



(51) International Patent Classification:

G06F 21/20 (2006.01) *G06F 15/16* (2006.01)
G06F 21/22 (2006.01) *G06F 17/00* (2006.01)

(21) International Application Number:

PCT/US2009/067790

(22) International Filing Date:

11 December 2009 (11.12.2009)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

12/361,548 29 January 2009 (29.01.2009) US

(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventor: **LEE, Henry**; One Microsoft Way, Redmond, Washington 98052-6399 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP,

KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

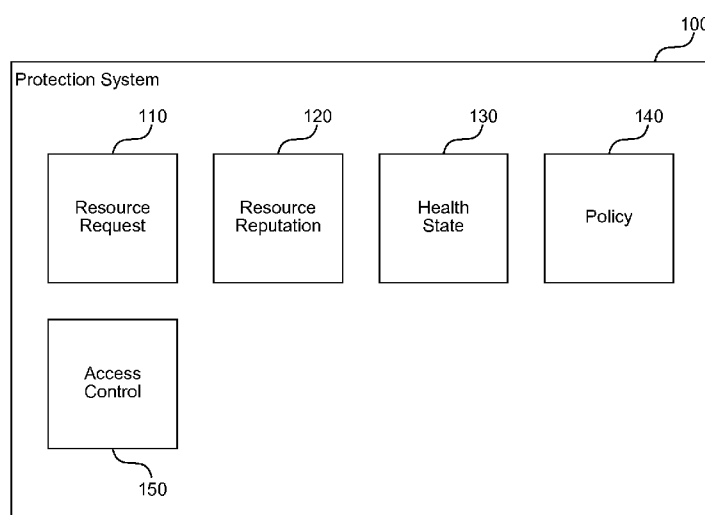
Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: HEALTH-BASED ACCESS TO NETWORK RESOURCES



(57) Abstract: A protection system is described herein that dynamically determines whether a computer system can access a particular resource based on a combination of a dynamic health state of the computer system and a dynamic reputation of the resource. When a user attempts to access a resource, the protection system intercepts the request. The protection system determines the reputation of the resource that the user is attempting to access and the health of the computer system through which the user is attempting to access the resource. Based on the determined resource reputation and the determined computer system health, the protection system determines whether to allow the requested access to the resource.

FIG. 1

HEALTH-BASED ACCESS TO NETWORK RESOURCES

BACKGROUND

[0001] Antivirus, antispyware, and other anti-malware applications seek to protect client computers by identifying harmful applications or other executable code and removing or at least neutralizing the harmful code. Current anti-malware applications (e.g., Microsoft Windows Defender, Microsoft Forefront Client Security, Microsoft OneCare, Microsoft Forefront Server for Exchange Server, and so forth) use a signature-based approach to detect viruses, worms, and spyware. One common type of malware prevention relies on inspecting the Uniform Resource Locators (URLs) or Internet Protocol (IP) addresses that a user requests (e.g., in a web browser) and blocking access to URLs that have been flagged as malicious or potentially malicious. Often, client computers are infected when a user visits a website and allows installation of “unknown software” (typically, users believe they are installing good software), so blocking access to sites known to be malicious can protect the user’s computer system from infection.

[0002] Machine health state is a technology used to determine a computer system’s health and determine the actions the computer system can perform based on the health of the computer system. For example, many corporate local area network (LAN) administrators define policies that each computer system meets before the policy allows a computer system to access the corporate network. For example, the administrator may define a policy that specifies a particular operating system patch or a particular antivirus definition version that is present on computer systems in the organization before the policy allows the computer systems to access the LAN.

[0003] Current web filtering technologies, either host-based (local machine) or edge-based (gateway device/server), use policies to control access to networking protocols or destinations. These policies typically use machine or user attributes known at policy authoring time, such as testing membership in a group, testing the destination or source site name, or imposing time of day restrictions for access to certain network resources.

[0004] One problem is that these protection mechanisms are often unnecessarily restrictive. Websites that allow users to create content may have many non-

malicious areas of the website in addition to several malicious areas of the website. In addition, a user's job may entail visiting harmful websites (e.g., to identify illegal distribution of copyrighted material for the corporation that employs the user), and the user may be taking other precautions to prevent the risk of infecting the user's computer system. For example, the user may run the browser in a protected or sandboxed mode to prevent websites from having an effect on other elements of the computer system. The user may also run stringent antivirus software to reduce the risk of infection. In the case of machine health, a user that rarely uses a laptop may take the laptop to a presentation and want to access a website but be prevented from doing so because the user has not recently updated the laptop with patches.

[0005] At other times, traditional protection mechanisms may not be restrictive enough. For example, URL-based blocking is only as good as the list of known malicious URLs. Malware authors constantly change the domain names that host malware and thus URL-based blocking may fail to identify a malicious website for a certain period until an administrator adds that website to a list of malicious websites.

SUMMARY

[0006] A protection system is described herein that determines whether a computer system can access a particular resource based on a combination of a health state of the computer system and a reputation of the resource. When a user attempts to access a resource, the protection system intercepts the request. The protection system determines the reputation of the resource that the user is attempting to access. The protection system also determines the health of the computer system through which the user is attempting to access the resource. Based on the determined resource reputation and the determined computer system health, the protection system determines whether to allow the requested access to the resource. Thus, the protection system allows an administrator to protect a computer system from potentially malicious resources based on the health of the computer system. The protection is dynamic because both the health of the computer system and the reputation of a particular resource may change over time, and the protection system protects the computer system based on the reputation and health at the time of the request.

[0007] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Figure 1 is a block diagram that illustrates components of the protection system, in one embodiment.

[0009] Figure 2 is a block diagram that illustrates an environment in which the protection system is employed, in one embodiment.

[0010] Figure 3 is a flow diagram that illustrates the processing of the system to handle a received request to access a resource, in one embodiment.

[0011] Figure 4 is a flow diagram that illustrates the processing of the access control component to allow or deny a request to access a resource, in one embodiment.

DETAILED DESCRIPTION

[0012] A protection system is described herein that determines whether a computer system can access a particular resource based on a combination of a health state of the computer system and a reputation of the resource. The system will allow a healthy machine to access a broader range of resources (e.g., less reputable) than an unhealthy machine. For example, a healthy machine can be one that has antivirus software installed with all of the latest virus definitions up to date. Such a machine is much less likely to be in danger from website content than a machine without antivirus software installed. The system may allow an unhealthy machine to access only resources that are known to be highly reputable.

[0013] When a user attempts to access a resource, the protection system intercepts the request. For example, a user may use a web browser to request a particular web page by specifying a URL. The protection system determines the reputation of the resource that the user is attempting to access. For example, if the resource is a website, then the system may send the URL to a trusted URL reputation service and receive a response that indicates a reputation score for the resource. The protection system determines the health of the computer system through which the user is attempting to access the resource. For example, the

protection system may check a Network Access Protection (NAP) state of the system (sometimes referred to as Network Access Control (NAC)) or run one or more health checks to determine a health score for the computer system. Based on the determined resource reputation and the determined computer system health, the protection system determines whether to allow the requested access to the resource. Thus, the protection system allows an administrator to protect a computer system from potentially malicious resources based on the health of the computer system. The protection is dynamic because both the health of the computer system and the reputation of a particular resource may change over time, and the protection system protects the computer system based on the reputation and health at the time of the request.

Health

[0014] As noted herein, the protection system determines the health of a computer system so that the protection system can apply differing levels of restrictions based on the health. There are many factors and processes that the protection system can use to determine a computer system's health or to define health checks. Some examples are described in further detail herein, and those of ordinary skill in the art will recognize other examples for determining health that can be used with the protection system. The health checks can include any measurement of a computer system that determines the computer system's level of defense to malicious content.

[0015] A health score as described herein can refer to a broad range of health indications. For example, a simple implementation of a health score includes a Boolean termination that the system is healthy or unhealthy. A more complex health score may have levels of health, such as low health, medium health, and high health. Alternatively, the health score may include a numerical range, such as 0-100 where one end of the range indicates system failure of all health policies defined by the administrator and the other end of the range indicates a system passing all health policies. In each case, the protection system allows the administrator to define a level of access to resources independently for each defined increment of the health score. The protection system considers the health score, the resource reputation score described further herein, and any defined policy to determine which resources a particular computer system can access.

[0016] Applications can implement health checks that the protection system invokes to determine a computer system's current health state. For example, products such as Windows Security Center, anti-virus products, patch management products, and other security products, or a combination of these products may provide health checks to the system. The system may also perform health checks at different times, such as by accessing a client health NAP agent, in-band by calling a NAP or other health-based server, or out of band by relying on a health certificate that certifies an earlier-determined health state. In addition to machine health (e.g., as provided by NAP), the protection system can also take user health into account when making resource access decisions. For example, in Microsoft Forefront Client's Security State Assessment, some checks are per-machine, and some checks are per-user. For per-user checks, like IE zone settings or Office macro settings, unknown resources can be blocked for a user, unless both machine-based and user-based health checks pass.

Resource Reputation

[0017] As noted herein, the protection system determines the reputation score of a particular resource (e.g., a URL) so that access to resources with lower reputations can be blocked or limited for computer systems with lower health. The protection system can use several methods to determine the reputation of a resource, which are described in further detail herein. Those of ordinary skill in the art will recognize that the protection system can be used with other resource reputation methods in addition to those described herein.

[0018] As with health score, the resource reputation score described herein can refer to a broad range of reputation indications. For example, a simple implementation of a reputation score includes a Boolean termination that a resource is malicious or safe. A more complex reputation score may have levels of reputation, such as low reputation, medium reputation, and high reputation. Alternatively, the reputation score may include a numerical range, such as 0-100 where one end of the range indicates known poor reputation and the other end of the range indicates a known high reputation. In each case, the protection system allows the administrator to define a level of access to resources independently for each defined increment of the reputation score. The protection system considers the health score described further herein, the resource reputation score, and any

defined policy to determine which resources a particular computer system can access.

[0019] In some embodiments, the protection system defines a resource reputation score that corresponds to an unknown resource the reputation of which has not been evaluated or is unavailable. For example, new websites may take some time to be discovered and for a resource reputation service to evaluate. During this time, it is still useful for the protection system to know that the resource is unknown. The protection system allows the administrator to define policies that determine whether a computer system can access a resource for which reputation information is unavailable. When the unknown state is not a Boolean, but rather a gradient level (e.g., from 0 to 100) the administrator can potentially customize the unknown state through policy. For example, one organization's administrator may feel a score greater than 50 is dangerous for unknown sites, while another organization's administrator may feel a score greater than 75 is dangerous. Thus, a cautious administrator can deny access to unknown resources or specify a particular minimum health state of computer systems that request access to unknown resources. For example, if a client system has antivirus real-time protection enabled and signatures and patches are up-to-date, the administrator may define a policy that specifies that the client system can browse unknown URLs. Previous systems typically block unknown URLs or allow access based on static criteria used in traditional edge-firewall rules.

[0020] In some embodiments, the reputation information that the protection system receives from a resource reputation service includes one or more categories that describe the type of content provided by the resource. For example, a website may belong to a category such as sports, news, pornography, and so forth. The protection system allows the administrator to define the categories that a client system can access and associate a threshold health state that the client system achieves to access resources of a particular category. In some embodiments, categories are assigned a confidence level, such as from 0 to 100, depending on how granular the data is returned from the reputation service. This allows policies based on "probable" category matches for dynamically rated content. For example, a policy can specify blocking URLs if the confidence of pornographic content is greater than 80.

System Components

[0021] Figure 1 is a block diagram that illustrates components of the protection system, in one embodiment. The protection system 100 includes a resource request component 110, a resource reputation component 120, a health state component 130, a policy component 140, and an access control component 150. Each of these components is described in further detail herein.

[0022] The resource request component 110 receives requests to access resources from user programs or other applications (e.g., an operating system service). In some embodiments, the protection system 100 employs a kernel-mode network filter to intercept network requests and perform the steps described herein for protecting the computer system. Many operating systems provide an application-programming interface (API) for adding networking filters into the networking stack. For example, Microsoft Windows provides APIs for adding several types of Network Driver Interface Specification (NDIS) filters that can perform the steps described herein. Alternatively or additionally, the protection system 100 can operate at a resource in between the requesting computer system and the requested resource, such as at a network router. By interfacing with the computer system at a level below any particular application, the protection system 100 can provide protection to the computer system regardless of the application a user uses to access a network (e.g., web browser, instant messaging application, email client, and so forth). The request received by the resource request component 110 identifies the resource to be accessed, such as by IP address, domain name, URL, or similar identification.

[0023] The resource reputation component 120 determines a reputation score of the requested resource for resource requests received through the resource request component 110. For example, the resource request component 110 may receive a URL and the resource reputation component 120 may send the URL to a resource reputation service hosted on the Internet to determine the reputation of the resource. The reputation service provides a response that indicates the reputation of the requested resource, or in some cases a response that indicates that the reputation of the requested resource is not known by the reputation service.

[0024] The health state component 130 determines a health score of the computer system making the resource request. The health score may include the results or

an aggregation of the results of administrator-defined health policies. Health policies can include a variety of types of information that the administrator defines as indicating computer system health, such as operating system patch level, virus definition level, whether drivers are signed, applications installed on the computer system, whether patches are up to date for each application, and so forth. The health state of a system may change over time based on user actions (e.g., installing applications), external factors (e.g., an antivirus software vendor releasing new virus definitions), and/or based on changes in policy defined by the administrator. Thus, the protection system 100 determines whether to allow or deny a particular access request based on the current health state of the computer system.

[0025] The policy component 140 receives and stores policies defined by an administrator that the protection system 100 uses to determine whether to allow or deny a particular access request. For example, the administrator may define a policy that allows access to high reputation resources by computer systems of any health level, access to medium reputation resources by computer systems of medium health level, and access to low reputation resources by computer systems of high health level. The administrator may also define a policy for resources for which reputation information is not available. Policies defined by the administrator may also include modifying the qualities of a computer system that determine the system's health. Policies can also include enabling an auto-remediation and/or auditing feature, as described further herein.

[0026] Although an administrator is described herein, typically acting in a large corporate setting, the administrator may also be a home user such as a parent that defines one or more policies for the computer systems of children in a household. Accordingly, an administrator is defined not by the size or type of an organization in which computer systems are used, but rather by the actions that a user acting in an administrative role can perform with respect to the protection system 100.

[0027] The access control component 150 determines whether to allow access to a requested resource by applying the defined policies to the determined health score of the requesting computer system and reputation score of the resource to which the computer system requests access. If the access control component 150 allows access to the resource, then the component 150 sends the received

request through the network stack as if the protection system 100 had not intercepted the request. The resource may then receive the request and provide an appropriate response (e.g., based on a protocol associated with the request).

[0028] In some embodiments, If the access control component 150 denies access

5 to a resource, the component 150 provides a response to the request (e.g., through the resource request component 110) that indicates failure to access the resource. For example, if the request is a Hypertext Transfer Protocol (HTTP) request received from a web browser, the access control component 150 may provide a 404 (resource not found) or other appropriate HTTP error code to the
10 web browser to inform the user that the system denied the request (and to prevent the application from waiting for a timeout and/or retrying repeatedly).

[0029] The computing device on which the system is implemented may include a central processing unit, memory, input devices (e.g., keyboard and pointing devices), output devices (e.g., display devices), and storage devices (e.g., disk
15 drives or other non-volatile storage media). The memory and storage devices are computer-readable storage media that may be encoded with computer-executable instructions (e.g., software) that implement or enable the system. In addition, the data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communication link. Various
20 communication links may be used, such as the Internet, a local area network, a wide area network, a point-to-point dial-up connection, a cell phone network, and so on.

[0030] Embodiments of the system may be implemented in various operating environments that include personal computers, server computers, handheld or
25 laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, digital cameras, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and so on. The computer systems may be cell phones, personal digital assistants, smart phones, personal
30 computers, programmable consumer electronics, digital cameras, and so on.

[0031] The system may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, and so on that perform particular

tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments.

[0032] Figure 2 is a block diagram that illustrates an environment in which the protection system is employed, in one embodiment. This diagram depicts host-based NAP enabled web filtering, where there is no edge gateway device. One or more client computer systems 210 communicate with three set of hosted servers: a URL reputation service 220 (shown as Microsoft Reputation Service Array and MRS Database), machine/user administration servers 230 for policy and telemetry (shown as Root Silo, Servicing Silo), and hosted NAP servers 240 (shown as NAP HRA, NAP CA, NAP NPS). Each of the computer system communicate via a network 250, such as one or more of the Internet, a LAN, a WAN, a wireless network, and so forth.

[0033] The NAP servers 240 provide a statement-of-health (in some cases cryptographically signed to prevent tampering) that the NAP servers 240 send to a NAP Agent on the client machines. The enforcement policies of the protection system can refer to the statement-of-health, cached by the NAP Agent, and can be used in combination with policies containing URL reputation categories, user group membership, or time-of-day restrictions to determine whether one of the client computer systems 210 can access a requested resource. For client machines that do not support a NAP Agent or are exempt from NAP checks, a NAP exemption certificate can be installed, which the protection system can use in evaluating resource policies, just like on NAP enabled clients.

[0034] The URL reputation service 220 may be provided by a third party for a variety of purposes (e.g., phishing filters, web access control, and so on). For example, Microsoft Live OneCare includes a URL reputation service 220 accessible via the Internet to which client computer systems 210 can send reputation requests. Similarly, NAP servers 240 may be hosted by a health provider that ensures the health of the client computer systems 210 and provides the statement of health as validation of the health. The protection system can be configured to provide its own independent URL reputation service and health servers or can uses third party services such as those described to perform the processes described herein.

[0035] Figure 3 is a flow diagram that illustrates the processing of the system to handle a received request to access a resource, in one embodiment. In block 310, the system receives a request to access a resource, wherein the request includes an identification of the resource. When a user attempts to access a resource, the protection system intercepts the request. For example, a user may use a web browser to request a particular web page by specifying a URL. Continuing in block 320, the system determines a reputation of the requested resource. For example, if the resource is a website, then the system may send the URL to a trusted URL reputation service and receive a response that indicates a reputation score for the resource. The reputation indicates a likelihood that accessing the resource will cause malicious content to be downloaded to or will otherwise negatively influence the computer system.

[0036] Continuing in block 330, the system determines the health of the computer system from which the request is received. For example, the protection system may check a network access protection (NAP) state of the system or run one or more health checks to determine a health score for the computer system.

Continuing in block 340, the system accesses an access policy that identifies one or more conditions based on the resource reputation and computer system health for which the computer system is allowed to access the requested resource.

[0037] Continuing in block 350, the system applies the access policy to determine whether to allow the computer system to access the requested resource based on the access policy. This step is described further with reference to Figure 4. The health of the computer system and the reputation of a particular resource may change over time, and the protection system allows an administrator to dynamically determine whether a computer system can access a resource based on the computer system's health and the reputation of the resource at the time of the request. After block 350, these steps conclude.

[0038] Figure 4 is a flow diagram that illustrates the processing of the access control component to allow or deny a request to access a resource, in one embodiment. In the illustrated embodiment, the protection system receives a reputation score that indicates high, medium, or low reputation and a health state that indicates high, medium, or low health. As described herein, the reputation and/or health state used by the protection system may be more or less granular than that illustrated.

[0039] In block 410, the component receives the health and reputation scores determined by the system based on the request. For example, the system may determine the health and reputation scores based on the process described with reference to Figure 3. Continuing in decision block 420, if the reputation of the requested resource is high (e.g., above a predetermined threshold), then the component continues at block 425 and allows the computer system to access the requested resource, else the component continues at block 430. For example, the requested resource may be of such high reputation (e.g., <http://www.microsoft.com/>) that the administrator-defined policy allows access to the resource regardless of the current health of the computer system.

[0040] Continuing in decision block 430, if the health of the requesting computer system is high, then the component continues at block 425 and allows the computer system to access the requested resource, else the component continues at block 440. For example, the administrator-defined health policies may be restrictive enough (e.g., all patches installed, browser in protected mode, JavaScript off, virus definitions up to date, and so forth) that the administrator is confident that high health machines will be safe regardless of the requested resource's reputation. Continuing in decision block 440, if the health of the requesting computer system is medium, then the component continues at block 450, else the component continues at block 460 and denies the computer system access to the requested resource.

[0041] Continuing in decision block 450, if the reputation of the requested resource is medium, then the component continues at block 425 and allows the computer system to access the requested resource, else the component continues in block 460 and denies the computer system access to the requested resource. Continuing in block 460, the component denies the computer system access to the requested resource. After block 460, these steps conclude.

[0042] In some embodiments, the protection system provides automated remediation and/or auditing of access to resources. For example, when denying a user access to a website, the protection system may display a user interface to the user information the user that the user can access the resource if the user updates his antivirus software or applies recommended operating system patches. The system may automatically apply such remedies without asking the user. In addition, when the protection system denies a user access to a resource, the

system may create an audit trail (e.g., a log) that an administrator can review to determine, for example, whether the policy could allow certain unknown sites or whether the system is denying access to legitimate sites. This allows a wide degree of security policy visibility through reporting, security compliance through blocking web access based on NAP health checks, and security remediation through NAP auto-remediation features. Instead of blocking web access first due to failed NAP health checks, the system can deny access as a last resort, after attempting NAP auto-remediation actions. For example, auto-remediation can include locking down firewall configuration, installing security patches, installing antivirus signature updates, and enabling security features automatically before access to a resource is denied.

[0043] In some embodiments, the protection system works in combination with, rather than in place of, traditional filtering techniques. For example, the system can use health checks in combination with user group membership, URL reputation category, or time-of-day restrictions to determine whether to allow access to a particular resource. The protection system may provide less restrictive access to users of certain groups (e.g., administrators or power users) than to users of other groups (e.g., guests) based on the expected sophistication of users in each group or other criteria.

[0044] In some embodiments, the protection system is provided as an online hosted service that protects client machines without substantial server infrastructure in an organization. For example, an organization may purchase a subscription to an online service through which an administrator can define access policies and to which clients send access requests and receive an allow/deny response without employing traditional firewalls, proxy servers, or other infrastructure within the organization. This can relieve the administrator of a substantial burden and save the organization on hardware costs.

[0045] In some embodiments, the protection system allows third parties to extend the system by adding additional health checks or defining new types of resource requests (e.g., by network protocol) that the system can monitor. For example, an application such as Apple iTunes may provide a plug-in to the system that informs the system when the application has all current patches installed. The protection system can use this information as a factor in determining the health score of a client system hosting the application.

[0046] From the foregoing, it will be appreciated that specific embodiments of the system have been described herein for purposes of illustration, but that various modifications may be made without deviating from the spirit and scope of the invention. For example, although web sites have been described as one type of resource, the protection system can be used with many types of resources including network shares, ftp sites, RSS feeds, video playlists, and so forth. Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

I/We claim:

1. A computer-implemented method for protecting a computer system from accessing malicious network resources, the method comprising:
 - 5 receiving 310 a request to access a resource, wherein the request includes an identification of the resource;
determining 320 a reputation of the requested resource, wherein the reputation indicates a likelihood that accessing the resource will cause malicious content to be downloaded to the computer system;
 - 10 determining 330 a health state of the computer system from which the request is received;
accessing 340 an access policy that identifies one or more conditions based on the resource reputation and computer system health state for which the computer system is allowed to access the requested resource; and
 - 15 applying 350 the access policy to determine whether to allow the computer system to access the requested resource based on the access policy and thereby protect the computer system from accessing malicious network resources, wherein the steps are performed by at least one processor.
2. The method of claim 1 wherein receiving a request comprises using a
20 network filter installed on the computer system to intercept the request from an application.
3. The method of claim 1 wherein receiving the request comprises receiving requests from multiple applications and protecting the computer system from accessing malicious network resources from the multiple applications.
- 25 4. The method of claim 1 wherein determining the reputation of the requested resource comprises sending the identification of the requested resource to a web-based reputation service and receiving from the reputation service a response that indicates a reputation score for the requested resource.
5. The method of claim 1 wherein determining the health state comprises
30 accessing a Network Access Protection (NAP) server to determine the health state of the computer system.
6. The method of claim 1 wherein the access policy provides more access if the computer system has a higher greater health state.

7. The method of claim 1 wherein the health state of the computer system and the reputation of the requested resource change over time, and the method dynamically determines whether to allow the computer system to access the requested resource based on the computer system's health state and the reputation of the resource at a time of the request.

8. A computer system for blocking access to a network resource based on the health of the computer requesting access to the network resource, the system comprising:

a processor and memory configured to execute software instructions;

a resource request component 110 configured to receive requests to access resources from user programs or other applications, wherein requests received by the resource request component identify the resource to be accessed;

a resource reputation component 120 configured to determine a reputation score of the network resource for resource requests received through the

resource request component;

a health state component 130 configured to determine a health score of the computer system making the resource request;

a policy component 140 configured to receive and store policies defined by an administrator to determine whether to allow or deny a particular resource

access request; and

an access control component 150 configured to determine whether to allow access to a requested resource by applying the defined policies to the determined health score of the requesting computer system and the reputation score of the network resource.

9. The system of claim 8 wherein the resource request component employs a kernel-mode network filter to intercept network requests through an operating system provided application-programming interface (API) for adding networking filters into a networking stack.

10. The system of claim 8 wherein the system operates at a resource in between a requesting computer system and the requested resource, to intercept resource requests and determine whether to allow the requesting computer system to access the requested resource.

11. The system of claim 8 wherein the resource reputation component is further configured to receive a Uniform Resource Locator (URL) and send the


URL to a URL reputation service hosted on the Internet to determine the reputation of the resource.

12. The system of claim 8 wherein the resource reputation component is further configured to indicate whether a reputation is known for the requested network resource and wherein the policy component is configured to apply a specific policy for resources without a known reputation.

13. The system of claim 8 wherein the health state component determines a health score based on an aggregation of the results of administrator-defined health policies that include one or more health checks selected from the group consisting of: an operating system patch level, a virus definition level, a driver signing state, an application installation state, an application patch state, and a firewall configuration.

14. The system of claim 8 wherein the policy component allows access to high reputation resources by computer systems of any health level and access to low reputation resources by computer systems of a high health level.

15. The system of claim 8 wherein the administrator includes a home user that defines policies for one or more computer systems in a household.

**FIG. 1**


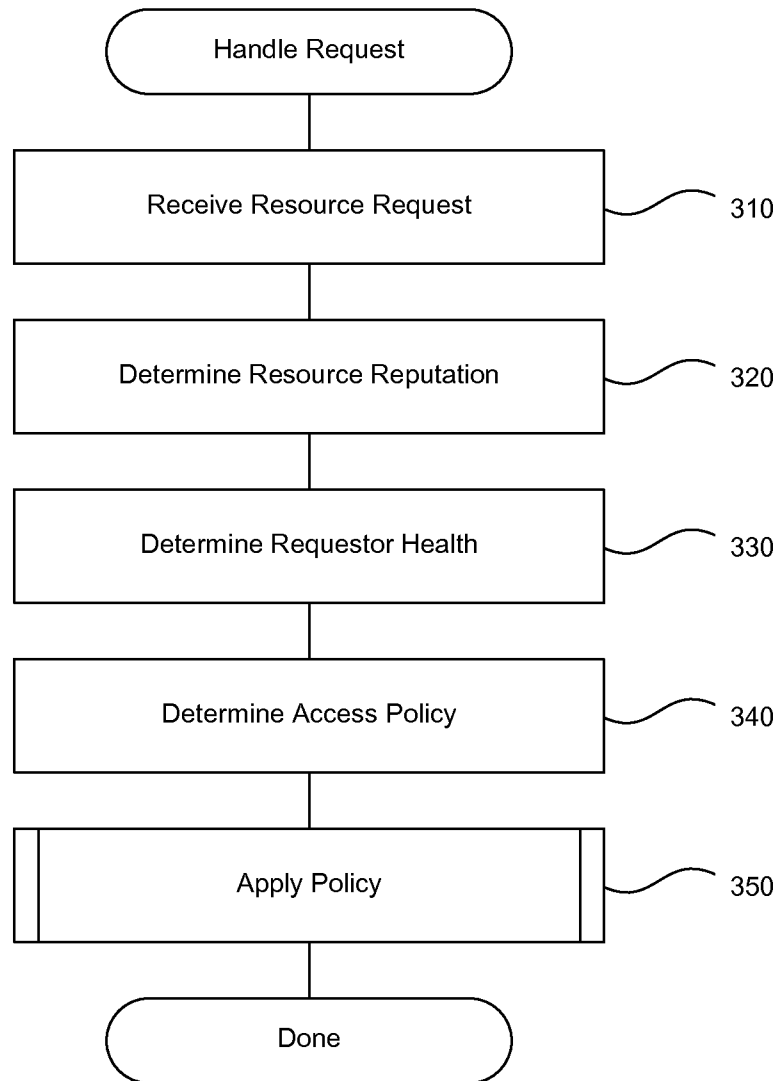
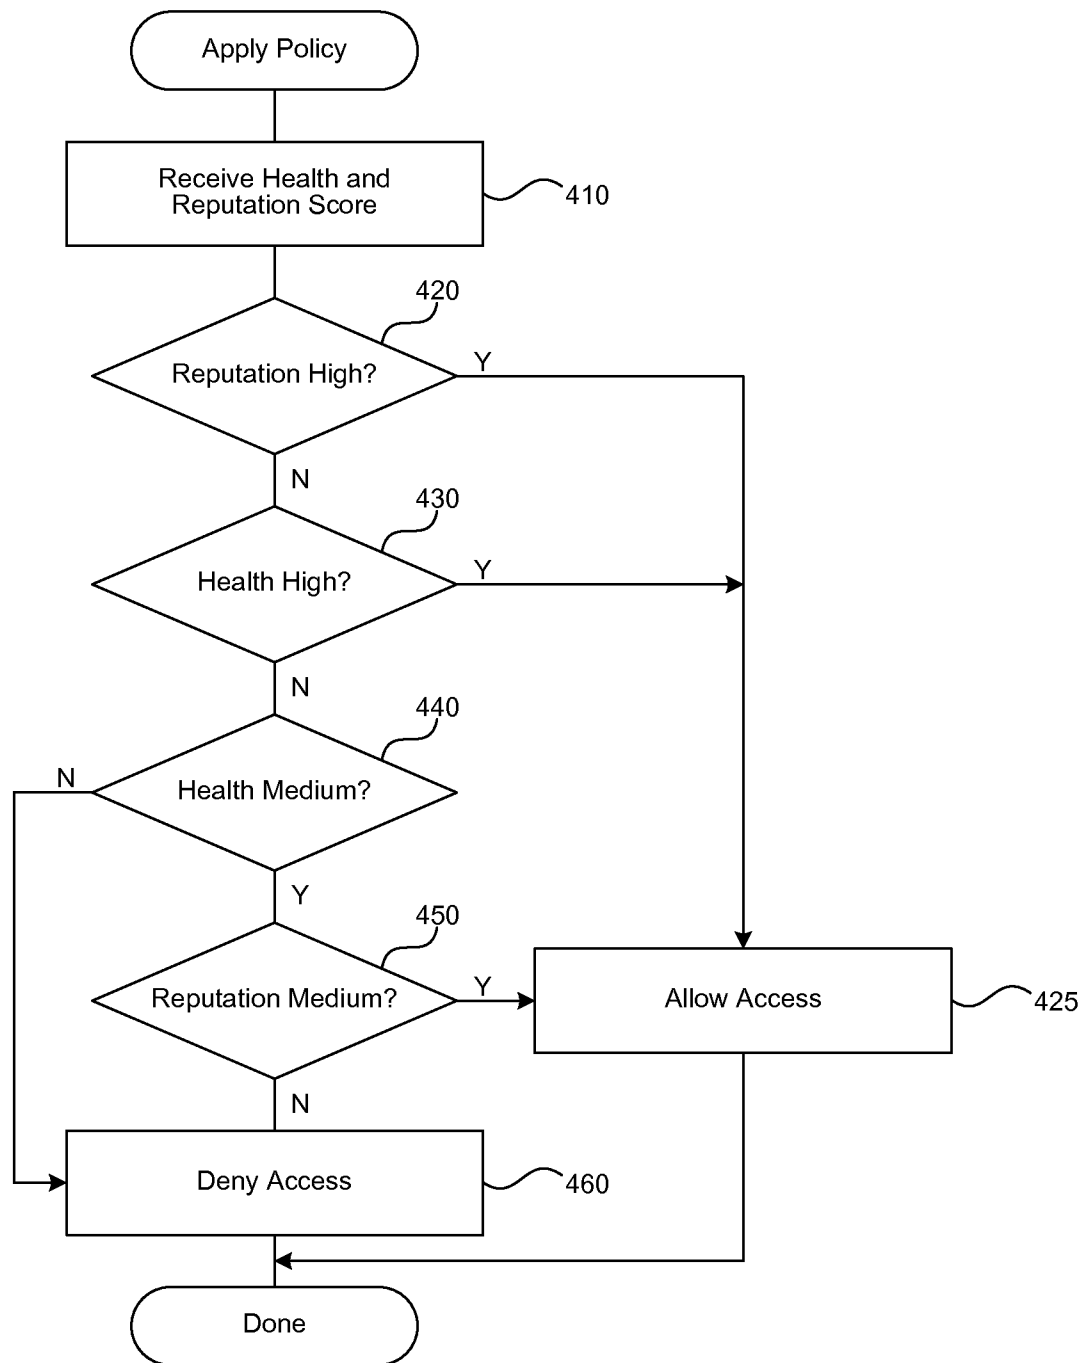


FIG. 2

3/4

**FIG. 3**

4/4

**FIG. 4**

A. CLASSIFICATION OF SUBJECT MATTER**G06F 21/20(2006.01)i, G06F 21/22(2006.01)i, G06F 15/16(2006.01)i, G06F 17/00(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 21/20; G06F 1/00; G06F 11/00; G06F 12/00; G06F 12/14; H04L 9/32; H04N 7/16

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: "file, access, protect, policy, malicious, reputation, likelihood"

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2002-0162013 A1 (RODNEY CARLTON BURNETT et al.) 31 October 2002 See abstract, paragraphs 25-33.	1-15
A	JP 2003-233521 A (HITACHI LTD) 22 August 2003 See abstract, claims 1-13.	1-15
A	US 2006-0120526 A1 (PETER BOUCHER et al.) 08 June 2006 See abstract, claim 1.	1-15
A	EP 1320016 A2 (PERVASIVE SECURITY SYSTEMS INC.) 18 June 2003 See abstract, claim 1.	1-15
A	JP 2006-120161 A (LUCENT TECHNOL INC) 11 May 2006 See the whole document.	1-15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of 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 JUNE 2010 (29.06.2010)

Date of mailing of the international search report

29 JUNE 2010 (29.06.2010)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
Government Complex-Daejeon, 139 Seonsa-ro, Seo-
gu, Daejeon 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

Soak, Sang Moon

Telephone No. 82-42-481-8470



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2009/067790

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2002-0162013 A1	31.10.2002	US 7260718 B2	21.08.2007
JP 2003-233521 A	22.08.2003	None	
US 2006-0120526 A1	08.06.2006	AU 2003-299729 A1	14.07.2004
		AU 2003-299729 A8	14.07.2004
		US 2004-0123150 A1	24.06.2004
		US 2004-0123153 A1	24.06.2004
		US 2005-0055578 A1	10.03.2005
		US 2006-0094400 A1	04.05.2006
		US 2008-0046965 A1	21.02.2008
		US 2008-0052395 A1	28.02.2008
		US 2008-0077971 A1	27.03.2008
		US 2008-0109679 A1	08.05.2008
		US 7308703 B2	11.12.2007
		US 7353533 B2	01.04.2008
		US 7478420 B2	13.01.2009
		US 7526800 B2	28.04.2009
		US 7636936 B2	22.12.2009
		WO 2004-057834 A2	08.07.2004
		WO 2004-057834 A3	08.07.2004
		WO 2006-076404 A2	20.07.2006
		WO 2006-076404 A3	20.07.2006
		WO 2006-076536 A2	20.07.2006
		WO 2006-076536 A3	20.07.2006
EP 1320016 A2	18.06.2003	EP 1320010 A2	18.06.2003
		EP 1320010 A3	25.06.2003
		EP 1320011 A2	18.06.2003
		EP 1320011 A3	26.11.2003
		EP 1320011 B1	07.03.2007
		EP 1320012 A2	18.06.2003
		EP 1320012 A3	09.07.2003
		EP 1320013 A2	18.06.2003
		EP 1320013 A3	08.06.2005
		EP 1320014 A2	18.06.2003
		EP 1320014 A3	01.06.2005
		EP 1320015 A2	18.06.2003
		EP 1320015 A3	27.04.2005
		EP 1320016 A3	27.04.2005
		EP 1320017 A2	18.06.2003
		EP 1320017 A3	15.06.2005
		EP 1320018 A2	18.06.2003
		EP 1320018 A3	06.07.2005
		EP 1324565 A1	02.07.2003
		EP 1326156 A2	09.07.2003
		EP 1326156 A3	23.11.2005
		EP 1326157 A2	09.07.2003
		EP 1326157 A3	10.12.2003

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2009/067790

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		JP 2003-218851 A	31.07.2003
		JP 2003-223353 A	08.08.2003
		JP 2003-228519 A	15.08.2003
		JP 2003-228520 A	15.08.2003
		JP 2003-242015 A	29.08.2003
		JP 2003-248658 A	05.09.2003
		US 2003-0108883 A1	12.06.2003
		US 2003-0110131 A1	12.06.2003
		US 2003-0110169 A1	12.06.2003
		US 2003-0110397 A1	12.06.2003
		US 2003-0120601 A1	26.06.2003
		US 2003-0120684 A1	26.06.2003
		US 2003-0154381 A1	14.08.2003
		US 2003-0217281 A1	20.11.2003
		US 2004-0103202 A1	27.05.2004
		US 2008-0034205 A1	07.02.2008
		US 6889210 B1	03.05.2005
		US 7178033 B1	13.02.2007
		US 7260555 B2	21.08.2007
		US 7380120 B1	27.05.2008
		US 7478418 B2	13.01.2009
		US 7562232 B2	14.07.2009
		US 7631184 B2	08.12.2009
		US 7681034 B1	16.03.2010
		WO 02-064840 A2	22.08.2002
		WO 02-064840 A3	22.08.2002
JP 2006-120161 A	11.05.2006	DE 602005011733 D1	29.01.2009
		EP 1650633 A1	26.04.2006
		EP 1650633 B1	17.12.2008
		US 2006-0090196 A1	27.04.2006