a one bit copied flag, a two bit class field, and a five bit option number. These taken together form an eight bit value for the option type field. In one aspect, this field may be used to specify additional source information, or to identify a link or pointer to a location where such information can be obtained. For example, up to 60 bytes of optional, additional data may be included in the header, and a label may be included in this additional data. For example, where the packet is used within an enterprise domain, a portion of the header may be used by any source-aware networking components to identify a source, or to provide a pointer to source information which may, for example, be within the data 910 of the packet 900, within a separate packet, or stored at a network-accessible location such as a gateway, firewall, or threat management facility. An internet protocol version '6' (IPv6) packet may include a 20-bit flow label field, as well as the availability to provide extension headers to indicate optional

information. For example, the flow label field may be used to include a 20-bit label. For example, a 'Destination Options' extension header may be used to include label information that is not required to be examined by intermediate networking devices, but may be examined by a gateway or a final destination.

**[00195]** In one aspect, the option information 908 of an IPv4 packet may be used to designate the packet as containing source information 912 within the data 910 of the packet 900. Thus, for example, the first byte or bytes of the data 910 may be used to specify source information at any desired level of granularity including a user, a machine, an application, a process, and so forth, as well as combinations of these. The source information 912 within the data 910 may also or instead provide a link to a resource within the enterprise network where source information for the packet 900 can be retrieved. It will be understood that source information, whether stored within source information 912 in the packet 900, within the option information 908 in the packet header 914, or stored in some remote resource, may in general be encrypted to secure source information against unauthorized access, and/or the source information may be digitally signed to permit verification of authenticity with reference to a trusted third party or internal trust authority.

**[00196]** In one aspect, source information 912 may include a reference to a process or other context information for an endpoint. The endpoint may communicate context information to a network device through a separate channel, for example using the heartbeat 314 channel (see, e.g., Fig. 3). The context information may be stored on the endpoint, and the source information 912 in the packet 902 may include a reference to the context information. Thus, context information may be provided via a first channel, and a reference to the context information may be provided in a packet label. For example, each process on an endpoint may be assigned an identifier, and the identifier communicated to the network device with the operating system process information when the process is detected. The identifier may be included in the label. In this way, machine-specific information may not be included in the packet, but only a reference to the information. In some implementations, the identifier is changed periodically.

**[00197]** A flag within the packet header 914 may also or instead be used to identify when (and where) source labeling information is present for a packet 900.

**[00198]** In one aspect, the packet 900 may be configured to be compatible with other IP network traffic, so that packets can move into and out of the enterprise network without additional handling, such as by inserting source information 912 into the data 910 of the packet

payload as illustrated in Fig. 9. In another aspect, packets may be processed at the perimeter of the enterprise network, e.g., at a VPN gateway, enterprise network gateway, or any other perimeter network device, or other suitable router, switch, or other network device inside the enterprise network, with non-IP compliant packet data and/or structures being removed for outbound traffic, and where possible or applicable, reinserted for inbound traffic. In another aspect, traffic from outside the enterprise network may use a predetermined flag or the like to indicate that a packet contains source information compliant with a format for source information used within the enterprise network. More generally, any suitable techniques may be used to permit communication of source context information with network traffic within a network, while ensuring compatibility with an external network architecture such as TCP/IP or the like, and all such techniques that would be apparent to one of ordinary skill in the art are intended to fall within the scope of this disclosure.

**[00199]** Fig. 10 shows a method for secure labeling of network flows. A system for labeling network flows is generally described herein (e.g., above). A network enterprise security system may be improved by instrumenting endpoints to explicitly label network flows with cryptographically secure labels that identify an application or other source of each network flow. Cryptographic techniques may be used, for example, to protect the encoded information in the label from interception by third parties or to support cryptographic authentication of a source of each label. As discussed above, a label may provide health, status, or other heartbeat information for the endpoint, and may be used to identify compromised endpoints, to make routing decisions for network traffic (e.g., allowing, blocking, or re-rerouting), to more generally evaluate the health of an endpoint that is sourcing network traffic, or for any other useful purpose.

**[00200]** As shown in step 1002, the method 1000 may include receiving a network message from a process executing on the endpoint. In general, the process may be based on an application on the endpoint. As described herein, the network message may be a packet or the like that includes a payload and a header, and that is addressed to a remote location (e.g., using an Internet Protocol address, Uniform Resource Locator (URL), or other address) accessible from the endpoint through a data network.

**[00201]** As shown in step 1004, the method 1000 may include generating a label for the network message. As described herein, the label may include information about a source of the network message on the endpoint, e.g., an application, a process, a user or the like on the endpoint that originated the network message, or an identity of the endpoint itself, such as by

reference to a globally unique name of the endpoint within an enterprise domain. Thus, for example, the label may include an identifier for the application that generated the network message, an identifier for the endpoint that generated the network message, the identifier of a user of the process on the endpoint that generated the network message, or any other useful source identifier. The label may also or instead encode useful information about the status of the endpoint. For example, the label may include a health status of the endpoint or any other useful information.

**[00202]** As shown in step 1005, the method 1000 may include cryptographically processing the label information. This may, for example, include cryptographically signing the label to provide a signed label information. This signed label may be used, for example, to verify an identity of the endpoint, an identity of an application on the endpoint, an identity of a process executing on the endpoint, or any combination of these that is useful for identifying a source of the network message. This signed label may be verified, for example, using a public key for the source that is signing the label, or through any other suitable relationship with a trusted third party capable of cryptographically authenticating the source. Cryptographically processing the network message may also or instead include encrypting the label with a cryptographic key to provide an encrypted label that is secured against interception by third parties.

**[00203]** As shown in step 1006, the method 1000 may include adding the cryptographically process label (e.g., the signed label and/or encrypted label) to the header of the network message.

**[00204]** As shown in step 1008, the method 1000 may include transmitting the network message from the endpoint to the remote location through the data network.

**[00205]** As shown in step 1010, the method may include receiving the network message from the endpoint at a network device such as a gateway, a firewall, a router, a switch, or a threat management facility, any of which may be a hardware device physically positioned between the endpoint and an external data network, or a cloud-based device accessible to the enterprise network for the endpoint through a remote service or resource. The network message may be any of the network messages described herein, and may include an encrypted label that identifies an application that generated the network message, and/or a signed label that can be used to verify a source or the security of the label.

**[00206]** As shown in step 1011, the method 1000 may include processing the network message on the network device to extract the label.

**[00207]** As shown in step 1012, the method may include cryptographically processing the label. This may, for example, include verifying a source of the label or decrypting the label with a cryptographic key to provide a decrypted label.

**[00208]** As shown in step 1014, the method 1000 may include processing the network message based on the label, for example, by caching label or payload contents on the network device, requesting a scan or remediation of the endpoint that provided the message, or any other responsive processing. This may, for example, include receiving an indication that the endpoint is compromised, e.g., either embedded within the label or through another communication channel. For example, the network device may receive a heartbeat from the endpoint, and the indication of compromise may be inferred from an absence of the heartbeat when expected, or the indication of compromise may be explicitly added by the endpoint (e.g., by a security agent executing on the agent) into the heartbeat. This indication of compromise may be used as a basis for preventing routing of additional network traffic from the compromised application through the network device. This routing or forwarding restriction may be maintained, for example, until the expected heartbeat resumes from the endpoint/application, or until an explicit remediation measure is completed.

**[00209]** By labeling network flows by application in this manner, the network device that is handling the network flow can advantageously make routing or blocking decisions on an application-by-application basis rather than for an entire endpoint, thus limiting network restrictions to particular applications that are compromised while permitting other applications on an endpoint to continue normal network communications, or optionally network communications with some heightened level of scrutiny or security.

**[00210]** As shown in step 1016, the method 1000 may include routing the network message based on the decrypted label, for example using any of the techniques described herein.

**[00211]** According to the foregoing, there is disclosed herein a system for secure labeling of network flows. The system may include an endpoint including a network interface configured to couple the endpoint in a communicating relationship with a data network, a memory on the endpoint, and a processor on the endpoint. The processor may be configured as described herein to execute instructions stored in the memory to perform the steps of receiving a network message from a process executing on the endpoint, where the process is based on an application and where the network message includes a payload and a header. The network message may be addressed to a remote location accessible from the endpoint through a data

network. The processor may further be configured to execute instructions stored in the memory to perform the steps of generating a label for the network message, where the label includes information about a source of the network message on the endpoint, encrypting the label with a cryptographic key, thereby providing an encrypted label, adding the encrypted label to the header of the network message, and transmitting the network message from the endpoint to the remote location through the data network. Encrypting the label may include cryptographically signing the label or otherwise cryptographically processing the label as described herein.

**[00212]** Fig. 11 shows a process for verifying user presence on an endpoint. In general, a gateway or other network device may be configured to monitor endpoint behavior, and to request a verification of user presence at the endpoint under certain conditions suggesting, e.g., malware or other endpoint compromise. For example, when a network request is directed to a low-reputation or unknown network address, user presence may be verified to ensure that this action was initiated by a human user rather than automatically by malware or the like. User verification may be implicit, based on local behavior such as keyboard or mouse activity, or the user verification may be explicit, such as where a notification is presented on a display of the endpoint requesting user confirmation to proceed.

**[00213]** As shown in step 1102, the method 1100 may include connecting an endpoint to a data network through a gateway.

**[00214]** As shown in step 1104, the method 1100 may include detecting a network request at the gateway by a process executing on the endpoint. The network request may, for example, include any of the network messages described herein, or any other network request or the like. By way of non-limiting example, the network request may include a request for a download of an executable from a data network, an HTTP PUT request or other a request to an address contained in a Uniform Resource Locator entered, e.g., in a browser or similar application, an FTP file transfer request, or any other manually or automatically generated network request. It will be understood that in the context of this method 1100, this step 1104 may refer specifically to the detection or other identification of a network request that presents a potential or actual security risk, or other indication of compromise or the like. While it is also possible to test for the presence of a human user under other circumstances, the remaining steps disclosed in this method 1100 relate to a context in which the network message poses a potential risk. Thus for example, the network request may be directed to a remote resource that presents a potential security risk such as a remote resource with a low or unknown reputation, or a remote

resource known to be associated with malware distribution or the like. This may also or instead include a network request directed to an unknown address or a request directed to a known or suspected source of malware. The potential security risk may also be identified contextually, such as by detecting a pattern of connections to a single resource, an unusual sequence of destination addresses, or some other network activity (or other endpoint behavior) indicative of potential compromise of the endpoint.

**[00215]** As shown in step 1105, once potentially harmful network traffic is detected, other steps may be taken to provide additional context for evaluation or to reach a conclusion on further action before evaluating user presence. For example, the method 1100 may include requesting a reputation of an application on an endpoint that initiated the network request. This may permit direct blacklist or whitelist determinations or the like without requiring additional queries to the endpoint. This may also be particularly useful for identifying circumstances where human presence may be highly relevant to the detection of a security risk. For example, where a connection is opened to a remote location of low or unknown reputation and an upload of data from the endpoint to the remote location is initiated, this pattern would suggest malicious data leakage or exfiltration, and it may be highly advantageous to determine whether a human user is controlling the activity by the endpoint. Similarly, certain protocols that are commonly used to download malware (e.g., peer networking protocols) or upload sensitive data (e.g., a file or data transfer protocol) may provide useful context for an initial decision whether to test for the presence of a human user.

**[00216]** As shown in step 1106, the method 1100 may include evaluating a status of the endpoint to determine whether a user is present at the endpoint. In particular, if there is a reason to suspect the network request poses a potential security risk, or if there are other indicia of compromise for the endpoint, then a procedure may be initiated to evaluate whether a human user is present on the endpoint and/or whether a human user initiated the network request. Thus for example, evaluating the status may include querying the endpoint about whether the user is present. Evaluating the status may also or instead include transmitting a request to the endpoint for a user input. For example, the request may be presented as a pop-up window or other notification with text requesting a response. This may include a simple request such as “click here to continue,” or a more instructive narrative such as: “A potential security issue has been detected. Please click here to confirm that you requested the following network activity.” This may also or instead include other user input besides clicking, e.g., typing characters into a text

box or the like. In another aspect, a blocking page may be presented in a web browser that requires human interaction before further network activity can be undertaken.

**[00217]** In another aspect, the presence of a human user may be inferred from endpoint activity or status. For example, information about user activity or user presence may be included in a secure heartbeat from the endpoint, and evaluating the status may include examining the secure heartbeat from the endpoint for information about whether the user is present. The status of the endpoint may include other state information or historical activity indicative of user presence. For example, the status may include whether a user is logged in to the endpoint (which suggests that a human user is presently using the endpoint) or whether a display of the endpoint is locked (which suggests that a human user is not presently using the endpoint). The status may also or instead include other information useful for inferring whether the user is present. For example, the status may include a record of keyboard or mouse activity within a predetermined time window such as a prior minute, a prior five minutes, a prior ten minutes, and so forth.

**[00218]** In another aspect, the inquiry may include evaluating whether a user was present when the application that generated the network request was launched, which may be relevant in a number of circumstances – particularly where a user might initiate a large online backup or other lengthy download or upload procedure – where a user might intentionally initiate network activity and then log out from an endpoint.

**[00219]** In another aspect, the inquiry may simply include verifying the intent of a user that is known to be present. For example, this method 1100 may be usefully employed to prevent a phishing attack where a legitimate-appearing electronic mail communication contains a link to a malicious site. In this use case, the application may be an electronic mail client (e.g., of presumed unsafe reputation), and the network request coming from the application may be addressed to a location of unknown reputation. In this context, even though a user is known to be present because of the manual link navigation, remediation may be appropriately undertaken, such as by presenting a pop-up window or blocking page that explicitly identifies the target page and requests confirmation of the navigation instruction. For example, the notification may state: “You have requested access to [website]. Are you certain that you wish to navigate to this potentially unsafe network location?”

**[00220]** This and other techniques may be used to evaluate whether a user is present, or whether a user intended to initiate a particular network request.

**[00221]** As shown in step 1108, when a user is determined to be present in step 1106, other processing may be performed. This may include returning to step 1104 where additional network requests may be detected. Of course, other processing, rules, or decision logic may be applied in this context. For example, where a destination address is blacklisted or whitelisted, a decision on forwarding the network flow may be made regardless of whether a user is present. All such existing techniques for conditionally forwarding traffic, including those described elsewhere herein may be used in addition to or instead of detecting user presence, and a method may provide various types of threat detection and security management, either concurrently, sequentially, or in some other order or manner, all without departing from the scope of this disclosure. Thus it should be appreciated that the techniques described with reference to Fig. 11 (and the other figures in this disclosure) may be intended for use either alone or in combination with other techniques for managing network flows and protecting against security threats in an enterprise network.

**[00222]** As shown in step 1110, where it is determined in step 1106 that no user is present, the method 1100 may include remediation. This may include any type of response or remediation suitable to the circumstances. For example, when no user is detected at the endpoint, this may include executing a security measure in response to the network request such as blocking the network request, initiating a scan or remediation of the endpoint, blacklisting the destination address, updating a reputation of the destination address, or any other suitable action. It will be understood that the remediation may be contextually selected based, for example, on other available information such as a reputation of the destination address for the network request, or a reputation of an application on the endpoint making the request.

**[00223]** In other circumstances, no remediation may be required. For example, where both the endpoint application and the remote location are updaters, then an endpoint that automatically requests a download without a user present may be a part of a legitimate, periodic software updating protocol. Similarly, where the endpoint application and the remote location both have a known, good reputation, then the activity may be permitted even in the absence of a human user of the endpoint.

**[00224]** When the remediation has been initiated or completed, the method 1100 may return to step 1104 where a next network request may be detected. In general, this method 1100 may be repeated indefinitely so long as an endpoint is active or powered on, or so long as the endpoint is connected to a data network or is otherwise generating network requests.

**[00225]** According to the foregoing, there is disclosed herein a system for verifying user presence associated with network flows. In one aspect, the system may include a network device including a network interface configured to couple in a communicating relationship with a data network that includes an endpoint, a memory on the network device, and a processor on the network device. The processor may be configured to execute instructions stored in the memory to perform the steps of connecting an endpoint to the network device through the network interface, detecting a network request by a process executing on the endpoint to a remote resource that presents a potential security risk, evaluating a status of the endpoint to determine whether a user is present at the endpoint, and executing a security measure in response to the network request when no user is detected at the endpoint. The network device may be any physical or virtual network device including without limitation a firewall, a gateway, a threat management facility, or the like.

**[00226]** In complementary fashion, a system may include an endpoint configured to provide information about user present to such a network device, for example, by including user presence information in a secure heartbeat or other communication to the network device. In one aspect, the system may include an endpoint including a network interface configured to couple the endpoint in a communicating relationship with a data network, a memory on the endpoint, and a processor on the endpoint. The processor may be configured to execute instructions stored in the memory to perform the steps of monitoring a status of the endpoint, periodically creating a status indicator characterizing the status, generating a user presence indicator containing an indication of whether a human user is present at the endpoint, creating a heartbeat containing the status indicator and the user presence indicator, and transmitting the heartbeat through the network interface to a gateway that couples the endpoint to a data network.

**[00227]** Fig. 12 shows a method for mitigating distributed denial of service attacks from an enterprise network.

**[00228]** In general, a distributed denial of service (DDoS) attack is a coordinated network attack using malicious bots that are installed on a large number of computers throughout a network and then concurrently launched to flood a target website with requests. As described herein, when a potential trigger for a DDoS attack is followed by an increase in network traffic from the endpoint to a high reputation network address, the endpoint is treated as a DDoS service bot and isolated from the network until remediation can be performed.

**[00229]** As shown in step 1202, the method 1200 may include monitoring outbound traffic from an endpoint in an enterprise network. The endpoint and the enterprise network may in general be any of the endpoints and enterprise networks described herein. In general, the method 1200 may be performed by a gateway or threat management facility for an enterprise network, or by an endpoint within the enterprise network, or some combination of this. Thus, for example, monitoring may include monitoring outbound traffic at a gateway for the enterprise network, or monitoring may include monitoring the outbound traffic includes monitoring the outbound traffic at the endpoint.

**[00230]** As shown in step 1204, the method 1200 may include detecting a potential trigger event for a distributed denial of service attack, e.g., detecting a potential anti-sandbox operation. For example, this may be a trigger event on the endpoint. In an aspect, the detected operation (e.g., the trigger event) may be specific to detecting a process monitoring clock and wall time drift. In this manner, the potential trigger event may be a detection of a process on the endpoint that monitors a clock, e.g., by periodically checking a system clock or a remote clock, or by creating and monitoring an internal timer or the like. The potential trigger event may also or instead be a detection of a process on the endpoint that monitors a file on the endpoint. For example, a launch instruction for a DDoS attack may be encoded in an innocent looking file that is maintained on an endpoint and updated with an instruction to launch by some other process or the like. This type of arrangement can obfuscate the presence of the infection and the launch instruction by permitting control of the DDoS service bot via a simple text file or other configuration file on the infected machine. In another aspect, the potential trigger event may be a receipt of a communication on the endpoint from a low reputation network address, or a receipt of a message from a (known or suspected) command and control site. The message in this context may thus include detecting network traffic, either directly or through an intermediary, between the endpoint and the low reputation network address, or from a known, suspected, or inferred command and control site. Thus, the potential trigger event may be a detection of network traffic between the endpoint and a low reputation network address, and/or the potential trigger event is a detection of network traffic from a command and control site.

**[00231]** In another aspect, the potential trigger event may be access to a publicly available communication channel, such as a group online chat (e.g., chat boards, internet relay chat (IRC), group instant message (IM), etc.) While these types of trigger events may be monitored directly on an endpoint, trigger events may also or instead occur at locations other

than the endpoint. Thus for example, the potential trigger event may include a detection of an increase in network traffic to a high reputation network address from one or more other endpoints in the enterprise network. Where, for example, this is known to be associated with a software update or the like, the event may be usefully filtered. However, where this increase in traffic is very large and specifically targeted at a particular domain or address, then there is a high likelihood that a DDoS attack is being initiated from devices within the enterprise network.

**[00232]** In one aspect, a conviction of a potential triggering event may be reached after aggregating small, suspicious activities over time, such as repeated contact to remote sites with unknown reputation coupled to use of a local, low-reputation process. This may include any of the techniques described herein for tracking, e.g., exposure of processes or files, use of low-reputation applications, processes, or websites, and so forth. In another aspect, this may include coloring a process as a clock watching process, e.g., using the techniques above, so that a relationship of the process with other processes and objects can be monitored over time. Similarly, a process that continuously checks for a file that is not present, or otherwise performs a repetitive read, search, file access or the like may be colored as a potential DDoS service bot. Other suspicious activity that might contribute to a multi-factor detection may include, for example, use of unattended network connections (e.g., connections without any corresponding keyboard or mouse activity), or use of IP addresses or ports that are typically reserved for interactive or support connections.

**[00233]** As shown in step 1206, the method 1200 may include observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event. While it is conceptually possible to launch a DDoS attack on low-reputation websites, a malicious attack of this type is typically directed toward a prominent business or website having a well-established and good reputation. This may, for example, include observing an increase in network traffic beyond a predetermined threshold, which may, for example be specified in terms of number of requests, number of outbound packets, or any other suitable objective measure for an amount of network communications. This latter technique may be particularly helpful for distinguishing malicious activity from legitimate activity such as increased traffic at an ecommerce site due to a sale or promotion, or increased traffic at a software vendor site due to a security update or other scheduled or newly released software update.

**[00234]** In some implementations, the potential triggering event(s) or an observation of an increase in network traffic may take place on a network device, such as a firewall or gateway. The network device may observe the triggering event if the event involves network activity. The network device may be able to observe that multiple devices are involved in the triggering event or increased traffic. For example, the activity of one or more endpoints communicating with a low reputation site followed by an increase in network activity by those same endpoints may be indicative of compromise.

**[00235]** As shown in step 1208, the method 1200 may include, in response to the increase in network traffic, identifying the endpoint as having a compromised state in which the endpoint serves as a distributed denial of service bot for the distributed denial of service attack.

**[00236]** As shown in step 1210, the method 1200 may include preventing network traffic from the endpoint until the compromised state is remediated. This may include preventing all network traffic, or only preventing network traffic directed to the network address that is being attacked. More generally, any technique for isolating the endpoint or otherwise preventing or inhibiting the endpoint from participating in the attack may be usefully employed in step 1210 to reduce or prevent attacks from originating within the enterprise network.

**[00237]** As shown in step 1212, the method 1200 may include remediating the endpoint. This may include any suitable techniques for removing the DDoS service bot, or removing or repairing corresponding infections of otherwise legitimate software executing on the endpoint. After the endpoint has been remediated, the endpoint may return to normal network communications, including legitimate use of the high-reputation network address that was subjected to the DDoS attack.

**[00238]** In one aspect, there is disclosed herein a threat management facility for managing an enterprise network. The threat management facility may generally include a network interface for coupling the threat management facility to the enterprise network, a memory, and a processor. The processor may be configured by computer executable code stored in the memory to protect against a distributed denial of service attack originating from within the enterprise network by performing the steps of monitoring outbound traffic from an endpoint in the enterprise network, detecting a potential trigger event on the network for a distributed denial of service attack, observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event, in response to the increase in network traffic, identifying the endpoint as having a compromised state wherein the endpoint

serves as a distributed denial of service bot, and preventing network traffic from the endpoint until the compromised state is remediated. The processor may be further configured to initiate a remediation of the endpoint using any suitable remediation techniques.

**[00239]** Fig. 13 shows a method for detecting endpoint compromise based on network usage history. In general, most applications will communicate to either a handful of URL's (e.g., for updates or access to remote data, file servers, mail servers, and so forth) or to hundreds or thousands of URLs (e.g., for a browser). This bifurcation of the connection history is useful and knowable at a per-user or per-device level, and application level and can be combined with an applications reputation score and the URL reputation score to better identify if a specific as yet un-classified URL should be considered malicious, suspect or benign. Thus, in the context of network activity by an endpoint in an enterprise network, malware detection can be improved by using a combination of reputation information for a network address that is accessed by the endpoint with reputation information for an application on the endpoint that is accessing the network address. This information, when combined with a network usage history for the application, provides improved differentiation between malicious network activity and legitimate, user-initiated network activity. The following method exploits this multi-factor technique. In general, the following method 1300 may be instantiated on an endpoint, on a threat management facility or gateway for an enterprise network used by the endpoint, or some combination of these.

**[00240]** As shown in step 1302, the method 1300 may include detecting a connection from an application executing on an endpoint in the enterprise network to a network address outside the enterprise network. In general, this may be any communication, interactive session, or the like that involves a network connection between the endpoint and the network address.

**[00241]** As shown in step 1304, the method 1300 may include determining a reputation of the application using the connection. This may include, for example, dynamically evaluating a reputation of the application, e.g., using application coloring or other techniques as described herein. This may also or instead include using a reputation lookup at the threat management facility (e.g., where the method 1300 is performed by the threat management facility) or requesting the reputation from the threat management facility (e.g., where the method 1300 is performed by an endpoint). In another aspect, this may include using a reputation lookup in a local reputation cache maintained on the endpoint, or retrieving reputation information from a local reputation cache on the endpoint to a threat management facility. In general, any suitable

method for characterizing reputation may be used. For example, any discrete or continuous method for characterizing reputation may be used. For example, the reputation may be expressed as a reputation score, or the reputation may be expressed as a category such as by specifying a reputation as one of known malicious, suspect, unknown, or known good. In one aspect, the reputation for the network address may be based on information from other devices about the network address, which may be accumulated, for example, at the threat management facility or any other network resource that receives periodic reporting of reputation information for network addresses and aggregates reported information into a reputation score or category. In another aspect, reputation may be established hierarchically, e.g., on a per-user basis, a global basis, or some intermediate basis such as for a group of users or a corporate entity, where patterns of usage might be expected to be fairly consistent or predictable.

**[00242]** As shown in step 1306, the method 1300 may include determining a reputation of the network address accessed by the endpoint through the connection. A variety of techniques may be employed to evaluate the reputation of the network address. For example, this may include using a reputation lookup at the threat management facility (e.g., where the method 1300 is performed by the threat management facility) or requesting the reputation from the threat management facility (e.g., where the method 1300 is performed by an endpoint). In another aspect, this may include using a reputation lookup in a local reputation cache maintained on the endpoint or on another local network device, such as a firewall, shared drive, etc.

**[00243]** As shown in step 1308, the method 1300 may include determining a network usage history for the application. This may be based on a monitored history of network activity, for example, by using a log of network activity maintained on the endpoint. Thus in one aspect, the method 1300 may include retrieving a network usage history for an application by retrieving the network usage history from a log on the endpoint. This may also include monitoring network activity by the application and storing the network usage history for the application in the log for subsequent use. In general, the log of network activity may be updated concurrently with the malware detection process contemplated herein. Thus for example, where a service on the endpoint uses the network address, the log may be updated accordingly based on an inferred relationship between the service and the application. More specifically, the method 1300 may include associating a network communication to the network address by a service executing on the endpoint with the application when the application controls the network communication, and adding the network communication to the network usage history for the application.

**[00244]** The history of network activity may also be categorized in any useful manner. For example, a particular application may have a predictable pattern of usage, for example, where the application only communicates with a specific, small group of URLs, or appears otherwise limited in the number and type of URLs, ports, etc. In this case, the application may be associated with that general pattern of usage, which may be associated with the particular application on a particular device, or all users of the application across an enterprise. This pattern may be used as a condensed representation of the history of usage in order to avoid a necessity of re-analyzing the complete usage history each time an application reputation is evaluated.

**[00245]** As shown in step 1310, the method 1300 may include evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application.

**[00246]** As shown in step 1312, the method 1300 may include remediating the endpoint or initiating a remediating action on the endpoint when a compromised condition is detected.

**[00247]** The method 1300 described above may be instantiated on a threat management facility. Thus in one aspect there is disclosed herein a threat management facility for managing an enterprise network, the threat management facility comprising a network interface for coupling the threat management facility to the enterprise network, a memory, and a processor. The processor may be configured by computer executable code stored in the memory to protect against malicious network activity in the enterprise network by performing the steps of detecting a connection from an endpoint in an enterprise network to a network address, determining a reputation of an application using the connection, determining a reputation of the network address, determining a network usage history for the application, evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application, and initiating remediation action on the endpoint when the compromised condition is detected. In this threat management facility, the reputation of the application may be at least one of known malicious, suspect, unknown, or known good, and the reputation of the network address may similarly be at least one of known malicious, suspect, unknown, or known good. The network usage history for the application may be retrieved from a log on the endpoint.

**[00248]** Fig. 14 shows useful malware detection insights based on the monitoring described with reference to Fig. 13. In a variety of circumstances, such as a known bad application contacting a known bad URL, there may be additional insight that may be obtained

from a history of a controlling application in combination with other information, such as application reputation and URL reputation. Such a combination may serve as the basis for additional monitoring or the deployment of remedial measures. In some implementations, an application with a positive reputation that accesses a URL with a positive reputation may be considered to be an acceptable access even if the new URL has not been accessed by the application previously. In other cases, where the application or the URL is unknown or suspect, the access of the URL outside of historical usage may be suspicious or indicative of compromise. In such cases, the application may be investigated or quarantined, the network access of the application may be restricted, etc.

**[00249]** In some implementations, an application with a positive reputation that accesses a new URL with a positive reputation may be suspect if the new URL access is one of multiple new URL accesses (e.g., a number of accesses meeting a threshold within a predetermined time window). For example, if the new URL involves an increase in network traffic from an endpoint directed to a high reputation network address after a potential trigger event, the traffic may be suspect or indicative of compromise. The access of a new URL may be a trigger 1204 (see Fig. 12). For example, more than a number (e.g., 3, 5, 10) of new URL accesses within a certain time period (e.g., a minute, hour, day, or week, etc.) may be indicative of compromise even if the new URL(s) accessed have a positive reputation.

**[00250]** Specifically, the chart 1400 of Fig. 14 shows the correlation between application reputation 1402, URL (and/or IP address) reputation 1404, a history of a controlling application 1406, and insights or conclusions when a new IP/URL connection is detected 1408. In a first example 1410 shown in the figure, if the reputation of the application is ‘known good,’ the reputation of the URL (or IP address) is ‘known malicious,’ and the history of a controlling application includes a constrained usage of the URL (e.g., a new URL not included in a list such that the URL is considered to be outside of the application’s historic usage), the insight or conclusion when this new IP/URL connection is detected may include that the ‘good’ application appears to have been exploited. In a second example 1420 shown in the figure, if the reputation of the application is ‘known good,’ the reputation of the URL (or IP address) is ‘suspect,’ and the history of a controlling application includes a constrained usage of the URL (e.g., a new URL not included in a list such that there is a new URL/IP address usage that is combined with an already suspect URL/IP address), the insight or conclusion when this new IP/URL connection is detected may include an indication of an exploit and probably a confirmed malicious URL/IP

address. In a third example 1430 shown in the figure, if the reputation of the application is ‘suspect,’ the reputation of the URL (or IP address) is ‘unknown,’ and the history of a controlling application includes a constrained usage of the URL (e.g., a new URL not included in a list), the insight or conclusion when this new IP/URL connection is detected may include a contribution to suspicion that there is a possible exploit. In a fourth example 1440 shown in the figure, if the reputation of the application is ‘known good,’ the reputation of the URL (or IP address) is ‘unknown,’ and the history of a controlling application includes a constrained usage of the URL (e.g., a new URL not included in a list), the insight or conclusion when this new IP/URL connection is detected may include that the ‘good’ application is going to a new location, but that this looks suspect. In a fifth example 1450 shown in the figure, if the reputation of the application is ‘known good,’ the reputation of the URL (or IP address) is ‘known good,’ and the history of a controlling application includes a constrained usage of the URL (e.g., a new URL not included in a list), the insight or conclusion when this new IP/URL connection is detected may include that the strange activity looks like a good application going to a new good location, e.g., good activity.

**[00251]** It should also be appreciated that the foregoing techniques may be usefully adapted for improved detection of malicious command and control activity from unknown network addresses, e.g., even when there is no reputation information available. Command and control IP addresses and domains are constantly being moved by malicious actors in order to avoid detection and remain active following discovery. For an unknown URL, aspects of the corresponding outbound or inbound packet communicated over the network connection may also be inspected instead of or in addition to the techniques described above, for example to determine whether the packet pattern, packet shape, packet size, header, payload profile, or other characteristics match a known command and control format. This latter combination – an unknown network address and pattern-matching to a known command and control packet characteristic – may also have independent diagnostic significance. That is, this combination may be useful for detection of an advanced persistent threat in combination with the foregoing methods or as an independent detection rule or technique.

**[00252]** For example, when an application communicates to command and control it is likely that the first few hundred packets sent are similar to other malicious communications sent to other addresses used by the attacker, or over different communication ports. If the command and control software is more stable than the infection malware – that is, the remote control

software remains consistent even as the distributed malicious code changes, then it may be possible to detect command and control handshake and data exchange even when the packets are encrypted.

**[00253]** Fig. 15 shows a method for local proxy detection. Protocol suites such as hypertext transfer protocol (HTTP) using secure socket layer (SSL) can facilitate secure network communications. In some instances, local proxies may be maliciously installed on an endpoint to support a so-called man-in-the-browser attack or other similar exploit in which the proxy intercepts these secure communications, steals a session key for communications, and then decrypts what are assumed to be secure communications with the remote resource for malicious use. For the secure communication protocols, network addresses are typically expressed as numeric internet protocol addresses rather than the human-readable uniform resource locators (URLs) that a user might enter into a browser address bar, e.g., during typical unsecure network browsing. This property can be exploited to differentiate between secure and insecure communications, and to detect certain instances where a malicious proxy has been deployed to intercept network traffic with an endpoint.

**[00254]** As shown in step 1502, a method 1500 for malicious proxy detection may include monitoring outbound traffic from an endpoint in an enterprise network. The outbound traffic may, for example, include communications with remote network addresses outside the enterprise network. Other malware detection techniques may also or instead be employed alongside the local proxy detection contemplated herein. Thus monitoring network traffic may usefully include looking up a reputation for destinations of outbound communications and/or monitoring outbound traffic at a gateway for the enterprise network.

**[00255]** As shown in step 1504, the method 1500 may include detecting use of a secure communication protocol with a request from the endpoint. The request may be directed from the endpoint to one of the remote network addresses. The secure communication protocol may, for example, include any protocol for securing communications between an endpoint and a remote resource using a session key or the like. This may, for example, include hypertext transfer protocol using secure socket layer or transport layer security.

**[00256]** As shown in step 1506, the method 1500 may include identifying a plaintext network address within the request. In a secure communication protocol, the network address is typically specified using a numeric address space such as that specified by IPv4 or IPv6. Thus any alphanumeric plaintext within the address may provide a useful indicator of a possible local

proxy. Thus the plaintext may generally be any alphanumeric address other than an internet protocol address consisting of a group of numbers. For example, this may be a uniform resource locator specifying a top level domain or other textual address, pathname or the like.

**[00257]** As shown in step 1508, the method 1500 may include, in response to identifying a plaintext network address in the request, initiating remediation of a potentially maliciously local proxy on the endpoint. This may include deploying a wide range of remediation measures or techniques. For example, initiating remediation may include quarantining the endpoint until the potentially malicious local proxy can be removed. Not all local proxies are malicious, thus initiating remediation may also or instead include reversing a malware identification for the potentially malicious proxy by identifying a non-malicious source of the local proxy. In another aspect, initiating remediation may include verifying the initial malware identification, e.g., by performing the steps of identifying a process that initiated the request, determining a reputation of the process, and then conditionally confirming the malware identification by performing the steps of: if the reputation of the process is low or unknown, confirming the malware identification, and if the reputation of the process is good, confirming the malware identification only when a calling process for the process that has a low or unknown reputation.

**[00258]** In another aspect, there is disclosed herein a device for malicious local proxy detection. The device may include a network interface, a memory, and a processor. The processor may be configured by computer executable code stored in the memory to perform the steps of monitoring outbound traffic from an endpoint in an enterprise network, detecting use of a secure communication protocol with a request from the endpoint, identifying a plaintext network address within the request, and in response to identifying a plaintext network address in the request, remediating a potentially maliciously local proxy on the endpoint. In general, the device may be the endpoint that is being monitored, e.g., where the detection system is locally deployed, or the device may be a gateway or threat management facility for the enterprise network, or some combination of these.

**[00259]** The above systems, devices, methods, processes, and the like may be realized in hardware, software, or any combination of these suitable for a particular application. The hardware may include a general-purpose computer and/or dedicated computing device. This includes realization in one or more microprocessors, microcontrollers, embedded microcontrollers, programmable digital signal processors or other programmable devices or processing circuitry, along with internal and/or external memory. This may also, or instead,

include one or more application specific integrated circuits, programmable gate arrays, programmable array logic components, or any other device or devices that may be configured to process electronic signals. It will further be appreciated that a realization of the processes or devices described above may include computer-executable code created using a structured programming language such as C, an object oriented programming language such as C++, or any other high-level or low-level programming language (including assembly languages, hardware description languages, and database programming languages and technologies) that may be stored, compiled or interpreted to run on one of the above devices, as well as heterogeneous combinations of processors, processor architectures, or combinations of different hardware and software. In another aspect, the methods may be embodied in systems that perform the steps thereof, and may be distributed across devices in a number of ways. At the same time, processing may be distributed across devices such as the various systems described above, or all of the functionality may be integrated into a dedicated, standalone device or other hardware. In another aspect, means for performing the steps associated with the processes described above may include any of the hardware and/or software described above. All such permutations and combinations are intended to fall within the scope of the present disclosure.

**[00260]** Embodiments disclosed herein may include computer program products comprising computer-executable code or computer-usable code that, when executing on one or more computing devices, performs any and/or all of the steps thereof. The code may be stored in a non-transitory fashion in a computer memory, which may be a memory from which the program executes (such as random access memory associated with a processor), or a storage device such as a disk drive, flash memory or any other optical, electromagnetic, magnetic, infrared or other device or combination of devices. In another aspect, any of the systems and methods described above may be embodied in any suitable transmission or propagation medium carrying computer-executable code and/or any inputs or outputs from same.

**[00261]** The elements described and depicted herein, including in flow charts and block diagrams throughout the figures, imply logical boundaries between the elements. However, according to software or hardware engineering practices, the depicted elements and the functions thereof may be implemented on machines through computer executable media having a processor capable of executing program instructions stored thereon as a monolithic software structure, as standalone software modules, or as modules that employ external routines, code, services, and so forth, or any combination of these, and all such implementations may be within

the scope of the present disclosure. Examples of such machines may include, but may not be limited to, personal digital assistants, laptops, personal computers, mobile phones, other handheld computing devices, medical equipment, wired or wireless communication devices, transducers, chips, calculators, satellites, tablet PCs, electronic books, gadgets, electronic devices, devices having artificial intelligence, computing devices, networking equipment, servers, routers and the like. Furthermore, the elements depicted in the flow chart and block diagrams or any other logical component may be implemented on a machine capable of executing program instructions. Thus, while the foregoing drawings and descriptions set forth functional aspects of the disclosed systems, no particular arrangement of software for implementing these functional aspects should be inferred from these descriptions unless explicitly stated or otherwise clear from the context. Similarly, it may be appreciated that the various steps identified and described above may be varied, and that the order of steps may be adapted to particular applications of the techniques disclosed herein. All such variations and modifications are intended to fall within the scope of this disclosure. As such, the depiction and/or description of an order for various steps should not be understood to require a particular order of execution for those steps, unless required by a particular application, or explicitly stated or otherwise clear from the context. Absent an explicit indication to the contrary, the disclosed steps may be modified, supplemented, omitted, and/or re-ordered without departing from the scope of this disclosure. Numerous variations, additions, omissions, and other modifications will be apparent to one of ordinary skill in the art. In addition, the order or presentation of method steps in the description and drawings above is not intended to require this order of performing the recited steps unless a particular order is expressly required or otherwise clear from the context.

**[00262]** The method steps of the implementations described herein are intended to include any suitable method of causing such method steps to be performed, consistent with the patentability of the following claims, unless a different meaning is expressly provided or otherwise clear from the context. So for example performing the step of X includes any suitable method for causing another party such as a remote user, a remote processing resource (e.g., a server or cloud computer) or a machine to perform the step of X. Similarly, performing steps X, Y and Z may include any method of directing or controlling any combination of such other individuals or resources to perform steps X, Y and Z to obtain the benefit of such steps. Thus method steps of the implementations described herein are intended to include any suitable

method of causing one or more other parties or entities to perform the steps, consistent with the patentability of the following claims, unless a different meaning is expressly provided or otherwise clear from the context. Such parties or entities need not be under the direction or control of any other party or entity, and need not be located within a particular jurisdiction.

**[00263]** It will be appreciated that the methods and systems described above are set forth by way of example and not of limitation. Numerous variations, additions, omissions, and other modifications will be apparent to one of ordinary skill in the art. In addition, the order or presentation of method steps in the description and drawings above is not intended to require this order of performing the recited steps unless a particular order is expressly required or otherwise clear from the context. Thus, while particular embodiments have been shown and described, it will be apparent to those skilled in the art that various changes and modifications in form and details may be made therein without departing from the spirit and scope of this disclosure and are intended to form a part of the invention as defined by the following claims, which are to be interpreted in the broadest sense allowable by law.

## CLAIMS

What is claimed is:

1. A computer program product for protecting against distributed denial of service attacks from an enterprise network, the computer program product comprising computer executable code embodied in a non-transitory computer readable medium that, when executing on one or more computing devices, performs the steps of:
  - monitoring outbound traffic from an endpoint in the enterprise network;
  - detecting a potential trigger event on the endpoint for a distributed denial of service attack;
  - observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event;
  - in response to the increase in network traffic, identifying the endpoint as having a compromised state in which the endpoint serves as a distributed denial of service bot for the distributed denial of service attack; and
  - preventing network traffic from the endpoint until the compromised state is remediated.
2. The computer program product of claim 1 wherein the potential trigger event is a detection of a process on the endpoint that monitors a clock.
3. The computer program product of claim 1 wherein the potential trigger event is a detection of a process on the endpoint that monitors a file on the endpoint.
4. The computer program product of claim 1 wherein the potential trigger event is a detection of network traffic between the endpoint and a low reputation network address.
5. The computer program product of claim 1 wherein the potential trigger event is a detection of network traffic from a command and control site.
6. The computer program product of claim 1 wherein the potential trigger event is a detection of an increase in network traffic to the high reputation network address from one or more other endpoints in the enterprise network.

7. The computer program product of claim 1 wherein monitoring the outbound traffic includes monitoring the outbound traffic at a gateway for the enterprise network.
8. The computer program product of claim 1 wherein monitoring the outbound traffic includes monitoring the outbound traffic at the endpoint.
9. A method comprising:
  - monitoring outbound traffic from an endpoint in an enterprise network;
  - detecting a potential trigger event for a distributed denial of service attack;
  - observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event;
  - in response to the increase in network traffic, identifying the endpoint as having a compromised state wherein the endpoint serves as a distributed denial of service bot; and
  - preventing network traffic from the endpoint until the compromised state is remediated.
10. The method of claim 9 wherein the potential trigger event is a detection of a process on the endpoint that monitors a clock.
11. The method of claim 9 wherein the potential trigger event is a detection of a process on the endpoint that monitors a file on the endpoint.
12. The method of claim 9 wherein the potential trigger event is a receipt of a message on the endpoint from a low reputation network address.
13. The method of claim 9 wherein the potential trigger event is a receipt of a message on the endpoint from a command and control site.
14. The method of claim 9 wherein the potential trigger event is a detection of an increase in network traffic to the high reputation network address from one or more other endpoints in the enterprise network.
15. The method of claim 9 wherein monitoring the outbound traffic includes monitoring the outbound traffic at a gateway for the enterprise network.

16. The method of claim 9 wherein monitoring the outbound traffic includes monitoring the outbound traffic at the endpoint.

17. The method of claim 9 wherein observing an increase in network traffic includes observing an increase in network traffic beyond a predetermined threshold.

18. The method of claim 9 further comprising remediating the endpoint.

19. A threat management facility for managing an enterprise network, the threat management facility comprising:

a network interface for coupling the threat management facility to the enterprise network;

a memory; and

a processor, the processor configured by computer executable code stored in the memory to protect against a distributed denial of service attack originating from within the enterprise network by performing the steps of monitoring outbound traffic from an endpoint in the enterprise network, detecting a potential trigger event on the network for a distributed denial of service attack, observing an increase in network traffic from the endpoint directed to a high reputation network address after detecting the potential trigger event, in response to the increase in network traffic, identifying the endpoint as having a compromised state wherein the endpoint serves as a distributed denial of service bot, and preventing network traffic from the endpoint until the compromised state is remediated.

20. The threat management facility of claim 19 wherein the processor is further configured to initiate a remediation of the endpoint.

21. A computer program product for protecting against malicious network activity in an enterprise network, the computer program product comprising computer executable code embodied in a non-transitory computer readable medium that, when executing on one or more computing devices, performs the steps of:

detecting a connection from an application executing on an endpoint in the enterprise network to a network address outside the enterprise network;

determining a reputation of the application using a first reputation lookup at a threat management facility for the enterprise network;

determining a reputation of the network address using a second reputation lookup at the threat management facility;

determining a network usage history for the application using a log of network activity maintained on the endpoint;

evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application; and

initiating remediation action on the endpoint when the compromised condition is detected.

22. The computer program product of claim 21 wherein the reputation of the application is at least one of known malicious, suspect, unknown, or known good.

23. The computer program product of claim 21 wherein the reputation of the network address is at least one of known malicious, suspect, unknown, or known good.

24. The computer program product of claim 21 wherein the reputation of the network address is based on crowd-sourced information about the network address.

25. The computer program product of claim 21 further comprising code that performs the step of monitoring network activity by the application and storing the network usage history for the application in the log on the endpoint.

26. The computer program product of claim 21 further comprising code that performs the step of associating a network communication to the network address by a service executing on the endpoint with the application when the application controls the network communication, and adding the network communication to the network usage history for the application.

27. The computer program product of claim 21 wherein the one or more computing devices include the endpoint.

28. The computer program product of claim 21 wherein the one or more computing devices include the threat management facility.

29. A method comprising:  
detecting a connection from an endpoint in an enterprise network to a network address;  
determining a reputation of an application using the connection;  
determining a reputation of the network address;  
determining a network usage history for the application;  
evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application; and  
initiating remediation action on the endpoint when the compromised condition is detected.
30. The method of claim 29 where the reputation of the application is requested from a threat management facility for the enterprise network.
31. The method of claim 29 wherein the reputation of the application is at least one of known malicious, suspect, unknown, or known good.
32. The method of claim 29 wherein the reputation of the network address is requested from a threat management facility for the enterprise network.
33. The method of claim 29 wherein the reputation of the network address is at least one of known malicious, suspect, unknown, or known good.
34. The method of claim 29 wherein the network usage history for the application is retrieved from a log on the endpoint.
35. The method of claim 29 wherein the reputation of the network address is based on crowd-sourced information about the network address.
36. The method of claim 29 further comprising monitoring network activity by the application and storing the network usage history for the application in a log on the endpoint.
37. The method of claim 36 further comprising associating a network communication to the network address by a service executing on the endpoint with the application when the

application controls the network communication, and adding the network communication to the network usage history for the application.

38. A threat management facility for managing an enterprise network, the threat management facility comprising:

- a network interface for coupling the threat management facility to the enterprise network;

- a memory; and

- a processor configured by computer executable code stored in the memory to protect against malicious network activity in the enterprise network by performing the steps of detecting a connection from an endpoint in an enterprise network to a network address, determining a reputation of an application using the connection, determining a reputation of the network address, determining a network usage history for the application, evaluating the endpoint for a compromised condition based on the reputation of the application, the reputation of the network address, and the network usage history for the application, and initiating remediation action on the endpoint when the compromised condition is detected.

39. The threat management facility of claim 38 wherein the reputation of the application is at least one of known malicious, suspect, unknown, or known good and wherein the reputation of the network address is at least one of known malicious, suspect, unknown, or known good.

40. The threat management facility of claim 38 wherein the network usage history for the application is retrieved from a log on the endpoint.

41. A computer program product for detecting a malicious proxy on an endpoint in an enterprise network, the computer program product comprising computer executable code embodied in a non-transitory computer readable medium that, when executing on one or more computing devices, performs the steps of:

- monitoring outbound traffic from the endpoint with remote network addresses outside the enterprise network;

- detecting use of a secure communication protocol with a request from the endpoint to one of the remote network addresses;

- identifying a plaintext network address within the request; and

in response to identifying a plaintext network address in the request, initiating remediation of a potentially maliciously local proxy on the endpoint.

42. The computer program product of claim 41 wherein the secure communication protocol is a hypertext transfer protocol using secure socket layer or transport layer security.

43. The computer program product of claim 41 wherein the plaintext network address includes an alphanumeric address other than an internet protocol address.

44. The computer program product of claim 41 wherein the plaintext network address includes a uniform resource locator.

45. The computer program product of claim 41 wherein monitoring outbound traffic includes looking up a reputation for destinations of outbound communications.

46. The computer program product of claim 41 wherein monitoring outbound traffic includes monitoring outbound traffic at a gateway for the enterprise network.

47. The computer program product of claim 41 wherein initiating remediation of the potentially malicious local proxy includes quarantining the endpoint until the potentially malicious local proxy can be removed.

48. The computer program product of claim 41 wherein initiating remediation of the potentially malicious local proxy includes reversing a malware identification for the potentially malicious proxy by identifying a non-malicious source of the local proxy.

49. The computer program product of claim 41 wherein initiating remediation of the potentially malicious local proxy includes verifying a malware identification for the potentially malicious local proxy by performing the steps of:

identifying a process that initiated the request;

determining a reputation of the process;

if the reputation of the process is low or unknown, confirming the malware identification; and

if the reputation of the process is good, confirming the malware identification only when a calling process for the process that has a low or unknown reputation.

50. A method comprising:
  - monitoring outbound traffic from an endpoint in an enterprise network;
  - detecting use of a secure communication protocol with a request from the endpoint;
  - identifying a plaintext network address within the request; and
  - in response to identifying a plaintext network address in the request, remediating a potentially maliciously local proxy on the endpoint.
51. The method of claim 50 wherein the secure communication protocol is a hypertext transfer protocol using secure socket layer or transport layer security.
52. The method of claim 50 wherein the plaintext network address includes an alphanumeric address other than an internet protocol address.
53. The method of claim 50 wherein the plaintext network address includes a uniform resource locator.
54. The method of claim 50 wherein monitoring outbound traffic includes looking up a reputation for destinations of outbound communications.
55. The method of claim 50 wherein monitoring outbound traffic includes monitoring outbound traffic at a gateway for the enterprise network.
56. The method of claim 50 wherein initiating remediation of the potentially malicious local proxy includes quarantining the endpoint until the potentially malicious local proxy can be removed.
57. The method of claim 50 wherein initiating remediation of the potentially malicious local proxy includes reversing a malware identification for the potentially malicious proxy by identifying a non-malicious source of the local proxy.

58. The method of claim 50 wherein initiating remediation of the potentially malicious local proxy includes verifying a malware identification for the potentially malicious local proxy by performing the steps of:

identifying a process that initiated the request;

determining a reputation of the process;

if the reputation of the process is low or unknown, confirming the malware identification; and

if the reputation of the process is good, confirming the malware identification only when a calling process for the process that has a low or unknown reputation.

59. A device comprising:

a network interface;

a memory; and

a processor, the processor configured by computer executable code stored in the memory to perform the steps of monitoring outbound traffic from an endpoint in an enterprise network, detecting use of a secure communication protocol with a request from the endpoint, identifying a plaintext network address within the request, and in response to identifying a plaintext network address in the request, remediating a potentially maliciously local proxy on the endpoint.

60. The device of claim 59 wherein the secure communication protocol is a hypertext transfer protocol using secure socket layer or transport layer security.

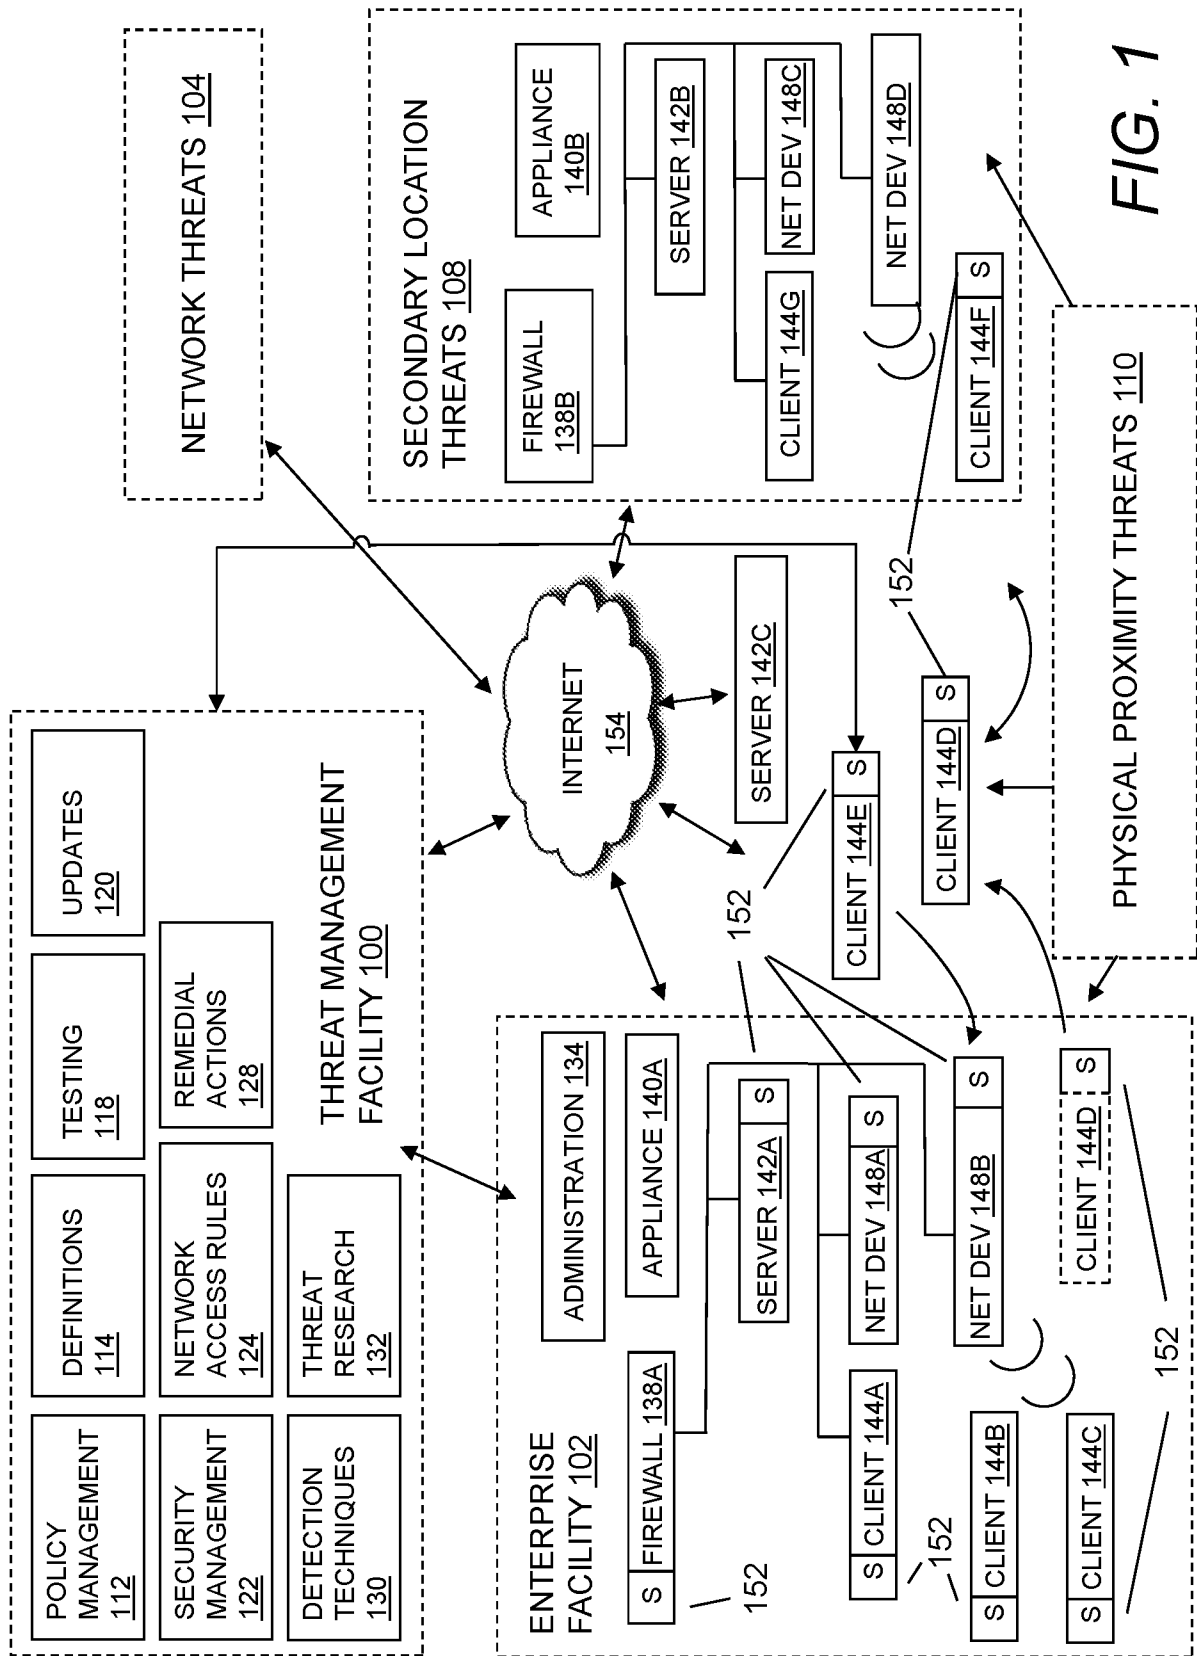


FIG. 1

FIG. 2





FIG. 3






FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




Fig. 9

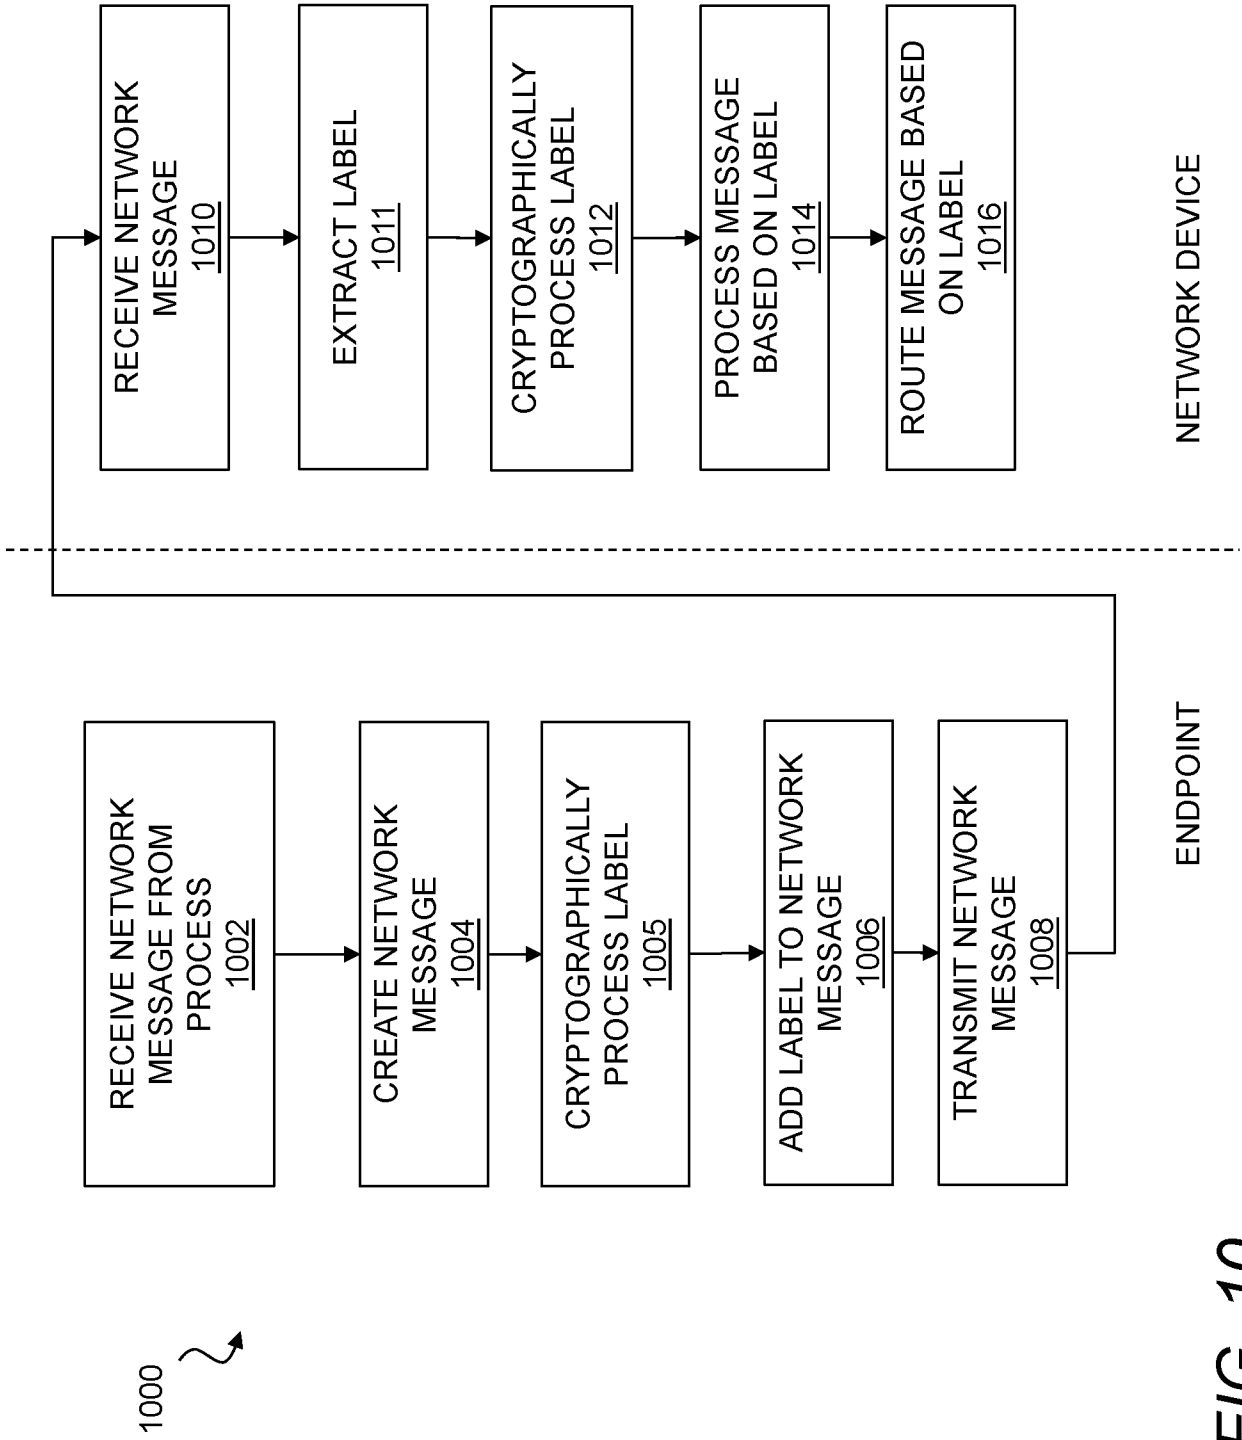


FIG. 10

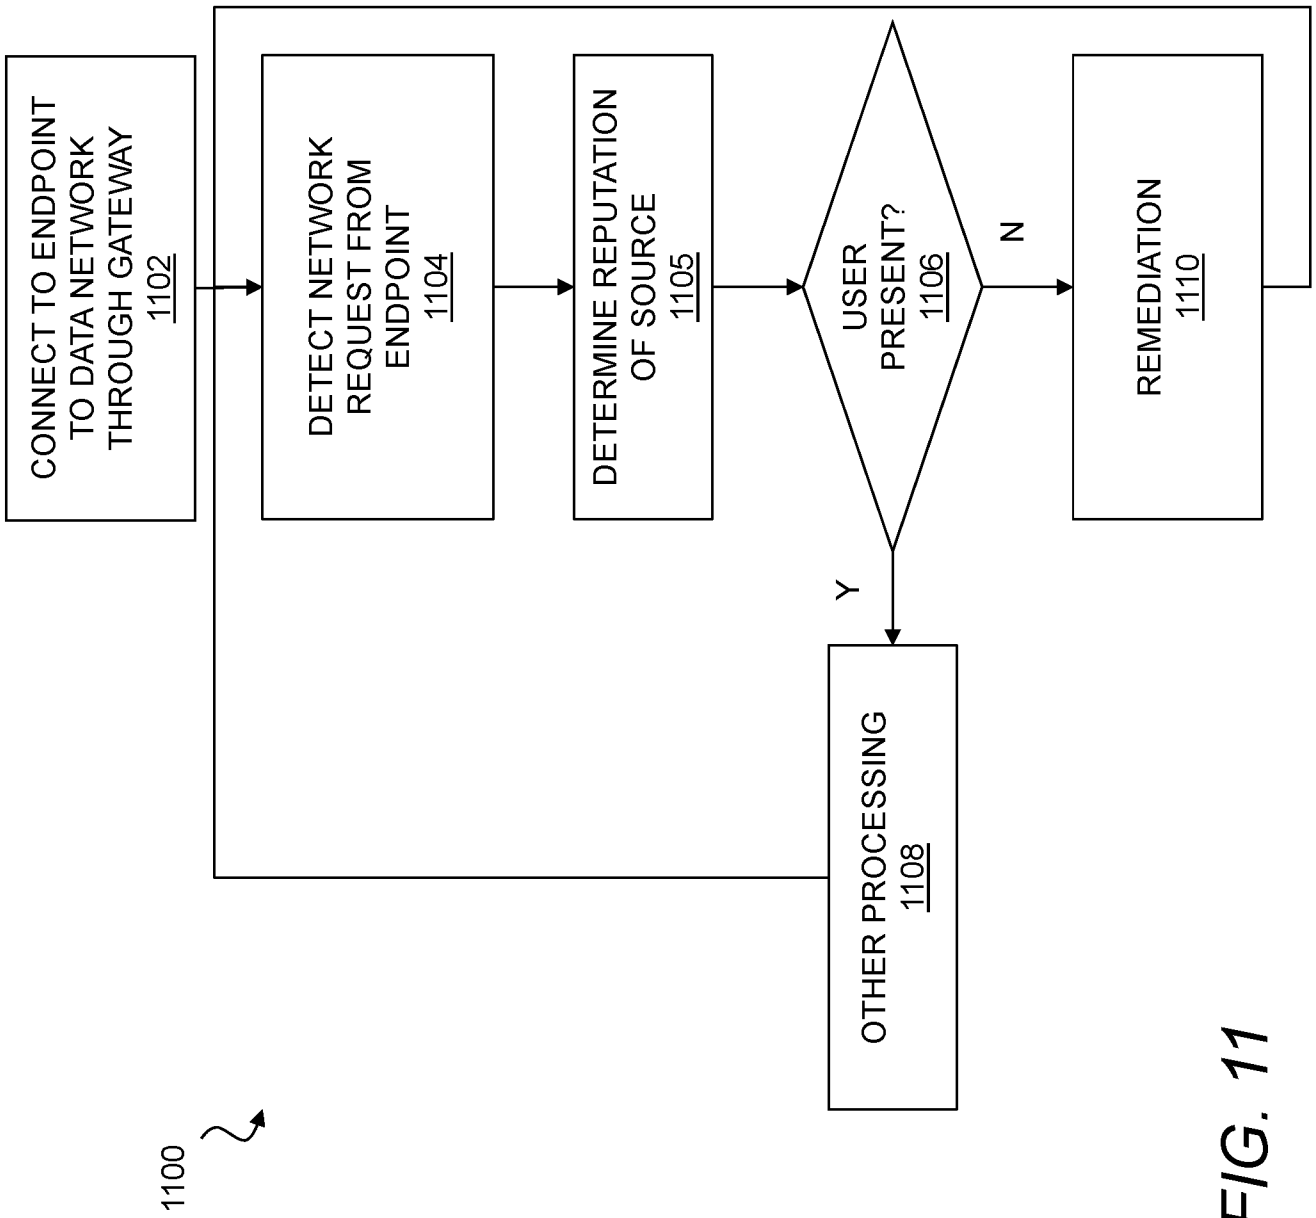


FIG. 11

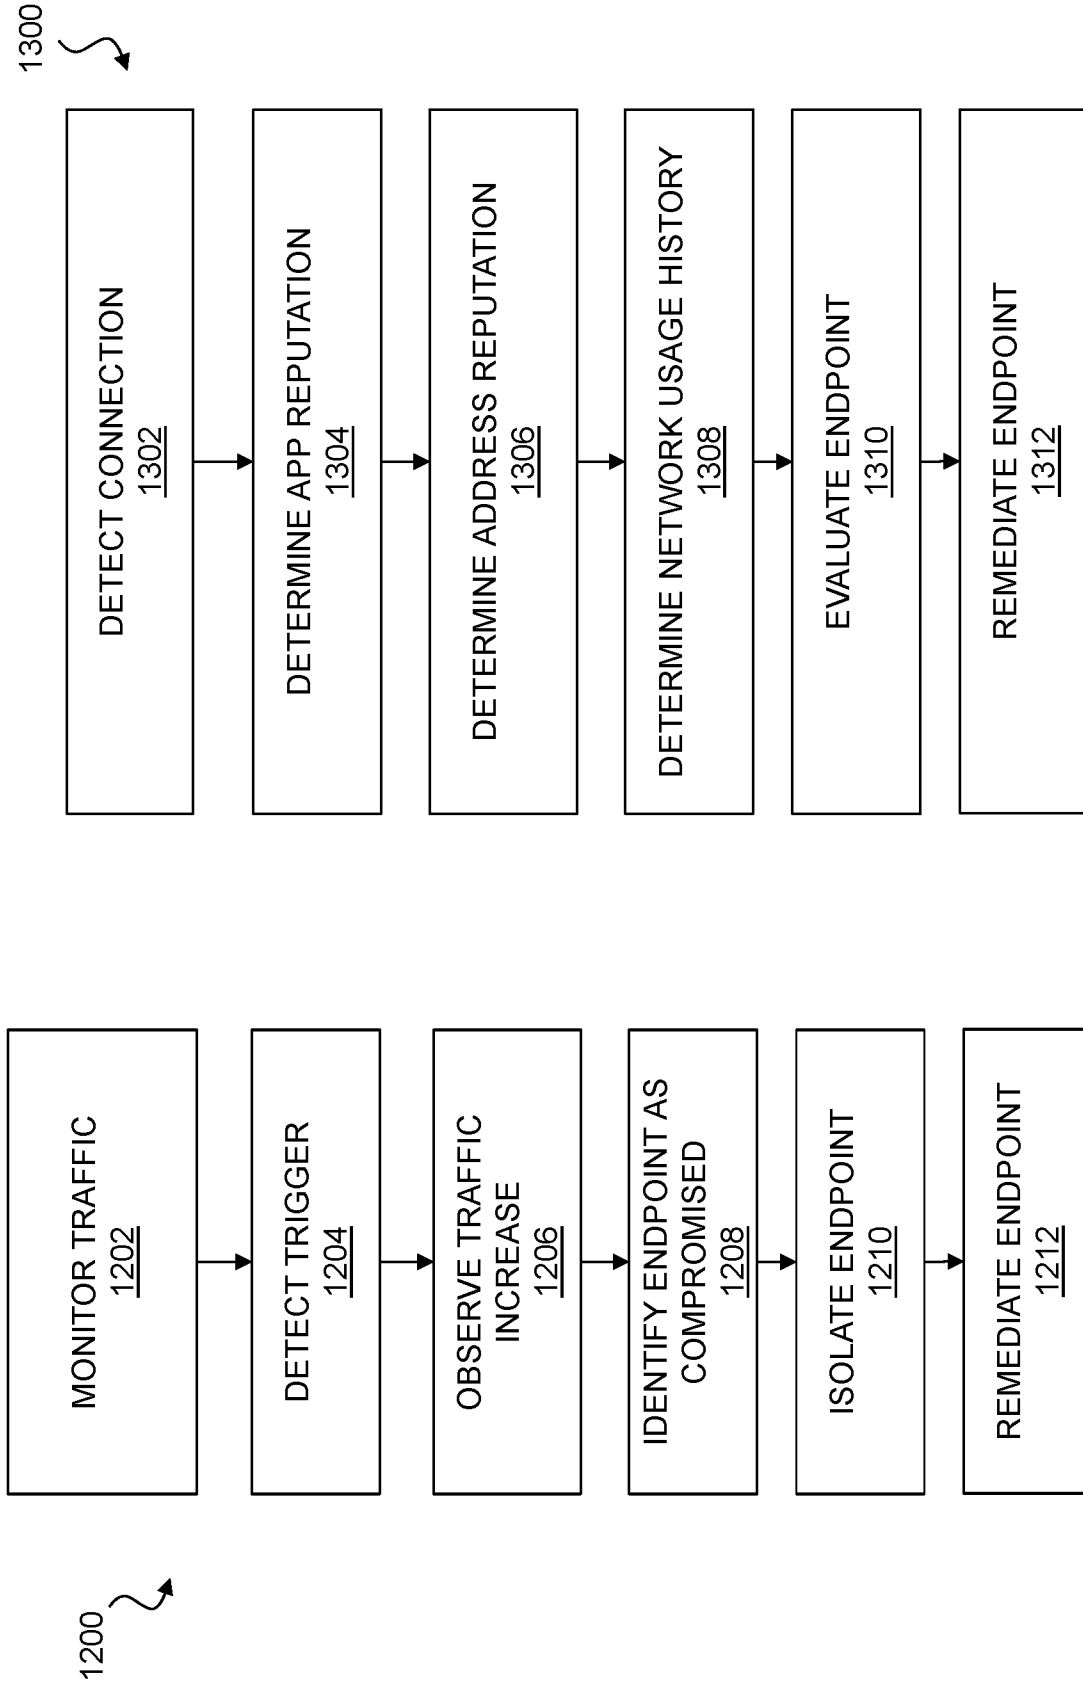


FIG. 12

FIG. 13

1400 ↗

<u>APP REPUTATION</u> 1402	<u>URL/IP REPUTATION</u> 1404	<u>HISTORY OF CONTROLLING APP</u> 1406	<u>INSIGHT/CONCLUSION WHEN A NEW IP/URL CONNECTION IS DETECTED</u> 1408
KNOWN GOOD	KNOWN MALICIOUS	CONSTRAINED URL USAGE (NEW URL NOT IN LIST)	URL IS OUTSIDE THE APP HISTORIC USAGE - THE "GOOD" APPLICATION WOULD APPEAR TO HAVE BEEN EXPLOITED
KNOWN GOOD	SUSPECT	CONSTRAINED URL USAGE (NEW URL NOT IN LIST)	THE NEW URL/IP USAGE COMBINED WITH AN ALREADY SUSPECT URL/IP INDICATES AN EXPLOIT AND PROBABLY CONFIRMED MALICIOUS URL/IP
SUSPECT	UNKNOWN	CONSTRAINED URL USAGE (NEW URL NOT IN LIST)	CONTRIBUTES TO SUSPICION, POSSIBLE EXPLOIT
KNOWN GOOD	UNKNOWN	CONSTRAINED URL USAGE (NEW URL NOT IN LIST)	GOOD APP MAY BE GOING TO A NEW LOCATION, BUT THIS LOOKS SUSPECT
KNOWN GOOD	KNOWN GOOD	CONSTRAINED URL USAGE (NEW URL NOT IN LIST)	STRANGE ACTIVITY LOOKS LIKE A GOOD APP IS GOING TO A NEW GOOD LOCATION.

1410 ↗


1420 ↗

1430 ↗

1440 ↗

1450 ↗

FIG. 14



*FIG. 15*

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/040094

A. CLASSIFICATION OF SUBJECT MATTER  
 IPC(8) - G06F 11/30, G06F 15/16, G06F 21/00, H04L 9/32, H04L 12/24, H04L 12/26 (2016.01)  
 CPC - G06F 21/316, G06F 21/32, G06F 21/566, G06F 2221/2133, H04L 41/0813, H04L 41/0816 (2016.05)  
 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)  
 IPC - G06F 11/30, G06F 15/16, G06F 21/00, H04L 9/32, H04L 12/24, H04L 12/26, H04L 29/06  
 CPC - G06F 21/316, G06F 21/32, G06F 21/566, G06F 2221/2133, H04L 41/0813, H04L 41/0816, H04L 41/0873, H04L 41/142  
 H04L 43/028, H04L 43/10, H04L 43/16, H04L 63/1416, H04L 63/145, H04L 63/1458, H04L 67/2852

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
 USPC - 713/171, 713/176, 726/22, 709/221, 709/224, 726/24, 726/3 (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
 Orbit, Google Patents, Google Scholar, Google  
 Search terms used: computer program product, protect against distributed denial of service attacks, enterprise network

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2008/0168559 A1 (TOUITOU et al) 10 July 2008 (10.07.2008), entire document	1-20, 41-60
Y	US 2007/0011740 A1 (DAVIS et al) 11 January 2007 (11.01.2007), entire document	1-20
Y	US 8,490,190 B1 (HERNACKI et al) 16 July 2013 (16.07.2013), entire document	1-20
Y	US 8,832,835 B1 (CHEN et al) 09 September 2014 (09.09.2014), entire document	2-4, 10-12, 41-60
Y	US 2016/0080417 A1 (SOPHOS LIMITED) 17 March 2016 (17.03.2016), entire document	21-40
Y	US 2010/0005291 A1 (HULTEN et al) 07 January 2010 (07.01.2010), entire document	21-40, 48, 49, 57, 58
A	US 2008/0244074 A1 (BACCAS et al) 02 October 2008 (02.10.2008), entire document	1-60
A	US 2015/0127790 A1 (HARRIS CORPORATION) 07 May 2015 (07.05.2015), entire document	1-60
A	US 2015/0215187 A1 (TUONONEN et al) 30 July 2015 (30.07.2015), entire document	1-60
A	US 2011/0320816 A1 (YAO et al) 29 December 2011 (29.12.2011), entire document	1-60

Further documents are listed in the continuation of Box C.  See patent family annex.

\* Special categories of cited documents:  
 "A" document defining the general state of the art which is not considered to be of particular relevance  
 "E" earlier application or patent but published on or after the international filing date  
 "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)  
 "O" document referring to an oral disclosure, use, exhibition or other means  
 "P" document published prior to the international filing date but later than the priority date claimed  
 "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention  
 "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone  
 "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art  
 "&" document member of the same patent family

Date of the actual completion of the international search 29 August 2016	Date of mailing of the international search report <b>08 SEP 2016</b>
---	--

Name and mailing address of the ISA/ Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, VA 22313-1450 Facsimile No. 571-273-8300	Authorized officer Blaine R. Copenheaver PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774
---	--