(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
18 September 2008 (18.09.2008)

PCT

(10) International Publication Number
WO 2008/113059 A1

(51) International Patent Classification:
G06F 21/00 (2006.01) *G06F 17/00* (2006.01)
G06F 15/16 (2006.01)

(21) International Application Number:
PCT/US2008/057164

(22) International Filing Date: 14 March 2008 (14.03.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/724,061 14 March 2007 (14.03.2007) US

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

(72) Inventors: **HUDIS, Efim**; c/o Microsoft Corporation, International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **HELMAN, Yair**; c/o Microsoft Corporation, International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **MALKA, Joseph**; c/o Microsoft Corporation, International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **BARASH, Uri**; c/o Microsoft Corporation, International Patents, 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, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, 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))*

[Continued on next page]

(54) **Title:** ENTERPRISE SECURITY ASSESSMENT SHARING

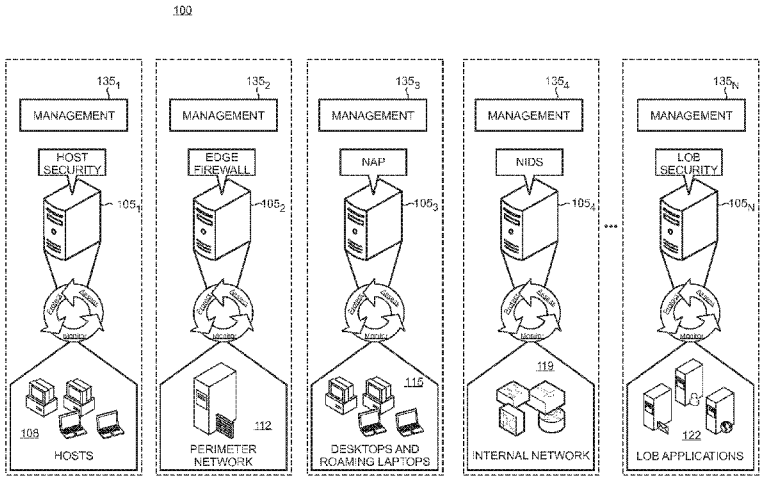


FIG. 1

(57) **Abstract:** An enterprise-wide sharing arrangement uses a semantic abstraction, called a security assessment, to share security-related information between different security products, called endpoints. A security assessment is defined as a tentative assignment by an endpoint of broader contextual meaning to information that is collected about an object of interest. Its tentative nature is reflected in two of its components: a fidelity field used to express the level of confidence in the assessment, and a time-to-live field for an estimated time period for which the assessment is valid. Endpoints may publish security assessments onto a security assessment channel, as well as subscribe to a subset of security assessments published by other endpoints. A specialized endpoint is coupled to the channel that performs as a centralized audit point by subscribing to all security assessments, logging the security assessments, and also logging the local actions taken by endpoints in response to security threats.

WO 2008/113059 A1

Published:
— *with international search report*
— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

ENTERPRISE SECURITY ASSESSMENT SHARING

BACKGROUND

[0001] In an enterprise computing environment, for example, an office of a business, a number of personal computers, workstations, servers and the like, along with other devices such as mass storage subsystems, internal network interfaces, and external network interfaces, are typically interconnected to provide an integrated environment in which information may be generated, accessed from external sources, and shared among various users. Commonly, users perform a variety of operations including order receipt, manufacturing, shipping, billing, inventory control, document preparation and management, e-mail, web browsing, and other operations in which creation, access, and sharing of data is beneficial.

[0002] Currently, security is typically provided for an enterprise using a variety of different security products that are each normally arranged to monitor only a partial portion of enterprise-wide data. That is, security products are arranged as separate local "islands" where each product monitors, assesses, and takes action with respect to different parts of the data within the enterprise. For example, an enterprise may utilize a combination of security products such as a product that protects host computers in the enterprise, an edge firewall product, a network intrusion detection system ("NIDS") product, a network access protection ("NAP") product, and other discrete security products in order to provide security for the various different parts of the enterprise.

[0003] While these security products often perform satisfactorily in many applications, detection of security incidents often suffers from undesirably high levels of false-positive and false-negative occurrences as a result of the monitoring of only partial enterprise security data. It has also been difficult to provide effective common management across all the enterprise security product islands. Current attempts to correlate enterprise-wide security data have high management and maintenance costs and have problems in scaling. More effective enterprise security management would be desirable to enable a single enterprise-wide view to enable security administrators to define and enforce clear, simple, and unified enterprise-wide policies for automatic responses to security incidents.

[0004] This Background is provided to introduce a brief context for the Summary and Detailed Description that follows. This Background is not intended to be an aid in determining the scope of the claimed subject matter nor be viewed as limiting the claimed

subject matter to implementations that solve any or all of the disadvantages or problems presented above.

SUMMARY

[0005] An enterprise-wide sharing arrangement called "ESAS – Enterprise Security Assessment Sharing" is provided in which a semantic abstraction, called a security assessment, is created to enable sharing of security-related information between different security products, called endpoints, in an enterprise security environment. A security assessment is defined as a tentative assignment by an endpoint of broader contextual meaning to information (i.e., data in some context) that is collected about an object of interest in the environment such as a computer, user, service (e.g., a website), data, or the enterprise as a whole. The security assessment utilizes a concise vocabulary for an endpoint to declare that an object in the environment falls into a particular assessment category such as "compromised" or "under attack" along with the severity (e.g., low, medium, high, critical) of the detected incident.

[0006] A security assessment is tentative because it is subject to some uncertainty and is valid for a limited period of time. The tentative nature of a security assessment is reflected in two of its components: a fidelity field which expresses the level of confidence the endpoint has in its assignment of contextual meaning, and a time-to-live ("TTL") field which reflects the endpoint's estimate of the time period for which the security assessment is expected to be valid. Thus, for example, a security assessment may be used by an endpoint to declare, in light of that endpoint's current understanding of one or more security incidents, that a particular machine is compromised, with a critical level of severity, with medium fidelity, and having a TTL of 30 minutes. A variety of types of security assessments may be used in any given enterprise security environment having, for example, various combinations of assessment category and object type.

[0007] Endpoints are enabled with functionality to publish security assessments onto a security assessment channel operating in the environment, as well as subscribe to a subset of available security assessments published by other endpoints. The security assessments existing in the environment that are active (i.e., those having a TTL which indicates the assessments are still valid) function to provide a security context that gives such ESAS-enabled endpoint with a new way to look at its own locally-available information. That is, the security context enables the ESAS-enabled endpoint to combine or correlate evidence from security assessments received from a variety of different sources, and across object

types, in order to significantly enhance the quality of its detection of potential security incidents. The ESAS-enabled endpoint then makes a decision as to what local action or response is appropriate for each type of security assessment (whether received from another endpoint or internally generated by the endpoint itself) in accordance with a set of response policies. Incident detection is both efficient and cost-effective because the security context enables distributed processing of enterprise-wide information, in the form of security assessments, without the burden of sharing large amounts of raw data throughout the enterprise (most of which is completely irrelevant due to the lack of any context). ESAS-enabled endpoints are further arranged to roll-back the local action upon expiration of the security assessment that prompted the local action (i.e., when the security assessment exceeds the time-to-live specified in the TTL field).

[0008] In an illustrative example, a specialized endpoint called an ESAS central server is coupled to the security assessment channel that performs as a centralized audit point by subscribing to all security assessments, logging the security assessments, and also logging the local actions taken by endpoints in response to security incidents in the environment. The ESAS central server provides administrators with a comprehensive view of the history and current status of the enterprise as a whole and of each ESAS-enabled endpoint. The utilization of the security assessments enables an administrator to compactly and efficiently configure response policies to incidents that are detected across the entire enterprise. The security assessments function as natural anchors, or starting points, to define enterprise-wide security response policies. A streamlined and consistent management interface is thus enabled to define the desired responses for each type of security assessment across the entire enterprise.

[0009] The present ESAS sharing arrangement provides a number of advantages. By employing a security assessment having a concise vocabulary, overall data complexity in the enterprise is drastically reduced and only meaningful information is shared between endpoints. Use of the security assessment also eliminates the need to collect large amounts of raw data in a central storage location, and thereby enables highly scalable enterprise security solutions to be built on a very cost effective basis. In addition, a new endpoint may be readily deployed with on-demand extensibility. Security assessments may be shared between the new endpoint and existing endpoints without the need to reconfigure any of the response policies within existing endpoints. The new endpoint simply functions as a new source of security assessments using a semantic abstraction that the existing endpoints already understand. The utilization of security assessments also enables enterprise-wide

security policies to be established using a very compact and clear methodology, without needing to understand all of the possible security events that every endpoint may generate in the enterprise, and then try to describe the responsive action for each event.

[0010] 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 as an aid in determining the scope of the claimed subject matter.

DESCRIPTION OF THE DRAWINGS

[0011] FIG 1 shows an illustrative enterprise security environment in which the present enterprise security assessment sharing may be implemented;

[0012] FIG 2 shows an illustrative enterprise security assessment sharing arrangement in which a channel is provided to enable a security assessment to be shared among multiple endpoints;

[0013] FIG 3 shows an illustrative terminology hierarchy that underlies a security assessment;

[0014] FIG 4 shows two illustrative endpoints and the complexity reduction enabled by the utilization of the present security assessment arrangement;

[0015] FIG 5 shows an illustrative example of functionality disposed in an endpoint that enables sharing of security assessments;

[0016] FIG 6 is a diagram of a first illustrative scenario in which a plurality of ESAS-enabled endpoints are coupled to a security assessment channel and a detected incident at one endpoint triggers responses at multiple other endpoints;

[0017] FIG 7 is a diagram of a second illustrative scenario in which a low fidelity security assessment is sent over the security assessment channel that triggers the generation of a new high fidelity assessment by a receiving ESAS-enabled endpoint which also performs cross-object mapping;

[0018] FIG 8 is a diagram of a third illustrative scenario that shows the targeted use of remediation techniques;

[0019] FIG 9 shows an illustrative screen that is provided by a graphical user interface ("GUI") which enables a user, such as an administrator, to manage and define the response policies of ESAS-enabled endpoints in the enterprise;

[0020] FIG 10 shows an illustrative screen that is provided by a GUI that is arranged to supplement, or be used as an alternative, to the GUI screen shown in FIG 9; and

[0021] FIG 11 shows an illustrative enterprise security arrangement in which the present ESAS feature set provides an enterprise security management layer functionality.

DETAILED DESCRIPTION

[0022] Analysis of current enterprise security solutions indicates that there are still significant opportunities for addressing customer needs. For example, each separate security product tends to have high rates of false positive and false negative detection of security incidents such as those produced through actions of malware or malicious users. Such low fidelity detection occurs because data from a single type of source (i.e., a subset of the enterprise-wide data) does not normally provide the context needed to make an accurate assessment of the security incident.

[0023] The use of automatic actions or responses is very infrequent as a result of the low fidelity detection since confidence in the validity of the detected incident is low. In addition, the typical response to a detected incident tends to be very harsh, for example, a user or machine may be disconnected from the network. Since such harsh actions generally impose significant costs to business activity in the enterprise, automation of such actions based on low fidelity detection is not generally performed.

[0024] Upon detection of an incident of interest, current security products typically perform investigation to determine the validity of the detection (i.e., whether the incident is true or false) and what action to take in response. Significant resources are expended on investigation to review the detailed data that is collected which may be relevant to the detected incident. Because it is not feasible to collect all data at all times, a security product collects only a subset of the available data through application of policies defined by an administrator. Such policies are often static and are commonly defined based on the storage capacity of the collection system, and not necessarily by the relevance of the incident data or the data source.

[0025] When an incident is detected, application of the policies typically results in a review of the data which triggered the detection. When this data is deemed insufficient to generate a high fidelity response, typically even more data is collected. For example, all of the data traffic into and out of a suspected compromised machine may be monitored. In many cases, a large amount of data is collected but is never used and has statistical significance only as noise. Consequently, many present security products collect an often overwhelming amount of noise, but not enough relevant data is collected.

[0026] Another area for improvement is the management and coordination of responses throughout the enterprise. Current enterprise security products inherently provide localized responses to incidents detected in the each separate island. Since the security products are isolated, the possible response options are limited to that part of the enterprise in which the particular security product operates. That is, actions and responses are capable of being defined in one security product island for separate incidents that are detected, but there is no ability to describe a desired action which may be more effective when it is applied in another part of the enterprise, or on a global basis. There is currently no single management point to enable enterprise-wide definition and enforcement of response policies to security incidents. Nor does a unified response channel and language / protocol exist by which each island can communicate to thereby notify the others that something has occurred or an action needs to be taken. The lack of management and coordinated responses results in significant costs being incurred for manual integration and correlation of data across the islands in the enterprise.

[0027] Turning now to the figures where like reference numerals indicate like elements, FIG 1 shows an illustrative enterprise security environment 100 in which a variety of security products 105 – 1, 2 … N, called endpoints, are deployed. It is emphasized that the number and type of endpoints 105 shown in FIG 1 are merely illustrative and the specific number of endpoints can be scaled up or down, and different types of security products/endpoints can be utilized, depending on the requirements of a specific application of enterprise security assessment sharing. For example, in addition to those shown in FIG 1 and described below, web application protection products, SEM / SIM (Security Event Management / Security Incident Management) products, operational heath monitoring and configuration management products (e.g., Microsoft Windows® Software Update Services, Microsoft Operations Manager), or identity management products (e.g., Microsoft Active Directory) are also usable in some applications.

[0028] In enterprise security environment 100, a host security endpoint 105_1 is deployed to protect, assess, and monitor a plurality of host computers 108 in the enterprise 100. A commercial example of the host security endpoint 105_1 is Microsoft Forefront Client Security® which provides unified malware protection for the enterprise's desktops, laptops, and server operating systems.

[0029] An edge firewall 105_2 is a security product that is arranged to protect the enterprise environment 100 from Internet-based threats while providing users with remote

access to applications and data through a perimeter network 112. Edge firewall 105_2 may be embodied by, for example, a Microsoft Internet Security and Acceleration® ("ISA") server.

[0030] A NAP security endpoint 105_3 performs computer health policy validation by ensuring ongoing compliance with health policies defined by an administrator. Typically, access is restricted for computers (e.g., desktops and roaming laptops 115) monitored by the NAP security endpoint 105_3 that do not comply with system health requirements.

[0031] A NIDS security endpoint 105_4 analyzes traffic inside the enterprise 100 over an internal network 119. The NIDS security endpoint 105_4 operates to detect malicious activity such as denial of service attacks port scans by monitoring network traffic on the internal network 119.

[0032] A line-of-business security endpoint 105_N protects various line-of-business applications 122. Line-of-business applications 122 include, for example, an e-mail application such as Microsoft Exchange® that is used in the enterprise 100. Security endpoint 105_N typically monitors e-mail to provide anti-virus and anti-spam protection.

[0033] Each of the security endpoints 105 in the enterprise 100 are normally arranged as individual islands, as indicated by the dashed rectangles in FIG 1. Accordingly, each security endpoint 105 is arranged for monitoring a subset of the available data in the enterprise 100 and for performing localized actions in response to a detected incident. In addition, each endpoint typically includes a local management function 135 – 1, 2 … N. As noted above, the individual local management functions are not generally integrated to provide a single point of management.

[0034] FIG 2 shows an illustrative ESAS arrangement 200 in which a channel 205 is provided to enable a semantic abstraction called a "security assessment" to be shared among multiple endpoints using a language / protocol that is commonly-utilized at each endpoint. The security assessment channel 205 facilitates a publish / subscribe model used by the endpoints for connecting the sources of security assessments (publishers) to the consumers of the security assessments (subscribers). As shown, both the publishers and subscribers on the security assessment channel 205 are endpoints 105.

[0035] The endpoints 105 are isolated from the mechanics of the actual transport and management of the publish/subscribe model through a semantic abstraction layer that is arranged to simplify interactions with the security assessment channel 205. The abstraction layer comprises tables describing the security assessment types to which the endpoints subscribe, and tables describing the security assessment types that endpoints publish (as described below, not all endpoints generally subscribe to all security assessment types). In

addition, the abstraction layer provides an API (application programming interface) for reading received security assessments, and an API for generating security assessments.

[0036] A specialized endpoint, ESAS central server 216, is coupled to the security assessment channel 205 and performs as a centralized audit point for the ESAS arrangement 200. Accordingly, the ESAS central server 216 subscribes to all security assessments and permanently logs them. ESAS central server 216 also receives and logs messages from the endpoints that indicate the local actions that are taken by an endpoint. The ESAS central server 216 thus provides administrators with security assessment monitoring functionality that gives a comprehensive view of the history and current status of the enterprise as a whole, and each ESAS-enabled endpoint.

[0037] FIG 3 shows an illustrative terminology hierarchy 300 that underlies a security assessment. A security assessment is defined as a tentative assignment of security meaning, or category, to information. Information, as used here, is defined as data with some context. Data is defined as discrete items devoid of context. These definitions may be further described by way of an example. As shown in FIG 3, a piece of data 305 is an event in an event log such as a failed login. Information 310 is data provided with context which, in this example, is that the failed login was the sixth such failure within 10 minutes on the same machine, a laptop named Laptop2. The security assessment 316, in this example, indicates that Laptop2 is categorized in a particular way, namely that it is assessed with a category of "Compromised," with high "severity," and where such assessment has low "fidelity" (these terms are defined and discussed below in more detail).

[0038] A security assessment may be performed on any object of interest in an enterprise security environment, such as a user or a device. In this illustrative example, assessments include four main object types: 1) Host – assessments about computers in an enterprise; 2) User – assessments about users or accounts in enterprise; 3) Service – assessments about a service provided to the enterprise such as a URL (Uniform Resource Locator) of a web site that has a reputation as being malicious; 4) Enterprise – assessments about the enterprise as a whole or a well-defined subset of the enterprise such as a department, subnet, site, or branch; and 5) Data – assessments about business-related data (e.g., as found in documents, e-mail, business data in a database etc.) that is present or accessed by objects in the enterprise.

[0039] It is emphasized that these object types are merely illustrative, and other object types may be used as required by specific scenarios. In most applications of enterprise security assessment sharing, endpoints only publish, and subscribe to, a subset of all of the

available security assessment types since particular endpoints are generally going to have interest in particular objects in the enterprise environment. In addition, while some endpoints will be both publishers and subscribers, there is no requirement for every endpoint to support both functionalities. For these reasons, the publish / subscribe model used herein is said to be loosely-coupled.

[0040] Table 1 below shows an illustrative set of assessment categories, and their mapping to specific object types, that may be contained in a typical security assessment:

Object Type	Assessment category	Description
Host	Vulnerable machine	Machine had vulnerable configuration or is missing some patches.
	Compromised machine	An endpoint detected some evidence that the machine might be compromised by a malicious software / user.
	Machine under attack	An attack attempt was detected without an evidence for success
	Machine of interest	An endpoint has a general suspicion about a machine without the ability to pin point what is wrong.
User	Compromised user	An endpoint detects some evidence that the user / account might be compromised.
	User under attack	An attack attempt was detected without an evidence for success
	Malicious user	An endpoint or an administrator detects that a user is a malicious one and actively (i.e., on purpose) performs illegal actions.
	User of interest	An endpoint has a general suspicion about a user / account without the ability to pin point what is wrong.
Enterprise	Enterprise under attack	An endpoint detects that an enterprise is under attack without evidence that a significant part of the enterprise is compromised.
	Compromised enterprise	An endpoint detects that significant part of the enterprise is compromised (machines / users).
Service (e.g., a web site)	Malicious	A URL (Uniform Resource Locator) has a malicious reputation.
Data	Compromised	An endpoint detects some evidence that some business-related data in the enterprise is compromised.
	Corrupted	An endpoint detects some evidence

		that some business-related data in the enterprise is corrupted.

Table 1

[0041] In the present illustrative ESAS arrangement, four levels of severity are typically utilized: low, medium, high, and critical. Three levels of fidelity are typically utilized: low medium, and high. Note that the number of levels for both severity and fidelity can be arranged to be different depending on the assessment category. For example, it is possible to use the three severity levels for the assessment category of "vulnerable machine" while using four severity levels for the assessment category of "compromised machine." The particular choice of number of levels to be utilized will depend on the requirements of a specific application of the present enterprise security assessment sharing.

[0042] A security assessment uses information that is available at the time the assessment is made and relies on the particular security expertise and knowledge that is resident in the endpoint that produces it. A security assessment is tentative because confidence in any particular event can never be absolute, and also because the assessment is temporary in nature as it relies on information that is present at the time it was produced. At some future time, other information will be available, so the security assessment may change.

[0043] The tentative nature of a security assessment is reflected in two fields included in each assessment – fidelity and time-to-live ("TTL"). The fidelity field provides a way for endpoints to express their confidence level in an assignment of a broader contextual meaning to information being analyzed. The TTL field enables endpoints to reflect the best estimate of the time period for which the security assessment is expected to be valid. Or alternatively, the TTL field provides the best estimate for a future security assessment update. When a TTL expires, an endpoint that takes actions based on a security assessment to which it subscribes is expected to roll-back such actions when the TTL of that assessment expires. Thus, the TTL provides a safety valve functionality to prevent a user or a machine from getting inappropriately trapped with restricted access due to a false positive, or the loss of a message somewhere in the enterprise. However, if such restricted access is indeed appropriate, then either a new security assessment may be generated to continue the restriction, or the TTL extended.

[0044] The security assessment is designed to enable precise semantics (i.e., the meaning imparted by the categorization used in the security assessment) using a compact vocabulary. As shown in FIG 4, two of the endpoints 105 in the enterprise log data about events that occur in their respective areas of interest. The host event logs 405 and firewall event logs 412 thus contain very large amounts of data. Typically, the data is processed in the respective endpoints using correlation rules 420 and 425 in order to identify events of interest. The correlation rules, which are often numerous, define the localized sponsors or actions taken responsibly to a detected event.

[0045] By comparison, the security assessments, indicated by reference numeral 432, contain only a relatively small amount of data. As security assessments are utilized to assign broad context to information, they provide answers to the questions: Who created the assessment? When? Why? For how long? And, on which object does the assessment apply? Thus, in order to make use of a security assessment, an endpoint need only understand the few assessment types of interest as compared with the unbounded number of information messages that result from application of the correlation rules. Accordingly, the complexity of the data collected by each endpoint is reduced by mapping information into one or more of the assessment types. Using security assessments thus enables relevant information to be provided to subscribing endpoints without requiring that large amounts of data or information be shared across the enterprise.

[0046] Table 2 below provides an illustrative set of fields that may be included in a typical security assessment.

Field	**Description**
Incident ID	A unique identifier that represents the incident for which a security assessment was created. Since there may be several assessments that are tied to the same incident (e.g., detection, response, approval, cancellation), this field also uses a correlation value
Instance ID	A unique identifier that represents a single security assessment message
Source	Endpoint type (e.g., host security, edge firewall, NAP, NIDS, etc.) and unique endpoint ID
Relates To	If an assessment was created responsively to other

	security assessments then this field will contain the list of all the assessment's incident IDs
Object type	host, user, reputation, enterprise or other object type
Object ID	An identifier for the object. Endpoints often describe the same object in different ways. For example, a host can be identified by its FQDN (Fully Qualified Domain Name), IP (Internet Protocol), MAC (Media Access Control), or SMTP (Simple Mail Transfer Protocol) addresses, etc.
Category	Compromised, Vulnerable, Under Attack etc.
Operation	Published (upon detection), response, approval, cancellation, etc.
Severity	Severity of the incident
Fidelity	Confidence of the endpoint in its detection of an incident
Creation time	GMT (Greenwich Mean Time) and local time
TTL	Time-to-Live in minutes
Description	A human-readable format that explains why the assessment was created
Data	Private information about why the assessment was created. Other endpoints may use this data for additional correlations
Response	(Optional) a compound node that contains the set of responses that were taken by the endpoint
By	(Optional) in case the assessment was approved / canceled / created by a user, this field will contain the name of the user
Version	Schema version of the assessment

[0047] Using the fields in Table 2, a security assessment is capable of expressing the following events:

1. Detection. An endpoint executes some analysis to deduce that some abnormal behavior has occurred (compromised machine, vulnerable machine, compromised user, etc.);
2. Response. Endpoints take actions as a result of security assessments. An endpoint should notify the system (specifically, the ESAS central server 216 in FIG 2) when actions are taken. Responses may include, for example, blocking traffic, triggering a scan, resetting a password, collecting more data about a machine, and similar actions. Note that some responses such as resetting a password or triggering a scan are intermittent, while other responses are persistent and need to be rolled back in order to be canceled;
3. Assessment approval. An administrator can approve assessments manually using an interface to the ESAS central server 216. Endpoints shall be notified upon such approval so they will perform the "required manual approval" responses;
4. Cancellation. An administrator or an endpoint can cancel an existing security assessment;
5. Response roll-back. An endpoint notifies the system (ESAS central server 216 in FIG 2) that it rolled-back all responses / actions that were taken due to a specific assessment;
6. Health information assessments such as connectivity verifiers, latency checking, and error information;
7. Request for investigation data. This is a request from an endpoint to another endpoint to send all its data about an object collected in a given time period; and
8. Request for investigation data completed. This is a method for the endpoint to acknowledge that it performed the request. The response to the request is sent after the data was stored / sent.

[0048] In this illustrative example of enterprise security assessment sharing, each endpoint is arranged to perform at least some of the tasks noted below. In some arrangements, each endpoint is enhanced with additional functionality as required to perform such tasks through use of a discrete ESAS agent. Alternatively, the enhanced functionality may be more tightly integrated into the core functionality provided by the

endpoint, and a separate or discrete agent may not necessarily be embodied in the endpoint. Such tasks include:

1. Generating new security assessments based on the locally available information about the monitored system and the security context;
2. Subscribing to a subset of available security assessments from other endpoints;
3. Processing incoming security assessments to thereby influence the security context. The processing can result in the generation of new security assessments;
4. Taking local actions according to response policies;
5. Rolling back (self-recovery) a local action when an assessment that caused it expires (i.e., expiration of the associated TTL).

[0049] FIG 5 shows an illustrative example of an ESAS agent 505 disposed in an endpoint which subscribes to a subset of available assessments from other endpoints 105 – 1, 2 … N (FIG 1) over the security assessment channel 205. As noted above, the functionality provided by ESAS agent 505 may alternatively be directly integrated with the core functionality of the endpoint.

[0050] A plurality of security assessments 506 are available for each of the security assessment types (i.e., hosts, users, reputations, and enterprise). As indicated by reference numeral 511, in this illustrative example, ESAS agent 505 subscribes to security assessments having a "host" object type with an assessment category of "vulnerable." It is emphasized that a particular combination of object types and assessment categories that is of interest can be different for different endpoints. Again, using the loosely coupled publish/subscribe model there is no requirement that every endpoint subscribes to every security assessment.

[0051] At process block 514, the endpoint processes the received security assessment using correlation rules 522 and locally-available data 527 that may be of some relevance. The outputs of such assessment process include the generation of the new assessment 530 and/or an invocation of a local action 535. As noted above, such local action is subject to roll-back 541 (i.e., self-recovery) when the received assessment expires according to the TTL field contained therein.

[0052] ESAS agent 505 interprets security assessments according to the following rules:

1. In generating a security assessment about a particular object, an endpoint can take into account any combination of the following:
   a) All of the locally-available information about the object or any other objects the endpoint monitors;
   b) All the currently active security assessments (i.e., those having an unexpired TTL) the endpoint has received;
   c) All the local actions the endpoint has taken in the past.
2. All of the endpoints in the present enterprise security assessment sharing arrangement comply with the principle that all of the sets of locally-available information in the endpoints are mutually exclusive. That is the present arrangement has no more than one endpoint processing a particular piece of local information.
3. Security assessments are typically interpreted as describing an endpoint's assessment about the current and future security state of an object.

[0053] It is important to note that the Rule 2 refers to the exclusivity of information and not data as these terms were defined above. Two endpoints can process the same or overlapping data sources provided that the information they extract from the data and later use to generate assessments is exclusive.

[0054] To illustrate the implications of Rule 3 consider the following example where a scheduled antivirus scan of a machine detects and removes a piece of known malware. Based on this detection, other locally available information, a received currently active assessment, and on the endpoint's embedded knowledge about current security incidents, the endpoint may conclude one of the following: 1) a machine was infected in the past, but is now clean and does not pose any additional future security risk; 2) The machine was infected and, while the particular malware was removed, it is possible or likely that it still poses a security risk. According to Rule 3, an endpoint should generate a security assessment about the machine in the latter case and not generate one in the former.

[0055] FIG 6 is a diagram of a first illustrative scenario in which a plurality of ESAS-enabled endpoints are coupled to the security assessment channel 205, and a detected incident at one endpoint triggers responses at multiple other endpoints. This illustrative scenario is described in three stages. As indicated by reference numeral 610, the edge firewall 105_2 first identifies a potentially compromised client, for example, because it

creates so many connections to the perimeter network 112 (FIG 1) that the most likely explanation for the behavior is the existence of a security compromise. Second, the edge firewall 105_2 sends a security assessment that indicates the particular client is "compromised" with high severity and high fidelity, as indicated by reference numeral 620, over the security channel 205 to subscribing endpoints.

[0056] Third, the subscribing endpoints 105 - 1, 3 … N and the ESAS central server 216 which receive the security assessment apply their specific security expertise through the application of their own correlation rules and locally-available data to trigger an appropriate action. As collectively indicated by reference numeral 630 in FIG 6, the host security endpoint 105_1 performs an on-demand scan. The NAP endpoint 105_3 revokes the IP security certificate for the identified compromised client and implements a port shutdown. The line-of-business security endpoint 105_N temporarily suspends instant messaging ("IM") traffic to the compromised client based on the received security assessment. The ESAS central server 216 raises an alert for a security analyst (e.g., an administrator) and also logs all of the security assessments and actions invoked.

[0057] The first illustrative scenario described above provides a case where the endpoint detecting the suspicious incident generates a security assessment with high severity and high fidelity (i.e., the endpoint has a high degree of confidence that it has validly detected a serious incident). By comparison, FIG 7 is a diagram of a second illustrative scenario in which a low fidelity security assessment is sent over the security assessment channel 205 that triggers the generation of a high fidelity assessment by a receiving endpoint which also performs cross-object mapping.

[0058] This second illustrative scenario is also described in three stages. As indicated by reference numeral 710, the edge firewall 105_2 first detects a large number of client connections to the perimeter network 112 (FIG 1). However, unlike the first illustrative scenario shown in FIG 6 and described in the accompanying text, the number of connections being established by the client is not so high that the edge firewall 105_2 can be absolutely certain that the client has been compromised. In current enterprise security systems, when an endpoint sees such data it generally just drops the data and takes no action as there is not enough evidence to warrant the typically harsh response such as disconnecting the machine. By comparison, in the present scenario the edge firewall 105_2 in the second stage sends a security assessment 715 over the security assessment channel 205 that indicates that the particular client is compromised with medium severity and low fidelity, as indicated by reference numeral 720.

[0059] Here, the subscribing endpoints to the particular object referenced in the security assessment 715 generated by the edge firewall 105_2 include the host security endpoint 105_1 and the ESAS central server 216. While such a low fidelity data normally does not trigger an action to be taken at an endpoint in current security products, in accordance with the present enterprise security assessment sharing, the host security endpoint 105_1 looks at its own local data differently in light of the received security assessment from the edge firewall 105_2. In this case, the local data resulting from an on-demand scan at the host security endpoint 105_1 and the information contained in the security assessment from the edge firewall 105_2 are used to generate new assessments 725 and 728. Thus, the host security endpoint 105_1 has information that, by itself, does not warrant the generation of a new security assessment, but when reinforced with even a low fidelity assessment from another endpoint, as in this case, there is sufficient evidence to justify the creation of the new security assessments 725 and 728 which each have high fidelity.

[0060] The host security endpoint 105_1 places the new security assessments 725 and 728 onto the security assessment channel 205. The new security assessments 725 and 728 are received over the security assessment channel 205 by the subscribing endpoints which, in this illustrative scenario, include the edge firewall 105_2, and the ESAS central server 216 for security assessment 725 and line-of-business endpoint 105_N for security assessment 728.

[0061] Note that the line-of-business endpoint 105_N was not a subscriber to the original security assessment 715 produced by the edge firewall 105_2 because the reference object type is a machine and the line-of-business endpoint 105_N, by virtue of its role in protecting e-mail, is typically concerned with the users. However, in this second illustrative scenario the host security endpoint 105_1 maps from a host object type to a user object type when it generates a new security assessment 728. Such cross object mapping capability may be beneficial in many scenarios as it is contemplated that a high severity incident, such as malware or malicious activity that potentially compromises the data confidentiality or integrity of a host computer, may also potentially compromise the user as well. A security assessment may be generated that cross maps the high severity incident from the host object type to a user object type with a certain degree of fidelity. Similarly, a critical severity incident in which malware or malicious activity has actually caused a loss of data integrity on a host computer, a security assessment for a user object type may be generated with even higher fidelity.

[0062] In stage three, the new security assessments 725 and 728 trigger a variety of respective actions at the receiving endpoints, as collectively indicated by reference numeral

730. Specifically, the edge firewall 105_2 blocks all access by the compromised client with the exception of software updates and/or mission-critical access. The line-of-business endpoint 105_N temporarily suspends outgoing e-mails. And, as with the first illustrative scenario, the ESAS central server 216 continues to log all assessments and actions. As noted above, such restrictions are enforced only during the period of time for which the TTL associated with the new security assessments 725 and 728 remains valid. When the new security assessments expire, the actions taken by the respective endpoints are rolled-back unless the TTL is extended or a new security assessment which invokes the restrictive actions is received.

[0063] FIG 8 is a diagram of a third illustrative scenario that shows the targeted use of remediation techniques. This third illustrative scenario is described in three stages. As indicated by reference numeral 810, the edge firewall 105_2 first detects a large number of client connections to the perimeter network 112 (FIG 1). Second, the edge firewall 105_2 sends a security assessment 815 that indicates the particular client is "compromised" with high severity and high fidelity, as indicated by reference numeral 820, over the security channel 205 to subscribing endpoints. The subscribing endpoints include the host security endpoint 105_1 the NAP endpoint 105_3 and the ESAS central server 216.

[0064] The host security endpoint 105_1 reviews the received security assessment and applies its specific security expertise using correlation rules and any relevant locally-available data. In this illustrative example, the host security endpoint 105_1 responsively generates a new security assessment 825 containing a user object type to which the line-of-business security endpoint 105_N subscribes.

[0065] In the third stage of the scenario, the remediation techniques employed by the endpoints are considered expensive in terms of their potential impact on business operations within the enterprise 100 (FIG 1). For example, as indicated by reference numeral 830, the line-of-business security endpoint 105_N implements a response policy that entails temporarily suspending outgoing e-mail. In addition, the host security endpoint 105_1 performs an on-demand scan and if no results are achieved, performs a deep scan. While such remediation techniques can be very effective in addressing malware, malicious users, and other problems, they typically impart significant expense to the enterprise. For example a user whose outgoing e-mail is suspended will be less productive, and deep scanning typically requires one or more reboots which will remove the machine from service for a period of time.

[0066] The present ESAS arrangement advantageously enables these effective, albeit expensive, remediation techniques to be applied in a targeted manner, and not merely in a general way or across-the-board which can be unjustified for some machines and/or users. Only objects in the environment that are deemed suspicious, using pre-defined criteria, will be subject to these particular remediation techniques.

[0067] FIG 9 shows an illustrative screen 900 that is provided by a graphical user interface ("GUI") which enables a user, such as an administrator, to manage and define the response policies of endpoints in the enterprise 100 (FIG 1). In some applications, the GUI is hosted on the ESAS central server 216 (FIG 2). Advantageously, the semantic abstraction layer embodied by the security assessments enables enterprise-wide security policies to be established using a very compact and clear methodology. That is, response policies may be configured using a security assessment as a defined starting point without being concerned as to which endpoint in the enterprise created the security assessment, or how the endpoint came to the conclusion reflected in the security assessment. The security assessments, with their compact taxonomy, thus function as natural anchors to enterprise-wide security response policies. Without the present ESAS arrangement to streamline the configuration of response policies, a user would need to consider every event and/or alert that every endpoint could possibly generate, and then define what to do with each such event.

[0068] Screen 900 is an illustrative example showing the configuration of enterprise-wide response policies, using fields 903 - 1, 2 … N for a number of different endpoints, for the case of an assessment category of a compromised machine with critical severity, which defines a starting point (i.e., "anchor" point) for the response policy configuration as indicated by reference numeral 906. It is emphasized that other user interface screens would be utilized for other assessment categories, object types, severity levels, etc., so that the user is enabled to define response policies for the plurality of different starting points that are likely to be used in a particular enterprise security environment. The response policies, in this particular example, are set depending upon the fidelity of a particular security assessment for a set severity level of "Critical." Fields 903 include a number of respective subfields that are arranged to reflect user-defined input, using for example, text entry boxes, drop-down menus, and the like that are employed in typical GUIs.

[0069] As indicated in subfield 910, for a security assessment indicating a compromised machine with critical severity, the edge firewall 105_2 (FIG 1) is configured to increase the amount of auditing (i.e., moving to a deep audit level that increases the amount of data collected as compared with a normal level of auditing) when a security assessment has low

fidelity. Subfield 913 shows that for an assessment having medium fidelity, the edge firewall 105_2 increases the auditing level and also restricts Internet access to the suspected compromised machine to only "white list" URLs which typically include sites that are known to not be malicious. When the fidelity is high, as shown by subfield 916, access to the Internet is blocked completely.

[0070] Field 903_2 shows the response policy configuration for the host security endpoint 105_1 (FIG 1). For a security assessment having low fidelity and indicating a compromised machine with critical severity, the host security endpoint 105_1 increases the amount of auditing to a deep audit level, as indicated by subfield 920. Subfield 923 indicates that for the cases of medium and high fidelity, the host security endpoint 105_1 increases its auditing, and also increases performs a deep scan of it hosts (where a "deep" scan may entail the computer to be rebooted one or more times).

[0071] Field 903_N shows the response policy configuration for the line-of-business security endpoint 105_N (FIG 1). For a security assessment having low fidelity and indicating a compromised machine with critical severity, the line-of-business security endpoint 105_N increases the amount of auditing to a deep audit level, as indicated by subfield 926. Subfield 932 indicates that for a security assessment having medium fidelity, the line-of-business security endpoint 105_N increases its data collection to a deep audit, and also restricts file attachments to e-mail. Subfield 935 indicates that for a security assessment having high fidelity, line-of-business security endpoint 105_N blocks all instant messaging ("IM") traffic.

[0072] Field 941 shows the response policy configuration for the ESAS central server 216 in FIG 2. For a security assessment having high fidelity, the ESAS central server 216, as indicated in subfield 943, executes a port shutdown for the affected machine and generates an alert that the associated user account has been suspended. As with the subfields discussed above, subfield 943 is typically arranged to accept user-defined input.

[0073] FIG 10 shows an illustrative screen 1000 that is provided by a GUI which enables a user, such as an administrator, to manage and define the response policies of endpoints in the enterprise 100 (FIG 1). This screen and GUI may be used to supplement the arrangement shown in FIG 9 and described in the accompanying text, or be used as an alternative arrangement. Screen 1000 provides a single view of response policy configuration for an assessment category of "Compromised Machine" for various levels of fidelity and for all levels of severity. In this illustrative example, a security assessment type 1006 having an assessment category of "Compromised" with any severity level is used as the anchor to the response policy configuration shown. It is emphasized that similar screens

are contemplated for use with other object types and assessment categories. As with the arrangement shown in FIG 9 and described in the accompanying text, the arrangement shown in FIG 10 provides a very compact management interface for defining response policies of endpoints throughout an enterprise.

[0074] FIG 11 shows an illustrative enterprise security arrangement 1100 in which the present ESAS feature set – including security assessment sharing, security assessment monitoring by ESAS central server (as described in the text accompanying FIG 2), and compact taxonomy for enterprise-wide response policy configuration (as described in the text accompanying FIGs 9 and 10) – functions as an enterprise security management layer 1105. That is, the ESAS feature set, indicated by reference numeral 1108, is commonly shared among all of the endpoints in the enterprise environment, and is not limited to being part of a single enterprise security product island. The semantic abstraction layer formed by the security assessments shared over the security assessment channel 205 (FIG 2) enables a single and consistent management interface to be utilized to thereby create a more integrated approach to enterprise security.

[0075] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

WHAT IS CLAIMED IS:

1. A security-related information sharing model by which security-related information is shareable among a plurality of endpoints in an enterprise security environment, the model facilitating use of a method comprising the steps of:

describing an object in the environment using a semantic abstraction of security-related information that is available to an endpoint, the semantic abstraction i) being categorized by type, and ii) being commonly utilizable by the endpoints; and

using a publish and subscribe model by which a publishing endpoint publishes the semantic abstraction to which a subscribing endpoint subscribes according to a subscription, the subscription being based on the semantic abstraction type.

2. The security-related information sharing model of claim 1 in which the semantic abstraction is a security assessment that is arranged for providing an assignment of context by an endpoint to the security-related information using a pre-defined taxonomy.

3. The security-related information sharing model of claim 2 in which the pre-defined taxonomy utilizes a schematized vocabulary comprising object types and assessment categories.

4. The security-related information sharing model of claim 3 in which the object types include at least one of host, user, service, data, or enterprise.

5. The security-related information sharing model of claim 3 in which the assessment categories include at least one of vulnerable, compromised, under attack, of interest, corrupted, or malicious.

6. The security-related information sharing model of claim 3 in which particular ones of assessment categories are mapped to particular ones of object types.

7. The security-related information sharing model of claim 2 in which the security assessment comprises a plurality of fields, at least one of which is a fidelity field that is arranged to express a degree of confidence an endpoint has in the security assessment.

8. The security-related information sharing model of claim 2 in which the security assessment comprises a plurality of fields, at least one of which is a time-to-live field that is arranged to express an estimate by an endpoint for a time period for which the security assessment is expected to be valid.

9. The security-related information sharing model of claim 1 in which at least one of the endpoints in the plurality of endpoints comprises a security solution object, the object selected from one of security product, security solution, management product, management solution, security service, or management service.

10. A method for enabling an endpoint to share security-related data within an enterprise security environment, the method comprising the steps of:

generating a security assessment to describe an event, in which the generating is based at least in part on locally-available information about a system being monitored by the endpoint, the security assessment being arranged to provide contextual meaning to the event and being defined with a time interval over which the security assessment is valid;

receiving a current security assessment in accordance with a subscription to a subset of available security assessments generated by other endpoints in the enterprise security environment; and

taking a response in accordance with a response policy on a per security assessment basis.

11. The method of claim 10 in which the response comprises a local action.

12. The method of claim 10 in which the security context comprises one or more previously received security assessments from the subset so long as the previous security assessments are valid.

13. The method of claim 11 in which the locally-available information further comprises one or more past local actions taken by the endpoint.

14. The method of claim 11 including a further step of rolling back the local action once the received security assessment is no longer valid.

15. A method for configuring security polices across an enterprise, the method comprising the steps of:

defining a security assessment schema in which assessments of security events are generated by a plurality of endpoints in the enterprise, the assessments i) using a pre-defined taxonomy to provide contextual meaning to the security events and, ii) being categorized by type; and

using an assessment as an anchor point to a rules matrix, the rules matrix describing a response to each assessment type.

16. The method of claim 15 including a further step of collecting the assessments generated by the plurality of endpoints in a central location.

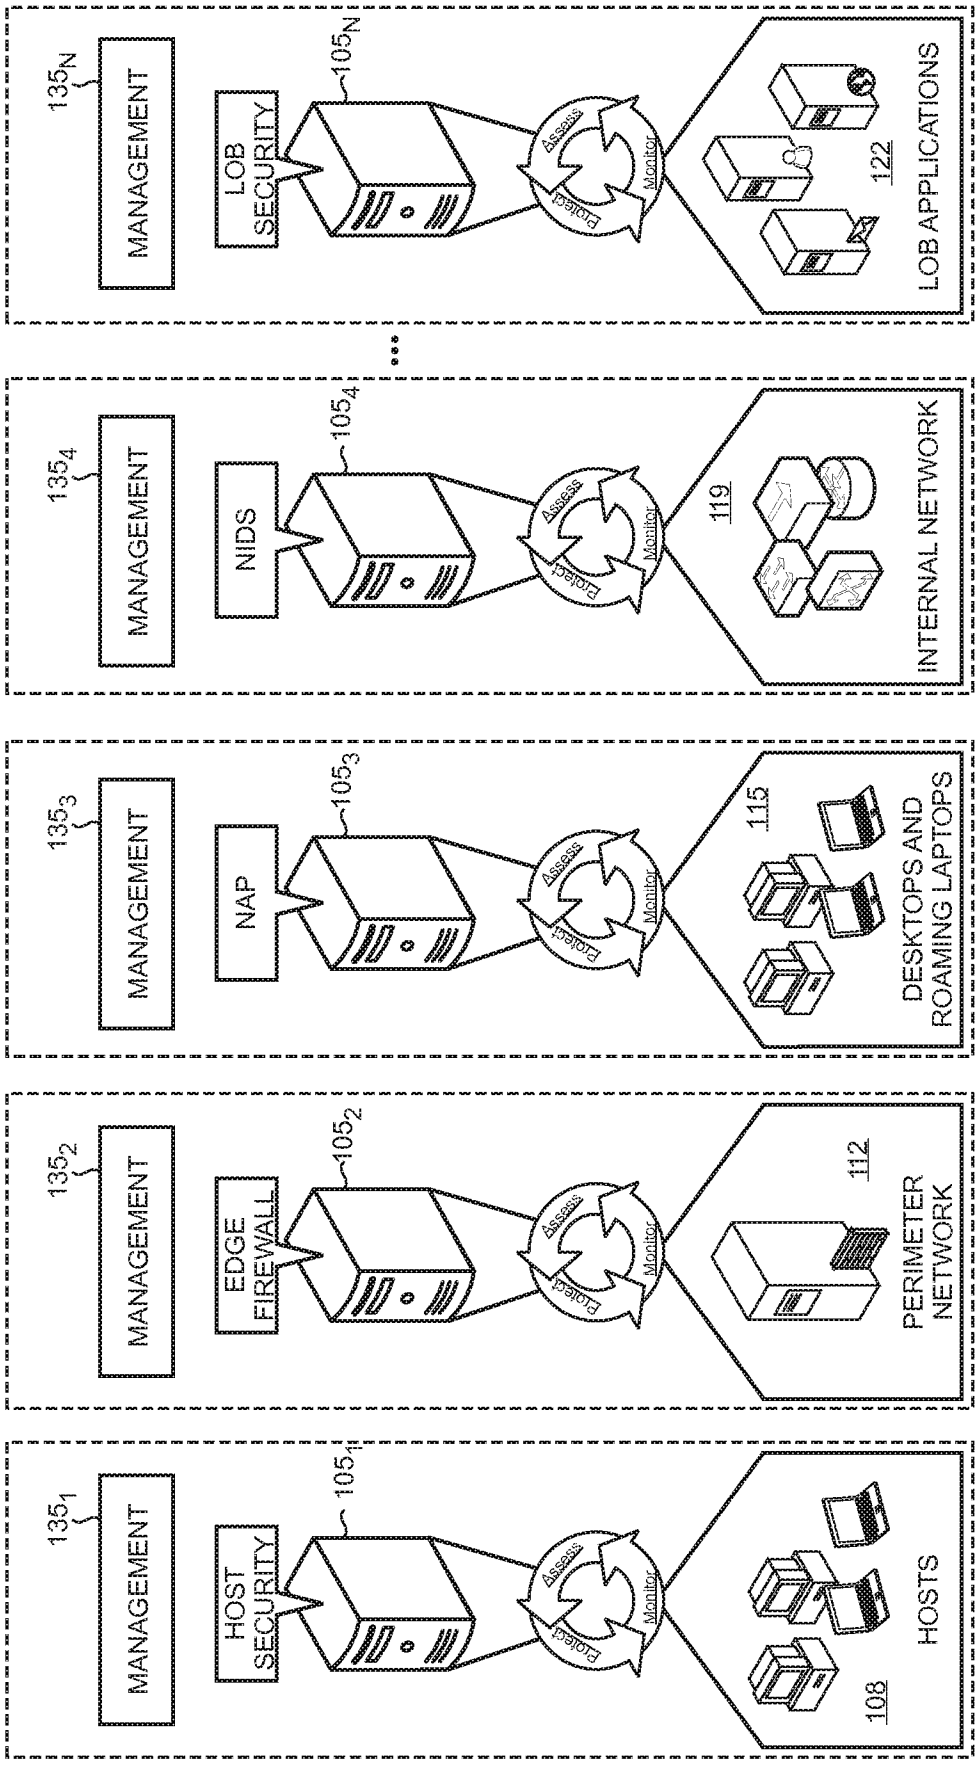
17. The method of claim 15 including a further step of collecting notifications generated by the endpoints, the notifications being associated with local actions taken by the endpoints.

18. The method of claim 15 in which an assessment type is defined by a combination of at least two of object type, assessment category, or event severity.

19. The method of claim 15 in which the rules matrix includes dimensions of assessment fidelity and event severity.

20. The method of claim 15 including a further step of providing a graphical user interface for displaying a representation of the anchor point and rules matrix.

FIG. 1


100
MANAGEMENT
135_1
HOST SECURITY
105_1
Assess
Protect
Monitor
108
HOSTS
MANAGEMENT
135_2
EDGE FIREWALL
105_2
Assess
Protect
Monitor
112
PERIMETER NETWORK
MANAGEMENT
135_3
NAP
105_3
Assess
Protect
Monitor
115
DESKTOPS AND ROAMING LAPTOPS
MANAGEMENT
135_4
NIDS
105_4
Assess
Protect
Monitor
119
INTERNAL NETWORK
MANAGEMENT
135_N
LOB SECURITY
105_N
Assess
Protect
Monitor
122
LOB APPLICATIONS

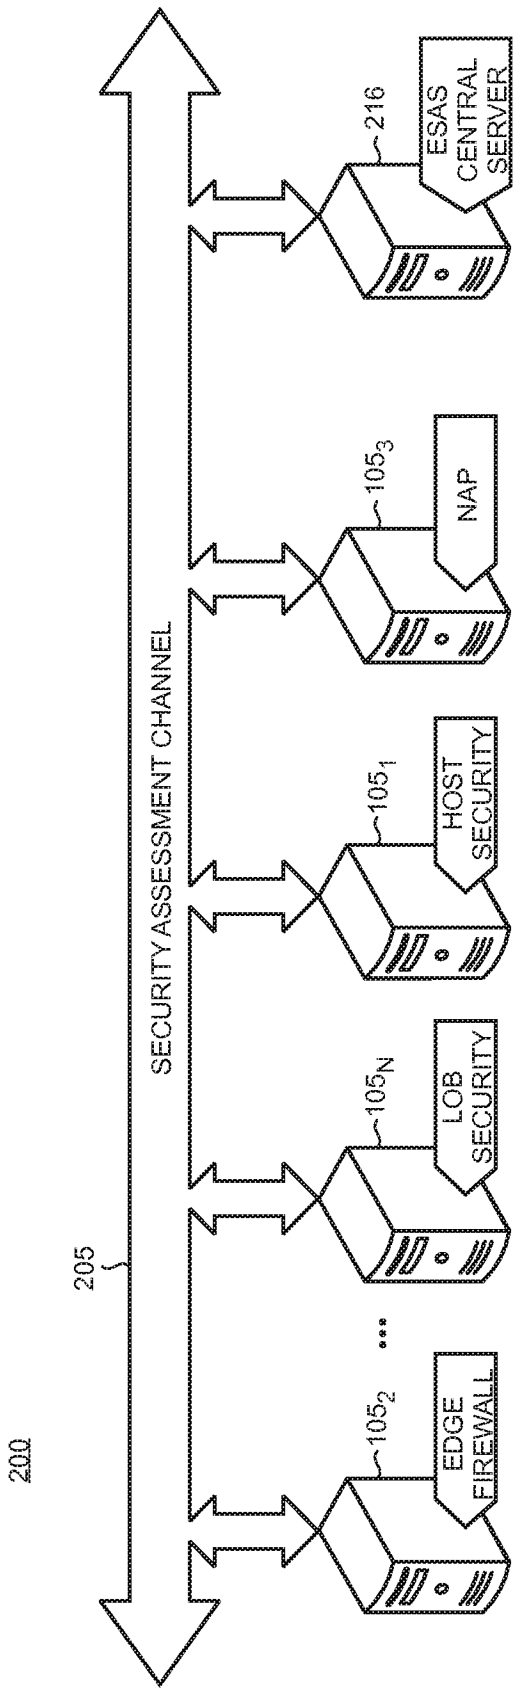

200
205
SECURITY ASSESSMENT CHANNEL
EDGE FIREWALL
105_2
LOB SECURITY
105_N
HOST SECURITY
105_1
NAP
105_3
ESAS CENTRAL SERVER
216

FIG. 2

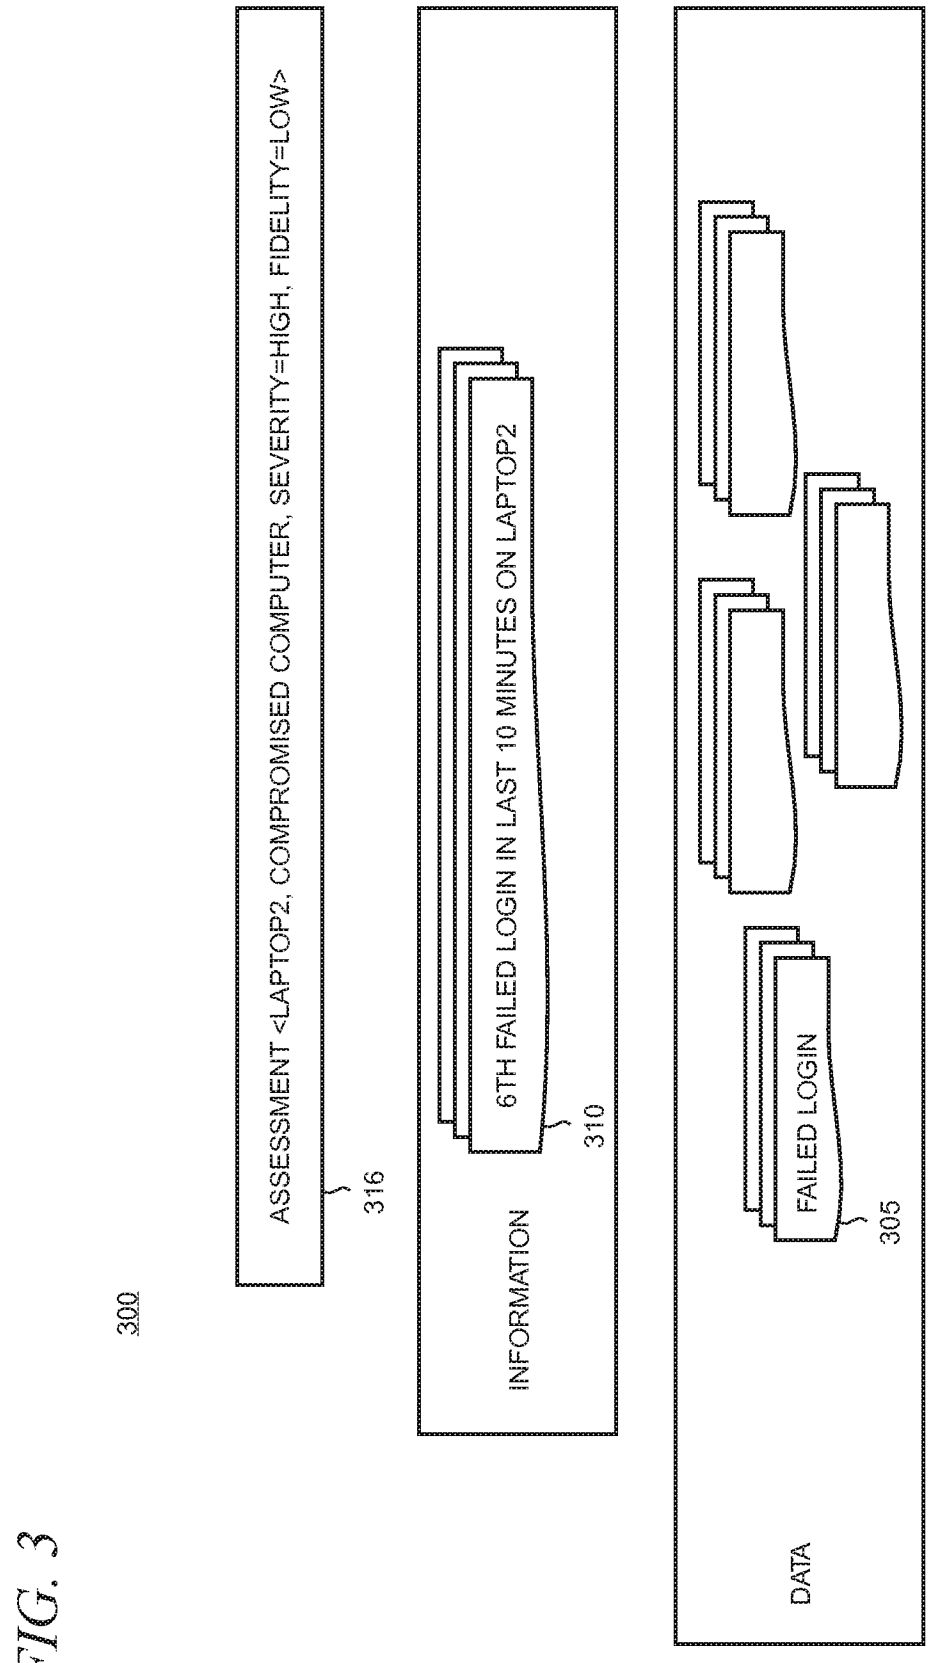

300
ASSESSMENT <LAPTOP2, COMPROMISED COMPUTER, SEVERITY=HIGH, FIDELITY=LOW>
316
INFORMATION
6TH FAILED LOGIN IN LAST 10 MINUTES ON LAPTOP2
310
DATA
FAILED LOGIN
305

FIG. 3

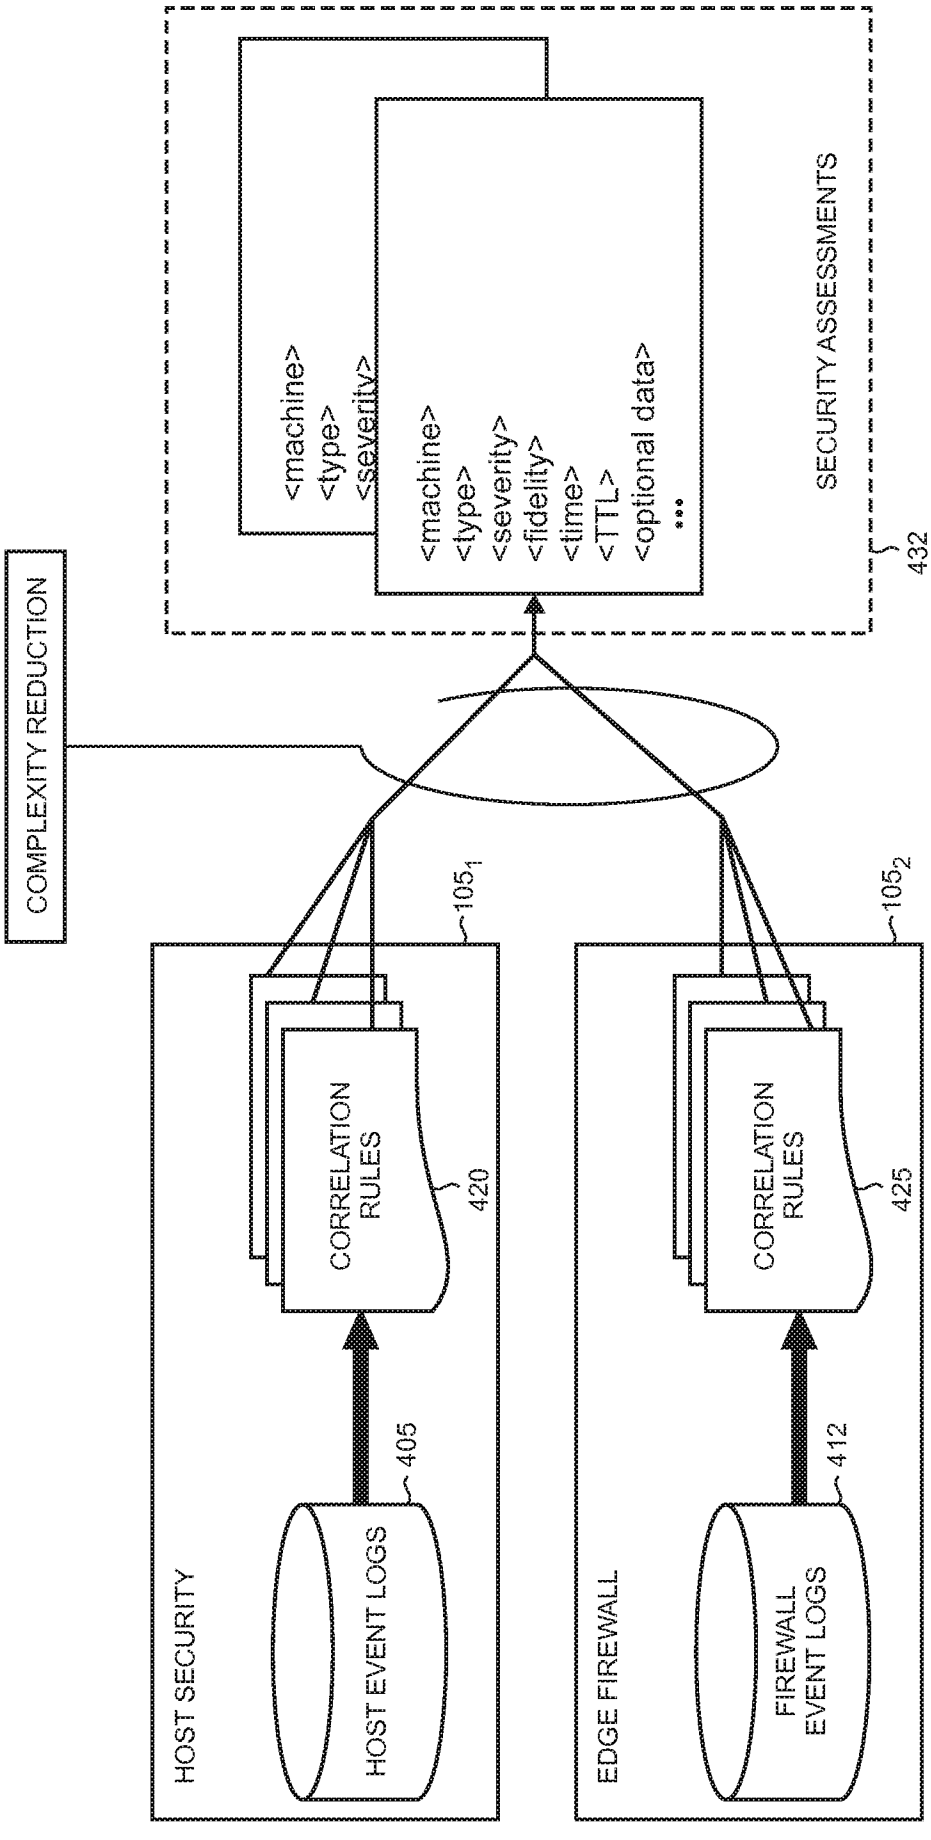

SECURITY ASSESSMENTS
432
<machine>
<type>
<severity>
<machine>
<type>
<severity>
<fidelity>
<time>
<TTL>
<optional data>
...
COMPLEXITY REDUCTION
HOST SECURITY
105_1
CORRELATION RULES
420
HOST EVENT LOGS
405
EDGE FIREWALL
105_2
CORRELATION RULES
425
FIREWALL EVENT LOGS
412

FIG. 4

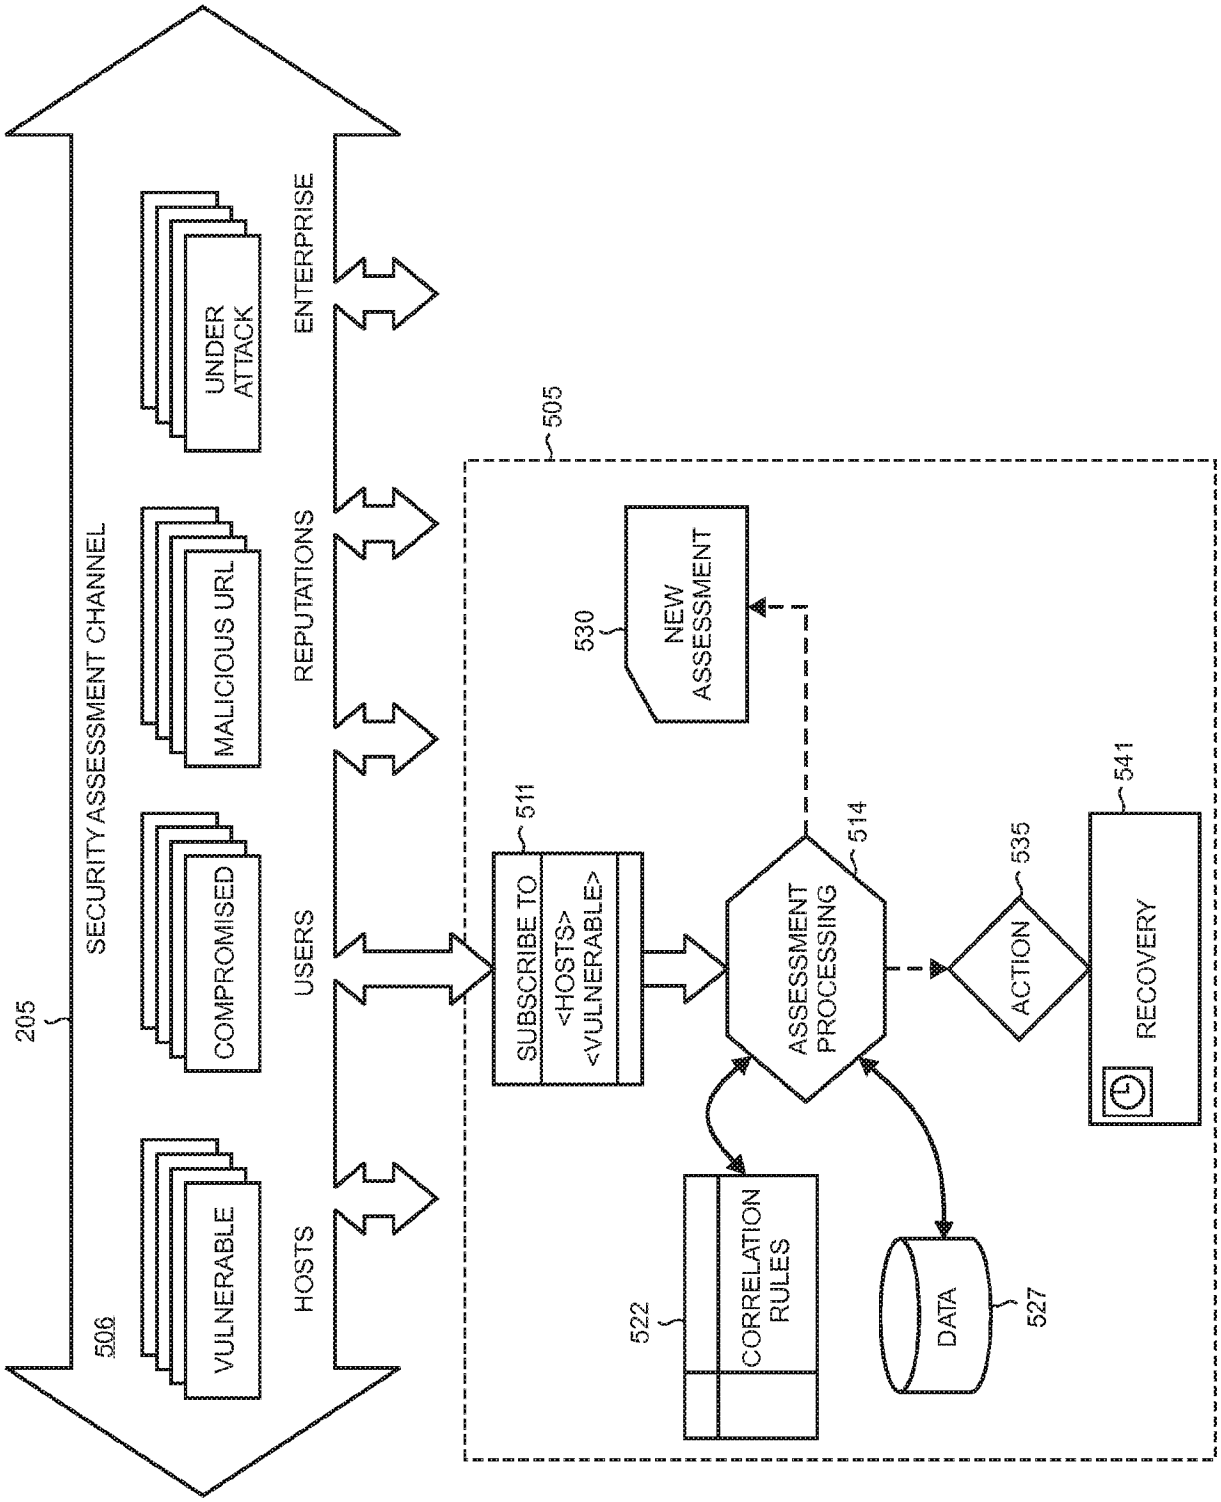

205
SECURITY ASSESSMENT CHANNEL
506
VULNERABLE
COMPROMISED
MALICIOUS URL
UNDER ATTACK
HOSTS
USERS
REPUTATIONS
ENTERPRISE
505
SUBSCRIBE TO
<HOSTS>
<VULNERABLE>
511
NEW ASSESSMENT
530
ASSESSMENT PROCESSING
514
ACTION
535
RECOVERY
541
CORRELATION RULES
522
DATA
527

FIG. 5

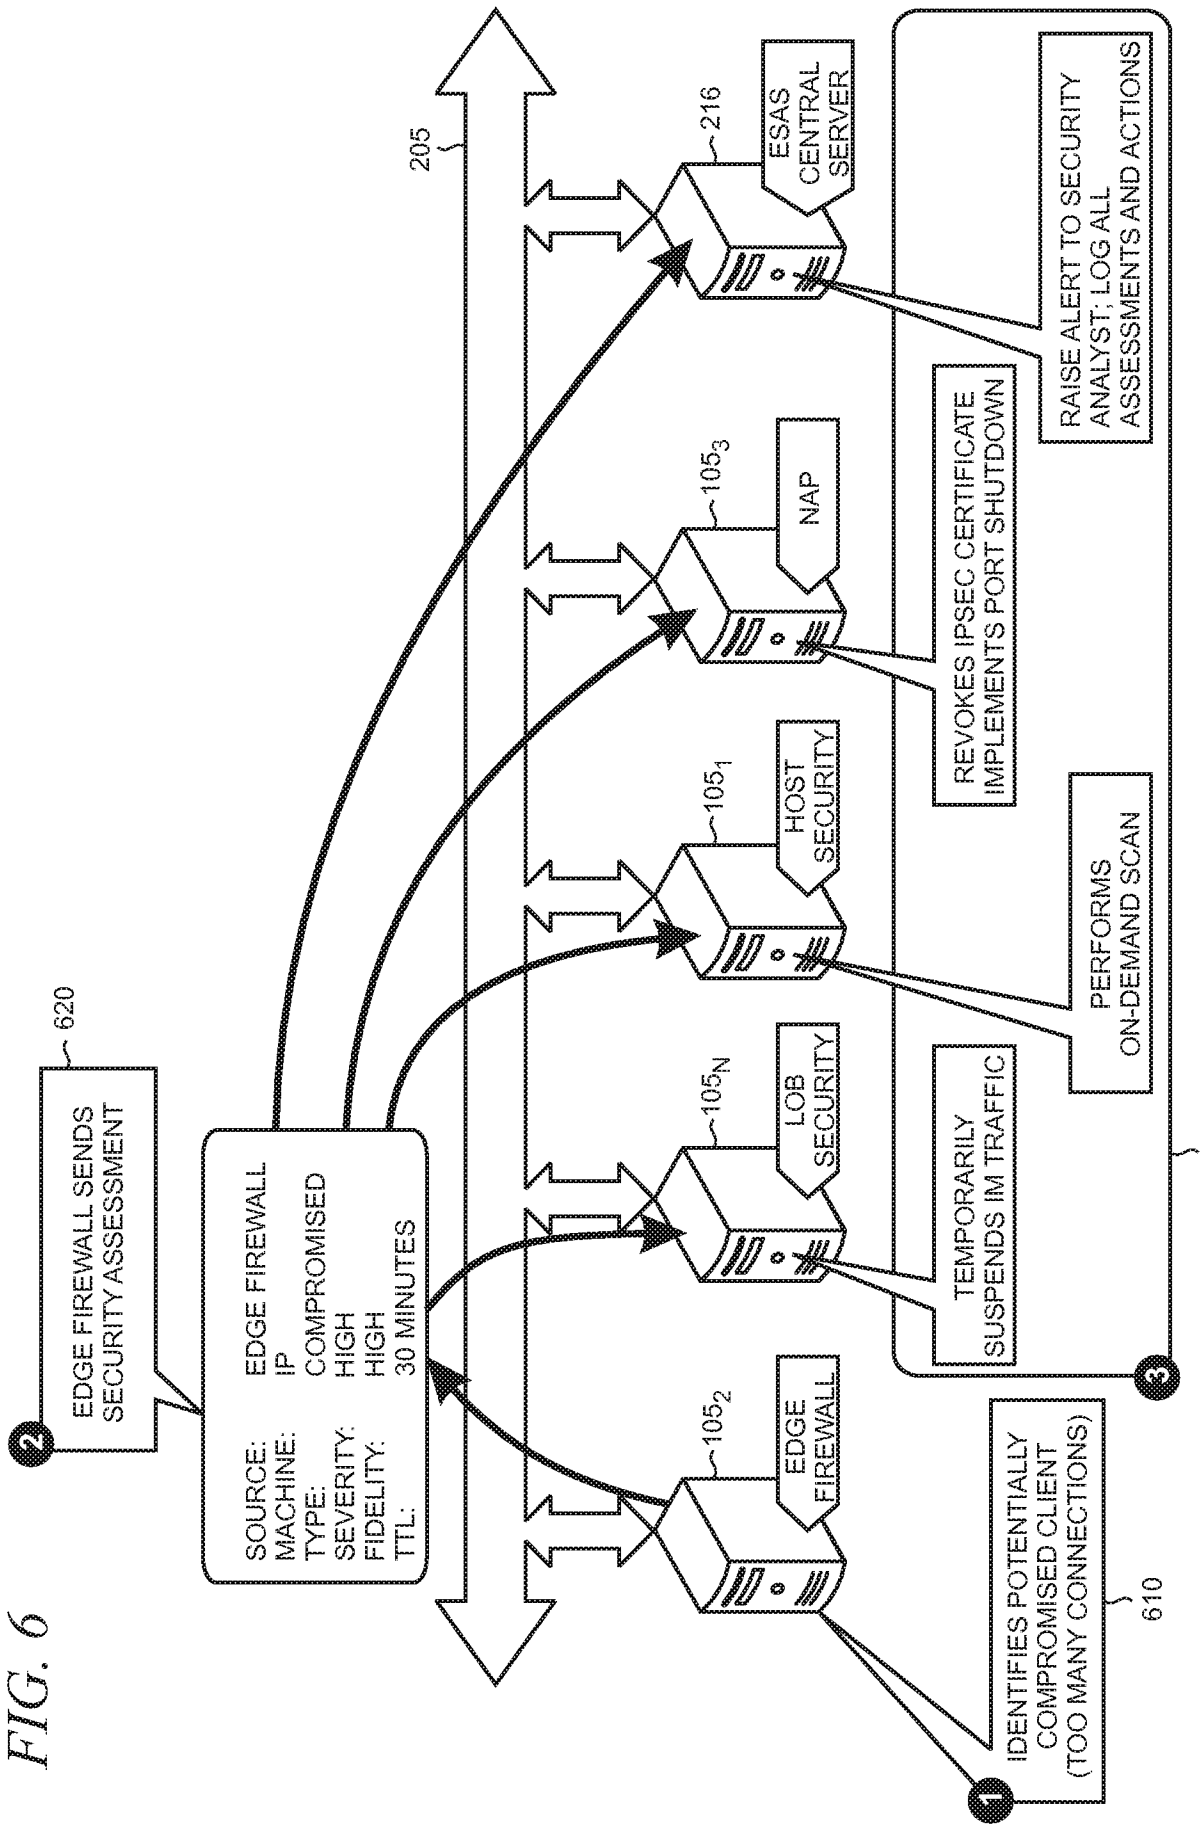

620
EDGE FIREWALL SENDS SECURITY ASSESSMENT
SOURCE: EDGE FIREWALL
MACHINE: IP
TYPE: COMPROMISED
SEVERITY: HIGH
FIDELITY: HIGH
TTL: 30 MINUTES
205
216
ESAS CENTRAL SERVER
105_3
NAP
105_1
HOST SECURITY
105_N
LOB SECURITY
105_2
EDGE FIREWALL
RAISE ALERT TO SECURITY ANALYST; LOG ALL ASSESSMENTS AND ACTIONS
REVOKES IPSEC CERTIFICATE IMPLEMENTS PORT SHUTDOWN
PERFORMS ON-DEMAND SCAN
TEMPORARILY SUSPENDS IM TRAFFIC
630
IDENTIFIES POTENTIALLY COMPROMISED CLIENT (TOO MANY CONNECTIONS)
610
1
2
3

FIG. 6

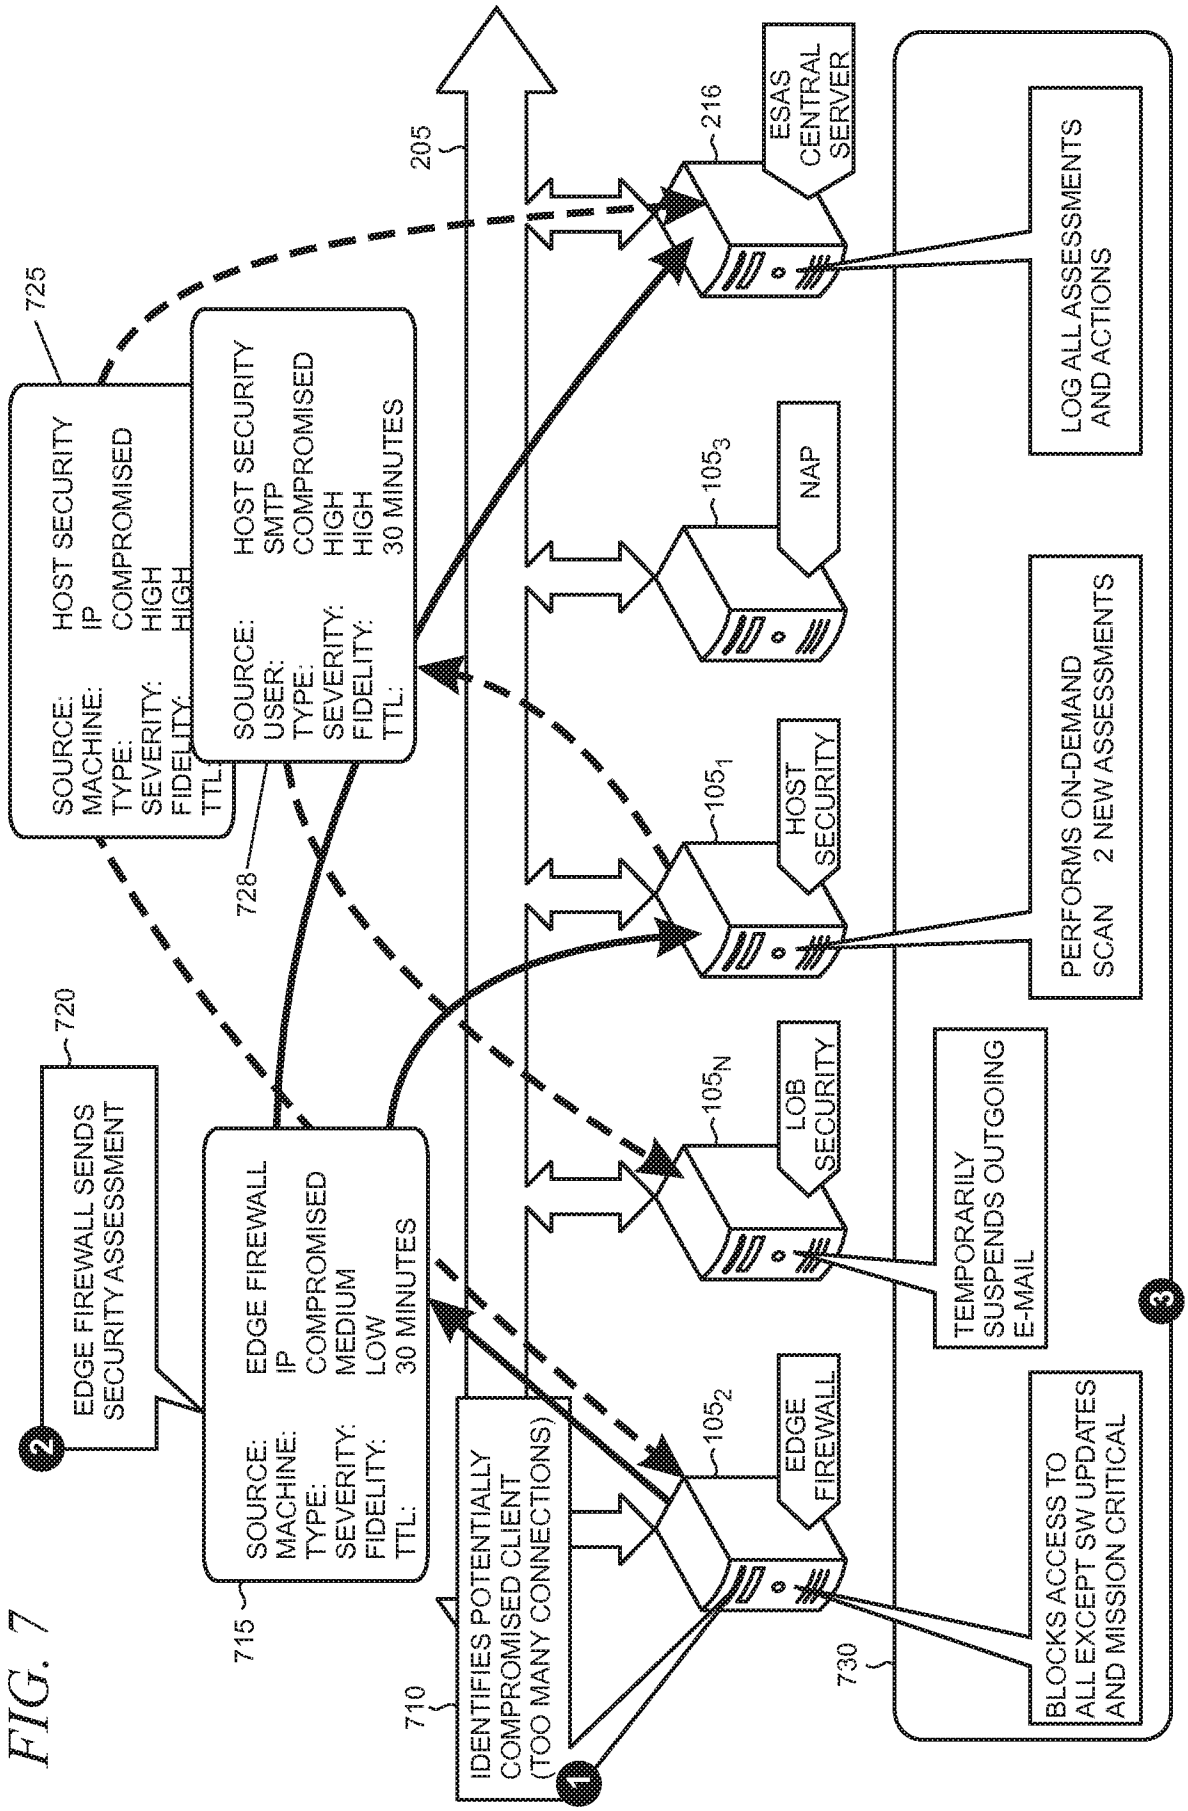

725
SOURCE: HOST SECURITY
MACHINE: IP
TYPE: COMPROMISED
SEVERITY: HIGH
FIDELITY: HIGH
TTL:
728
SOURCE: HOST SECURITY
USER: SMTP
TYPE: COMPROMISED
SEVERITY: HIGH
FIDELITY: HIGH
TTL: 30 MINUTES
205
216
ESAS CENTRAL SERVER
LOG ALL ASSESSMENTS AND ACTIONS
105 3
NAP
105 1
HOST SECURITY
PERFORMS ON-DEMAND SCAN 2 NEW ASSESSMENTS
105 N
LOB SECURITY
TEMPORARILY SUSPENDS OUTGOING E-MAIL
105 2
EDGE FIREWALL
BLOCKS ACCESS TO ALL EXCEPT SW UPDATES AND MISSION CRITICAL
730
720
2
EDGE FIREWALL SENDS SECURITY ASSESSMENT
715
SOURCE: EDGE FIREWALL
MACHINE: IP
TYPE: COMPROMISED
SEVERITY: MEDIUM
FIDELITY: LOW
TTL: 30 MINUTES
710
IDENTIFIES POTENTIALLY COMPROMISED CLIENT (TOO MANY CONNECTIONS)
1
3

FIG. 7

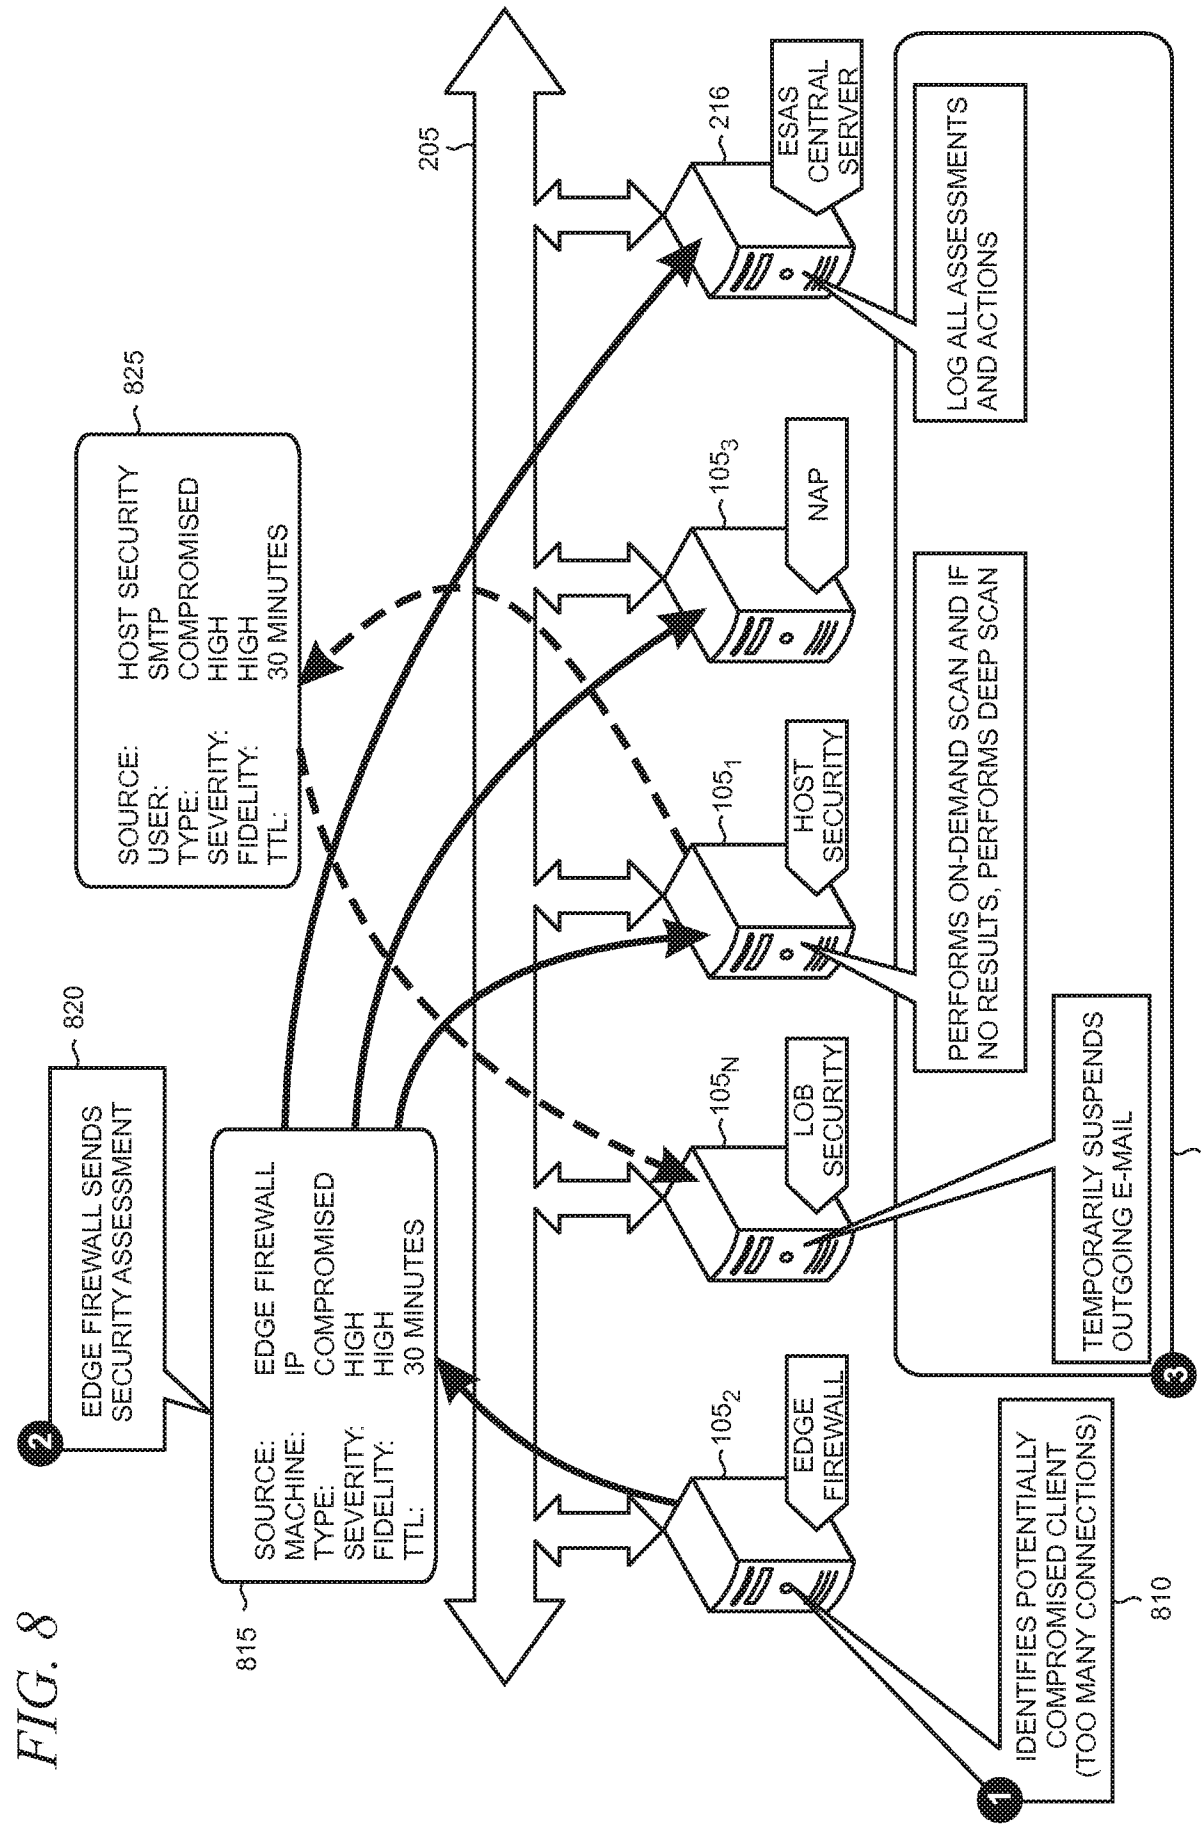

1
IDENTIFIES POTENTIALLY COMPROMISED CLIENT (TOO MANY CONNECTIONS)
810
2
EDGE FIREWALL SENDS SECURITY ASSESSMENT
820
3
TEMPORARILY SUSPENDS OUTGOING E-MAIL
830
SOURCE: EDGE FIREWALL
MACHINE: IP
TYPE: COMPROMISED
SEVERITY: HIGH
FIDELITY: HIGH
TTL: 30 MINUTES
815
SOURCE: HOST SECURITY
USER: SMTP
TYPE: COMPROMISED
SEVERITY: HIGH
FIDELITY: HIGH
TTL: 30 MINUTES
825
205
EDGE FIREWALL
105 2
LOB SECURITY
105 N
HOST SECURITY
105 1
NAP
105 3
ESAS CENTRAL SERVER
216
PERFORMS ON-DEMAND SCAN AND IF NO RESULTS, PERFORMS DEEP SCAN
LOG ALL ASSESSMENTS AND ACTIONS

FIG. 8

FIG. 9

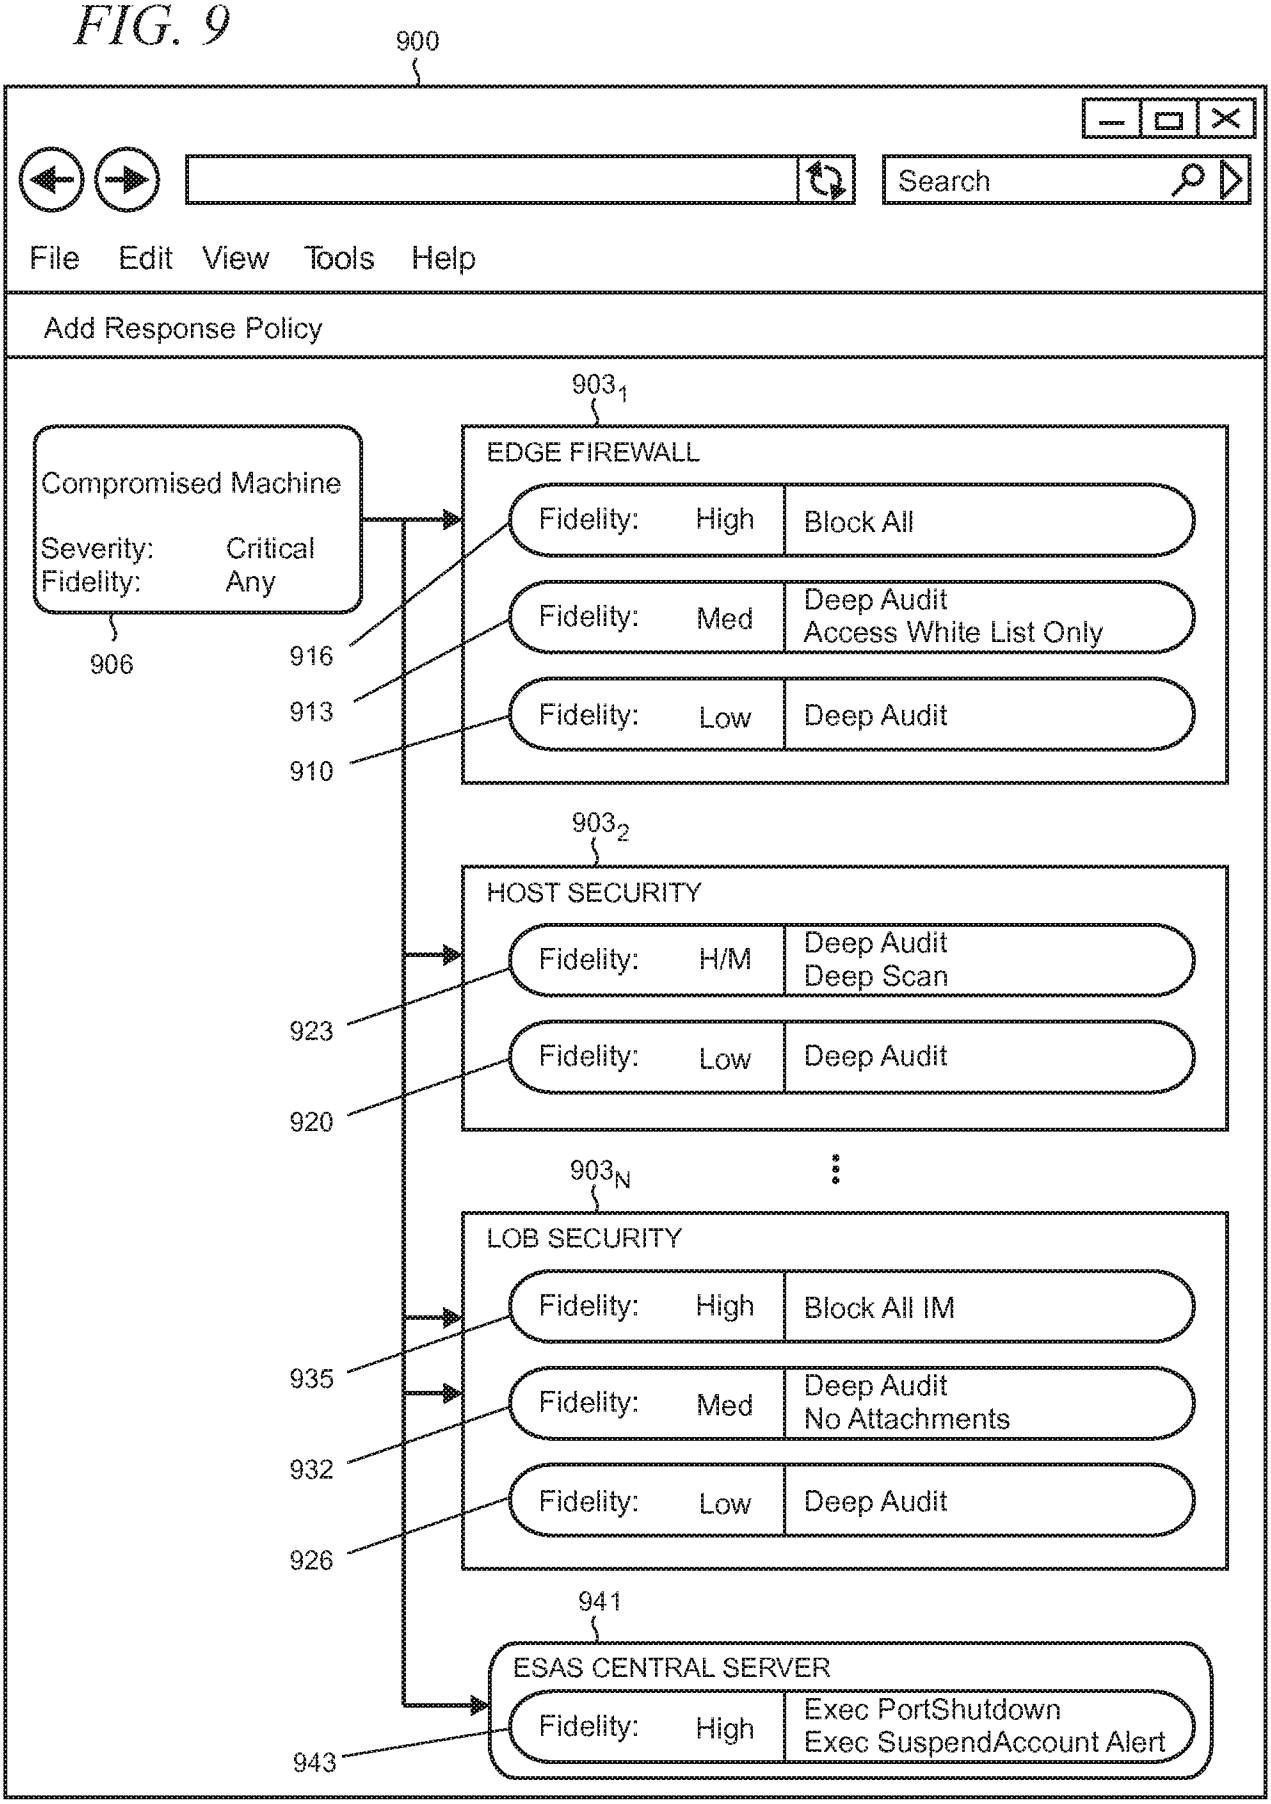

900
Search
File Edit View Tools Help
Add Response Policy
903_1
EDGE FIREWALL
Compromised Machine
Severity: Critical
Fidelity: Any
906
916
913
910
Fidelity: High
Block All
Fidelity: Med
Deep Audit
Access White List Only
Fidelity: Low
Deep Audit
903_2
HOST SECURITY
923
920
Fidelity: H/M
Deep Audit
Deep Scan
Fidelity: Low
Deep Audit
903_N
LOB SECURITY
935
932
926
Fidelity: High
Block All IM
Fidelity: Med
Deep Audit
No Attachments
Fidelity: Low
Deep Audit
941
ESAS CENTRAL SERVER
943
Fidelity: High
Exec PortShutdown
Exec SuspendAccount Alert

FIG. 10

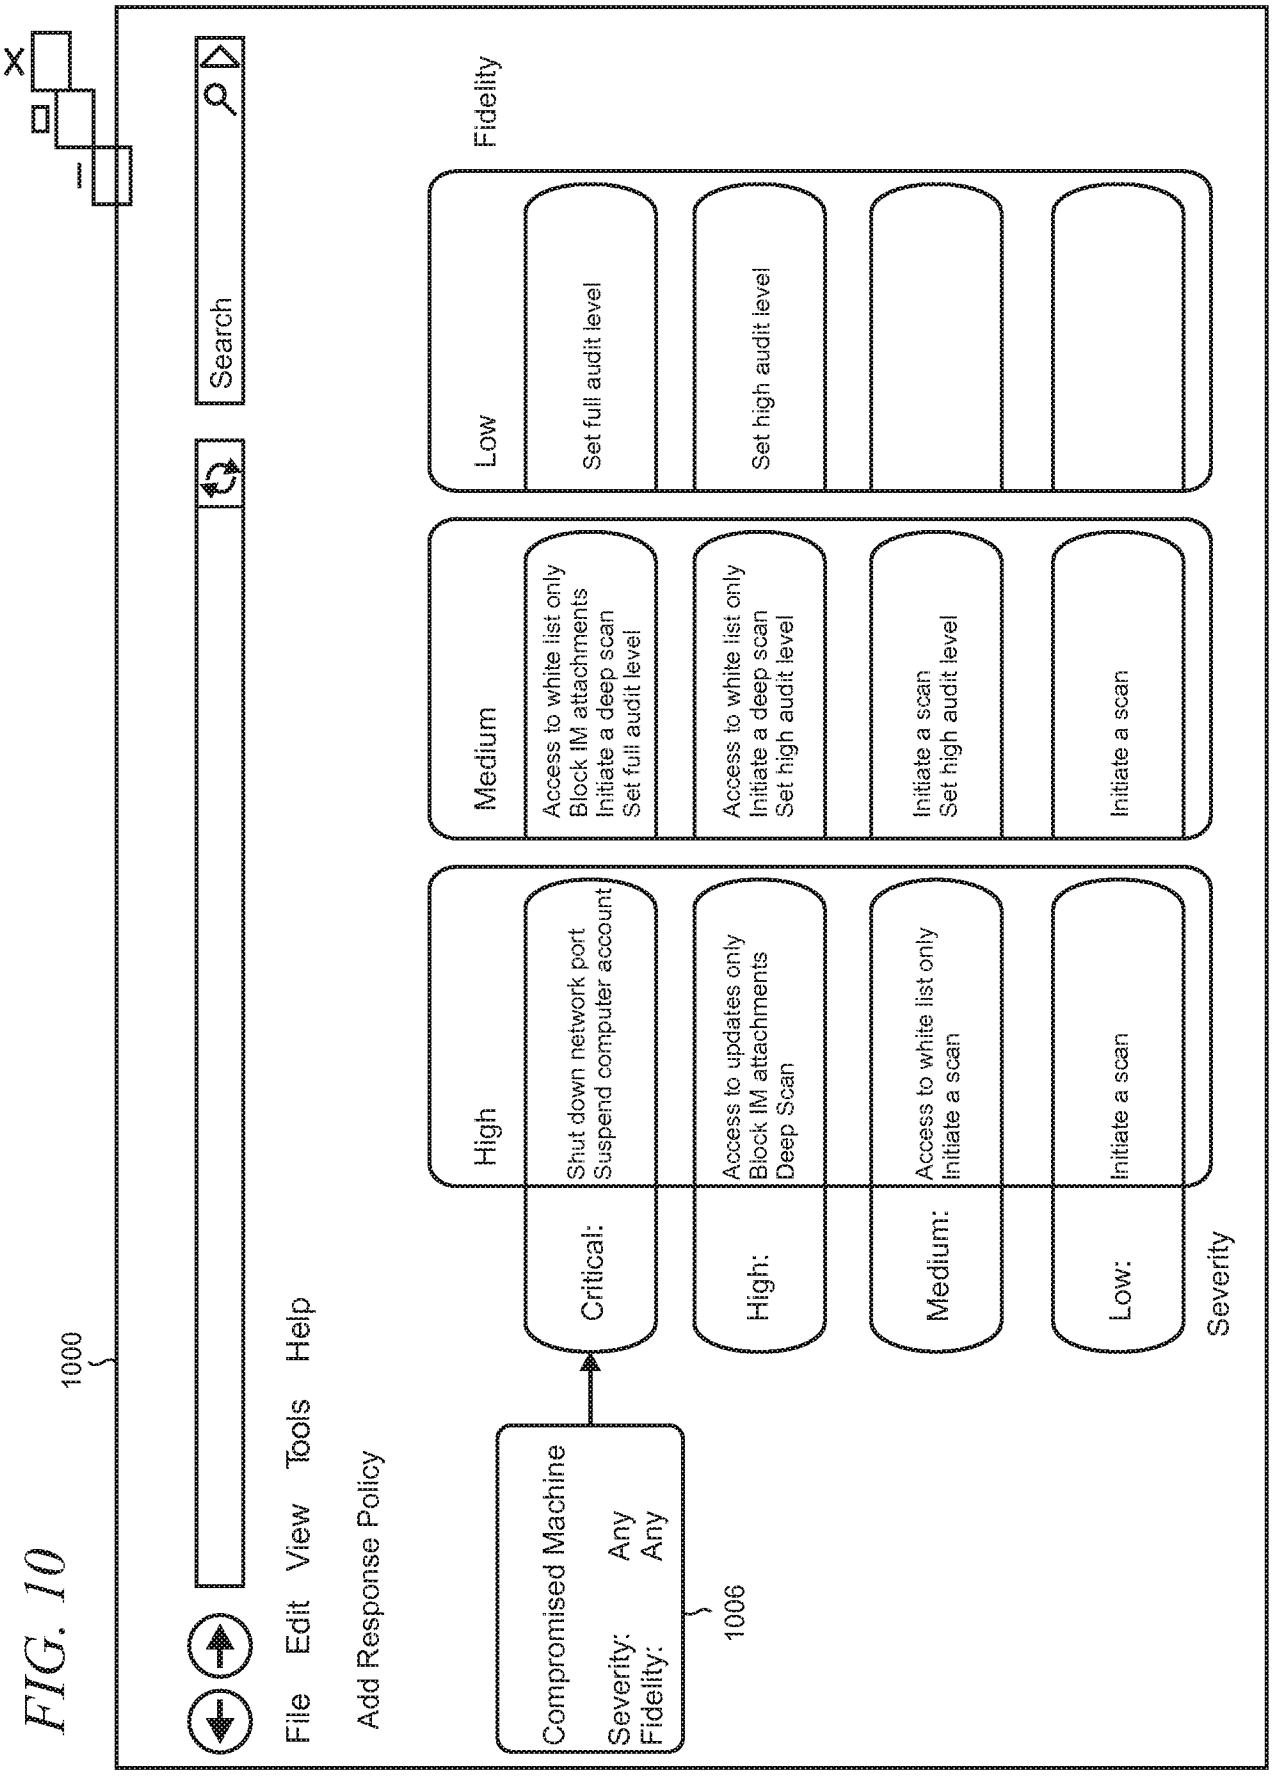
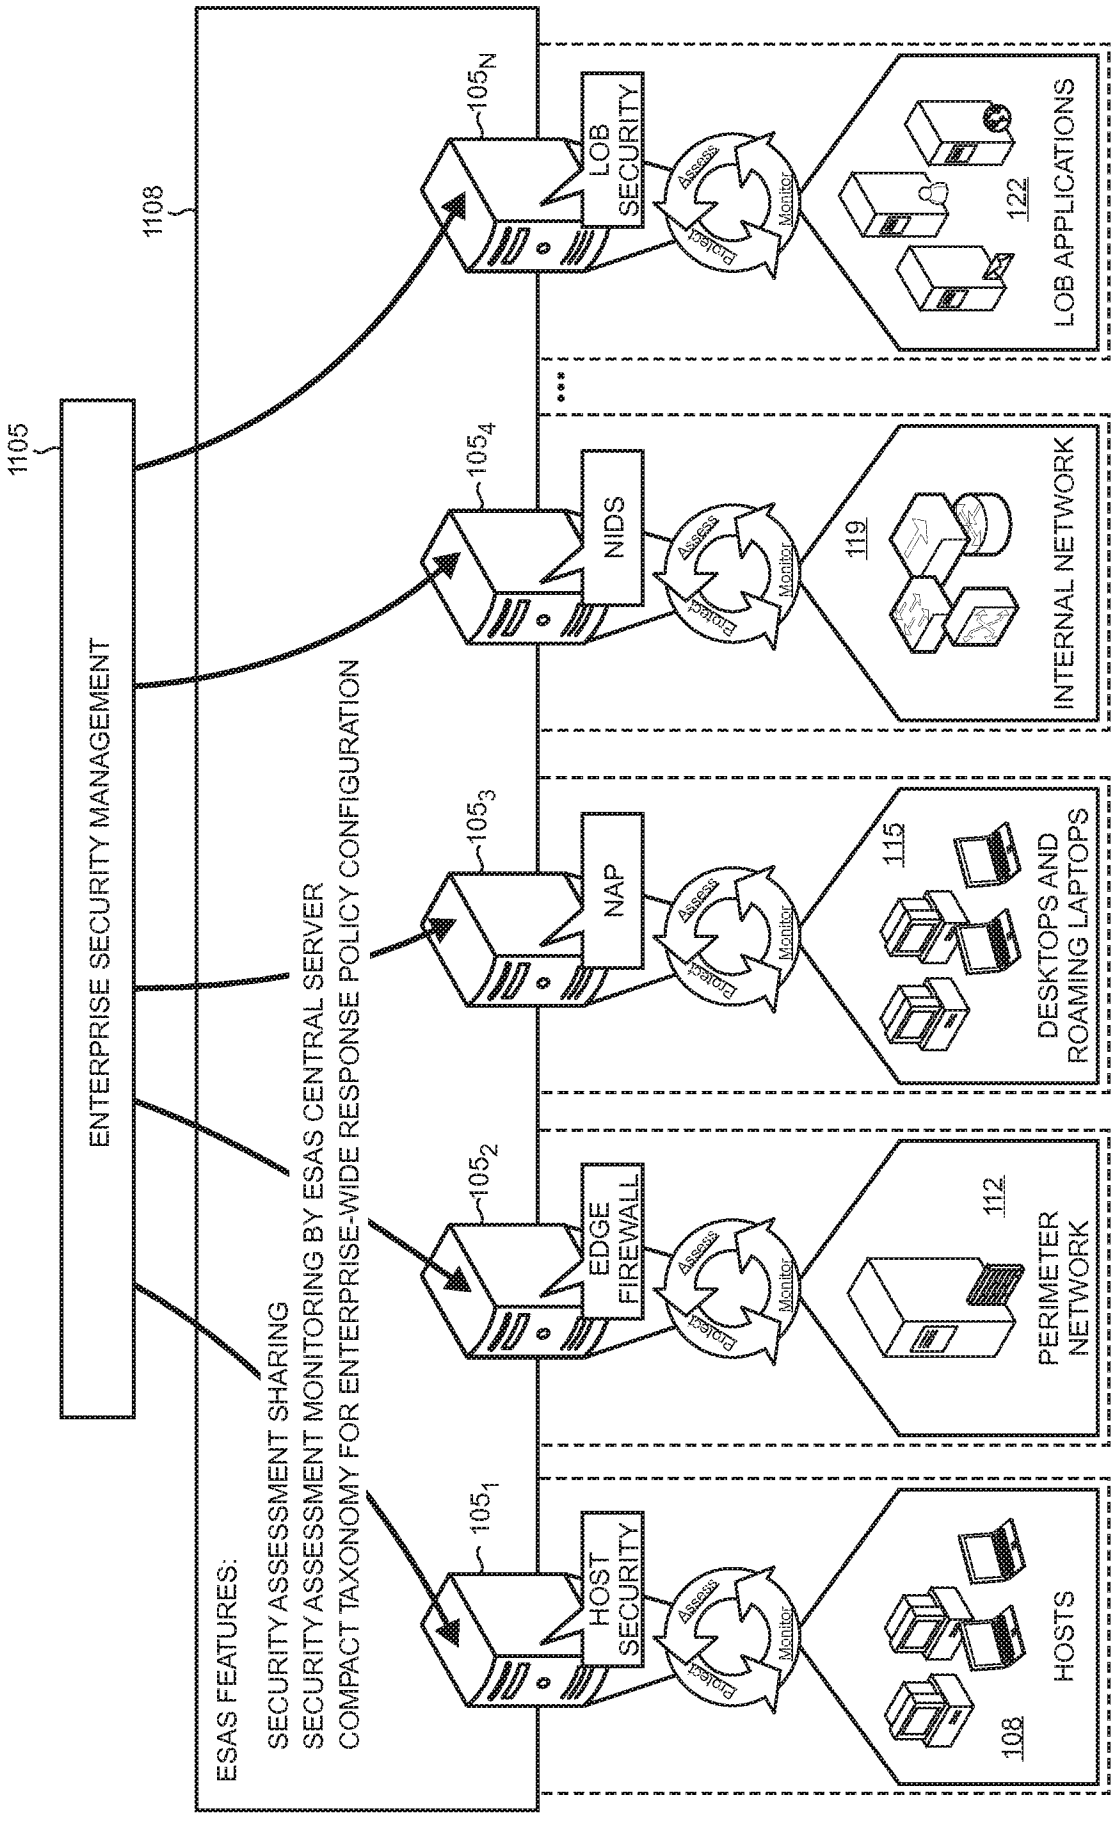

1000
File Edit View Tools Help
Search
Add Response Policy
Compromised Machine
Severity: Any
Fidelity: Any
1006
Critical:
High:
Medium:
Low:
Severity
High
Medium
Low
Fidelity
Shut down network port
Suspend computer account
Access to white list only
Block IM attachments
Initiate a deep scan
Set full audit level
Set high audit level
Access to updates only
Block IM attachments
Deep Scan
Initiate a scan
Set high audit level
Access to white list only
Initiate a scan
Initiate a scan

FIG. 11


1100
1105
ENTERPRISE SECURITY MANAGEMENT
1108
ESAS FEATURES:
SECURITY ASSESSMENT SHARING
SECURITY ASSESSMENT MONITORING BY ESAS CENTRAL SERVER
COMPACT TAXONOMY FOR ENTERPRISE-WIDE RESPONSE POLICY CONFIGURATION
105_1
HOST SECURITY
105_2
EDGE FIREWALL
105_3
NAP
105_4
NIDS
105_N
LOB SECURITY
Assess
Isolate
Monitor
108
HOSTS
112
PERIMETER NETWORK
115
DESKTOPS AND ROAMING LAPTOPS
119
INTERNAL NETWORK
122
LOB APPLICATIONS

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2008/057164

A. CLASSIFICATION OF SUBJECT MATTER

G06F 21/00(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)

IPC 8 G06F 1/00, 9/00, 15/16, 17/00, 17/30, 17/60, 21/00, H04L 9/00, 9/32

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 since 1975

Japanese Utility models and applications for Utility models since 1975

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

e-KIPASS(KIPO internal) "security","semantic","taxonomy","share"

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 03/093960 A2 (SCHLUMBERGER OMNES, INC.) 13 NOVEMBER 2003 See abstracts, claims 1 - 17, figures 1 - 4	1 - 20
A	US 6,986,060 B1 (Daniel ManHung Wong) 10 JANUARY 2006 See abstracts, claims 1 - 4, figures 1 - 4	1 - 20
A	US 6,925,443 B1 (Charlie C. Bagget, Jr., et al.) 02 AUGUST 2005 See abstracts, claims 1 - 4, figure 1	1 - 20
A	Gonzalo Alvarez, et al., 'A new taxonomy of Web attacks suitable for efficient encoding', In: Computer & Security, July 2003, Vol.22, Issue.5, P. 435-449	1 - 20
A	Douglas Harris, et al., 'Standards for secure data sharing across organizations', In: Computer Standards & Interfaces, January 2007, Vol.29, Issue 1, P. 86-96	1 - 20

☐ 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	Date of mailing of the international search report
22 AUGUST 2008 (22.08.2008)	**22 AUGUST 2008 (22.08.2008)**
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 KANG, YOON SEOK Telephone No. 82-42-481-5722

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2008/057164

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2003/093960 A2	13.11.2003	AU 2003-243177 A1	17.11.2003
		BR 200304653 A	21.12.2004
		CN 1714327 A	28.12.2005
		EP 1565797 A2	24.08.2005
		PA 04010698 A	17.05.2005
		RU 2004134589 A	27.06.2005
		US 2004-010709 A1	15.01.2004
		US 7290275 BB	30.10.2007
		WO 2003-093960 C1	21.05.2004
		WO 2003-093960 C2	30.06.2005
		ZA 200409477 A	08.06.2005
US 6986060 B1	10.01.2006	NONE	
US 6925443 B1	02.08.2005	AU 2001-53901 A1	07.11.2001
		WO 2001-82205A 1	01.11.2001