



- (51) **International Patent Classification:**
G06F 11/00 (2006.01)
- (21) **International Application Number:**
PCT/US2012/063569
- (22) **International Filing Date:**
5 November 2012 (05.11.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/288,905 3 November 2011 (03.11.2011) US
13/288,917 3 November 2011 (03.11.2011) US
- (71) **Applicant:** CYPHORT, INC. [US/US]; 2025 Gateway Place, Suite 300, San Jose, CA 95110 (US).
- (72) **Inventors:** GOLSHAN, Ali; 3595 Granada Avenue, Apartment 416, Santa Clara, CA 95051 (US). BINDER, James, S.; 6342 Skywalker Drive, San Jose, CA 95135 (US).
- (74) **Agents:** KLOKE, Daniel, C. et al.; Sheppard, Mullin, Richter & Hampton LLP, 379 Lytton Avenue, Palo Alto, CA 94301 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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

(54) **Title:** SYSTEMS AND METHODS FOR VIRTUALIZED MALWARE DETECTION

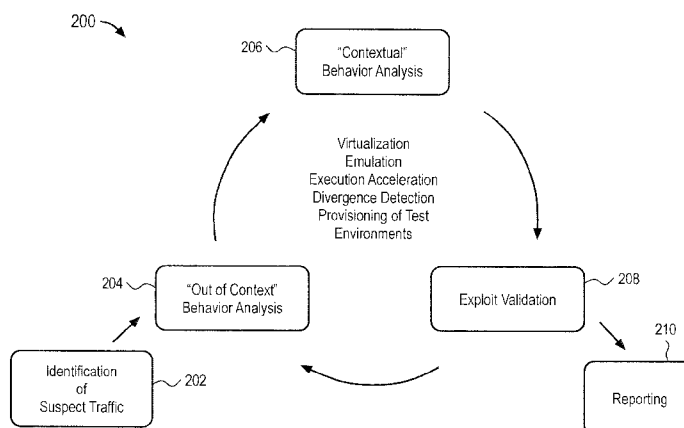


FIG. 2

(57) **Abstract:** Systems and methods for virtualized malware enabled detection are described. In some embodiments, a method comprises intercepting an object provided from a first digital device, determining one or more resources the object requires, instantiating a virtual environment with the one or more resources, processing the object within the virtual environment, tainting operations of the object within the virtual environment, monitoring the operations of the object, identifying an additional resource of the object while processing that is not provided in the virtual environment, re-instantiating the virtual environment with the additional resource, monitoring the operations of the object while processing within the re-instantiated virtual environment, identifying untrusted actions from the monitored operations, and generating a report identifying the operations and the untrusted actions of the object.

SYSTEMS AND METHODS FOR VIRTUALIZED MALWARE DETECTION**BACKGROUND****1. Field of the Invention**

[001] The present invention(s) generally relate to malware detection. More particularly, the invention(s) relate to systems and methods for virtualization and emulation assisted malware detection.

2. Description of Related Art

[002] Malware and advanced persistent attacks are growing in number as well as damage. In 2010, the rise of targeted attacks included armored variations of Conficker.D and Stuxnet (which was referred to as the most advanced piece of malware ever created). Targeted attacks on Google, Intel, Adobe, Boeing, and an estimated 60 others have been extensively covered in the press. The state of the art security defenses have proved ineffective.

[003] Cyber-criminals conduct methodical reconnaissance of potential victims to identify traffic patterns and existing defenses. Very sophisticated attacks involve multiple “agents” that individually appear to be legitimate traffic, then remain persistent in the target’s network. The arrival of other agents may also be undetected, but when all are in the target network, these agents can work together to compromise security and steal targeted information. Legacy security solutions use a structured process (e.g., signature and heuristics matching) or analyze agent behavior in an isolated context, without the ability to detect future coordinated activity. As a result, legacy security solutions are not able to detect sophisticated malware that is armored, component based, and/or includes different forms of delayed execution.

SUMMARY OF THE INVENTION

[004] Systems and methods for virtualized malware detection are described. In some embodiments, a method comprises intercepting an object provided from a first digital device to a second digital device, determining one or more resources the object requires when the object is executed, instantiating a virtual environment with the one or more resources, processing the object within the virtual environment, tainting operations of the object within the virtual environment, monitoring the operations of the object while processing within the virtual environment, identifying an additional resource of the object while processing that is not provided in the virtual environment, re-instantiating the virtual environment with the additional resource as well as the one or more resources, monitoring the operations of the object while processing within the re-instantiated virtual environment, identifying untrusted actions from the monitored operations, and generating a report identifying the operations and the untrusted actions of the object.

[005] The object may comprise an executable file, a batch file, or a data file.

[006] The method may further comprise performing a heuristic process on the object and determining the one or more resources the object requires based on the result of the heuristic process. Determining the one or more resources the object may be based on metadata associated with the object. The one or more resources may include one or more applications.

[007] Generating the report identifying the operations and the untrusted actions of the object may comprise generating a signature to be used to detect malware. In some embodiments, generating the report identifying the operations and the untrusted actions of the object may comprise identifying a vulnerability in an application based on the operations and the untrusted actions of the object.

[008] Re-instantiating the virtual environment with the additional resource as well as the one or more resources may comprise instantiating a second instance of a virtual environment with at least one resource that is different than a resource available in the prior virtual environment. Further, the method may comprise comparing identified monitored operations of the prior virtual environment to operations monitored in the second instance of the virtual environment. Generating the report may comprise generating the report based, at least in part, on the comparison.

[0009] The method may further comprise increasing or decreasing a clock signal within the virtual environment. In some embodiments, the method may comprise logging a state of the virtual environment while monitoring the operations of the object. Further, re-instantiating the virtual environment with the additional resource as well as the one or more resources may comprise halting the virtual environment and re-instantiating the virtual environment with the logged state.

[0010] An exemplary system may comprise a collection module, a virtualization module, a control module, and a report module. The collection module may be configured to receive an object provided from a first digital device to a second digital device. The virtualization module may be configured to instantiate a virtual environment with the one or more resources, to process the object within the virtual environment, to identify an additional resource of the object while processing that is not provided in the virtual environment, re-instantiate the virtual environment with the additional resource as well as the one or more resources, and to taint operations of the object within the virtual environment. The control module may be configured to determine one or more resources the object requires when the object is processed, to monitor the operations of the object while processing within the virtual environment, to monitor the operations of the object while processing within the re-instantiated virtual environment, and to identify untrusted actions from the monitored operations. The report module may be configured to generate a report identifying the operations and the untrusted actions of the object.

[0011] An exemplary computer readable medium may comprise instructions. The instructions may be executable by a processor for performing a method. The method may comprise intercepting an object provided from a first digital device to a second digital device, determining one or more resources the object requires when the object is executed, instantiating a virtual environment with the one or more resources, processing the object within the virtual environment, tainting operations of the object within the virtual environment, monitoring the operations of the object while processing within the virtual environment, identifying an additional resource of the object while processing that is not provided in the virtual environment, re-instantiating the virtual environment with the additional resource as well as the one or more resources, monitoring the operations of the object while processing within the re-instantiated virtual environment, identifying untrusted actions from the

monitored operations, and generating a report identifying the operations and the untrusted actions of the object.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] FIG. 1 is a diagram of an environment in which some embodiments may be practiced.

[0013] FIG. 2 is a flow diagram of an exemplary process for detection of malware and subsequent reporting in some embodiments.

[0014] FIG. 3 is a block diagram of an exemplary security server in some embodiments.

[0015] FIG. 4 is a conceptual block diagram of a virtualization module in some embodiments.

[0016] FIG. 5 is a block diagram of an exemplary virtualization module in some embodiments.

[0017] FIG. 6 is an exemplary virtualization environment for detection of malware in some embodiments.

[0018] FIG. 7 is a flow diagram of an exemplary malware detection method.

[0019] FIG. 8 is a flow diagram of an exemplary method of controlling a virtualization environment to detect malware.

[0020] FIG. 9 is a flow diagram of an exemplary model to detect malware through multiple virtualization environments.

[0021] FIG. 10 is a block diagram of an exemplary digital device.

DETAILED DESCRIPTION OF THE INVENTION

[0022] Some embodiments of systems and methods described herein describe appliance-based solutions to protect enterprises, governments, and cloud infrastructures against targeted sophisticated attacks with corporate espionage or possibly cyber warfare objectives. By watching patterns of abnormal traffic, various systems and methods described herein may predict interactions, identify vulnerabilities, and predictably deny particular protocols, data, or network paths to developing malware.

[0023] An exemplary system comprises a heuristics engine, an instrumented execution infrastructure, and an intelligent engine. The heuristics engine may identify payloads that require further static and dynamic analysis. The dynamic and instrumented execution infrastructure may combine both virtualization and emulation environments. The environments may be constantly updated dynamically to enable “suspect” traffic to execute to its fullest extent through divergence detection and distributed interaction correlation. The intelligent engine may exchange and cross-reference data between “on the fly” spawned virtual environments and emulated environments allowing, for example, the implementation of such resources as modified nested page tables. As a result, the virtualization environment may recreate all or part of the end-user environment as well as a fully optimized environment to extract the full execution and behavior of potential malware. Contextual environment may also be created to allow analysis of targeted malware built with armoring capabilities such as anti-virtualization, or anti-debugging technologies.

[0024] FIG. 1 is a diagram of an environment 100 in which some embodiments may be practiced. Systems and methods embodied in the environment 100 may detect malicious activity, identify malware, identify exploits, take preventive action, generate signatures, generate reports, determine malicious behavior, determine targeted information, recommend steps to prevent attack, and/or provide recommendations to improve security. The environment 100 comprises a data center network 102 and a production network 104 that communicate over a communication network 106. The data center network 102 comprises a security server 108. The production network 104 comprises a plurality of end user devices 110. The security server 108 and the end user devices 110 may comprise digital devices. A digital device is any device with a processor and memory. An embodiment of a digital device is depicted in FIG. 10.

[0025] The security server 108 is a digital device configured to identify malware and/or suspicious behavior by running virtualized and emulated environments and monitoring behavior of suspicious data within the virtualized and emulated environments. In various embodiments, the security server 108 receives suspicious data from one or more data collectors. The data collectors may be resident within or in communication with network devices such as Intrusion Prevention System (IPS) collectors 112a and 112b, firewalls 114a and 114b, ICAP/WCCP collectors 116, mailer mail plug-in collectors 118, switch collectors 120, and/or access points 124. Those skilled in the art will appreciate that a collector and a network device may be two separate digital devices (e.g., see F/W collector and IDS collector).

[0026] In various embodiments, data collectors may be at one or more points within the communication network 106. A data collector, which may include a tap or span port (e.g., span port / IDS at switch 120) for example, is configured to intercept network data from a network. The data collector may be configured to identify suspicious data. Suspicious data is any data collected by the data collector that has been flagged as suspicious by the data collector and/or any data that is to be processed within the virtualization environment.

[0027] The data collectors may filter the data before flagging the data as suspicious and/or providing the collected data to the security server 108. For example, the data collectors may filter out plain text but collect executables or batches. Further, in various embodiments, the data collectors may perform intelligent collecting. For example, data may be hashed and compared to a whitelist. The whitelist may identify data that is safe. In one example, the whitelist may identify digitally signed data or data received from a known trusted source as safe. Further, the whitelist may identify previously received information that has been determined to be safe. If data has been previously received, tested within the environments, and determined to be sufficiently trustworthy, the data collector may allow the data to continue through the network. Those skilled in the art will appreciate that the data collectors (or agents associated with the data collectors) may be updated by the security server 108 to help the data collectors recognize sufficiently trustworthy data and to take corrective action (e.g., quarantine and alert an administrator) if untrustworthy data is recognized. In some embodiments, if data is not identified as safe, the data collectors may flag the data as suspicious for further assessment.

[0028] Those skilled in the art will appreciate that one or more agents or other modules may monitor network traffic for common behaviors and may configure a data collector to collect data when data is directed in a manner that falls outside normal parameters. For example, the agent may determine or be configured to appreciate that a computer has been deactivated, a particular computer does not typically receive any data, or data received by a particular computer typically comes from a limited number of sources. If data is directed to a digital device in a manner that is not typical, the data collector may flag such data as suspicious and provide the suspicious data to the security server 108.

[0029] Network devices include any device configured to receive and provide data over a network. Examples of network devices include, but are not limited to, routers, bridges, security appliances, firewalls, web servers, mail servers, wireless access points (e.g., hotspots), and switches. In some embodiments, network devices include IPS collectors 112a and 112b, firewalls 114a and 114b, Internet content adaptation protocol (ICAP)/ web cache communication protocol (WCCP) servers 116, devices including milter mail plug-ins 118, switches 120, and/or access points 124.

[0030] The IPS collectors 112a and 112b may include any anti-malware device including IPS systems, intrusion detection and prevention systems (IDPS), or any other kind of network security appliances.

[0031] The firewalls 114a and 114b may include software and/or hardware firewalls. In some embodiments, the firewalls 114a and 114b may be embodied within routers, access points, servers (e.g., web servers), or appliances.

[0032] ICAP/WCCP servers 116 include any web server or web proxy server configured to allow access to a network and/or the Internet. Network devices including milter mail plug-ins 118 may include any mail server or device that provides mail and/or filtering functions and may include digital devices that implement milter, mail transfer agents (MTAs), sendmail, and postfix, for example.

[0033] Switches 120 include any switch or router. In some examples, the data collector may be implemented as a TAP, SPAN port, and/or intrusion detection system (IDS). Access points 124 include any device configured to provide wireless connectivity with one or more other digital devices.

[0034] The production network 104 is any network that allows one or more end user devices 110 to communicate over the communication network 106. The communication network 106 is any network that may carry data (encoded, compressed, and/or otherwise) from one digital device to another. In some examples, the communication network 106 may comprise a LAN and/or WAN. Further, the communication network 106 may comprise any number of networks. In some embodiments, the communication network 106 is the Internet.

[0035] FIG. 1 is exemplary and does not limit systems and methods described herein to the use of only those technologies depicted. For example, data collectors may be implemented in any web or web proxy server and is not limited to only the servers that implement ICAP and/or WCCP. Similarly, collectors may be implemented in any mail server and is not limited to mail servers that implement milner. Data collectors may be implemented at any point in one or more networks.

[0036] Those skilled in the art will appreciate that although FIG. 1 depicts a limited number of digital devices, collectors, routers, access points, and firewalls, there may be any kind and number of devices. For example, there may be any number security servers 108, end user devices 110, IPS collectors 112a and 112b, firewalls 114a and 114b, ICAP/WCCP collectors 116, milner mail plug-ins 118, switches 120, and/or access points 124. Further, there may be any number of data center networks 102 and/or production networks 104.

[0037] FIG. 2 is a flow diagram of an exemplary process 200 for detection of malware and subsequent reporting in some embodiments. In step 202, suspect traffic is identified. In various embodiments, any network device may be used to monitor and/or collect network traffic for further assessment. In various embodiments, the network device and/or another digital device (e.g., the security server 108) applies heuristics and/or rules (e.g., comparison of data to a whitelist and/or a blacklist) to identify suspicious data. Those skilled in the art will appreciate that any technique may be used to flag network traffic as suspicious. For example, the security server 108 may flag data as suspicious if the data is directed towards a known infected computer, a disabled account, or any untrustworthy destination. Further, for example, the security server 108 may flag data as suspicious if the data came from a suspected source of malware or a source that is known to be untrustworthy (e.g., a previously identified botnet server). In another example, the data collector and/or agent associated with the data collector

may perform packet analysis to identify suspicious characteristics in the collected data including the header, footer, destination IP, origin IP, payload and the like.

[0038] In step 204, suspect data and/or suspect processes are tested in one or more virtualization environments for “out of context” behavior analysis of the suspicious data and suspect processes. In some embodiments, the suspect data and/or processes are initially virtualized in a set of virtualization environments. Each different virtualization environment may be provisioned differently (e.g., each different virtualization environment may comprise different resources). The initial set of resources for a virtualization environment may be predetermined based on common resources required for processing the data and/or metadata associated with the data. If the suspect data and/or suspect process are determined to be behaving suspiciously in the virtualization environment, the suspect data and/or process may also be processed in an emulation environment as discussed here.

[0039] In various embodiments, the suspect data and/or process is analyzed with multiple virtualization environments to extend predictive analysis to distributed and application interactions as described further herein. The suspect data and/or process may be identified as malware or may behave in an untrusted manner in the virtualized environment. In order to further assess the data and/or process, the data and/or process may be processed in a plurality of different virtualization environments with different resources and different limitations. Those skilled in the art will appreciate that the suspicious data and/or process may or may not be further tested after the initial set of environments.

[0040] In step 206, contextual behavioral analysis is conducted on the suspect data and suspect processes using one or more emulation environments. In some embodiments, if the suspicious data acts suspiciously in one or more virtualization environments (e.g., halting execution without performing functions, storing data without using the data, and the like), the data is processed in one or more emulation environments. The emulation environment may be provisioned based on commonly needed resources, metadata associated with the suspicious data, and/or resources identified as needed during processing of the suspicious data within the virtualization environment. The suspicious may have direct access to memory data in the emulation environment. The behavior of the suspicious data may be monitored within the emulation environment.

[0041] In step 208, exploits are identified and validated based on the behavior of the suspect data or suspect process in the environments. For example, the virtualization and/or emulation environments may be provisioned with various applications and operating systems in order to monitor the behavior of the suspect data or suspect process. As a result, the environments may test suspect data or suspect processes against network resources and/or applications to determine vulnerabilities and malicious actions. As a result, the assessment of the suspect data and/or process may extend predictive analysis to applications for a fuller or complete identification of targeted vulnerabilities.

[0042] In some embodiments, when a divergence is detected between the behavior of suspect data and/or process in the virtualization environment and the emulation environment, the virtualization environment may be dynamically re-instantiated and re-provisioned (e.g., the process returns to step 204 with the re-instantiated and/or re-provisioned virtualization environment(s)). Data from the emulation environment (e.g., responses from within the emulation environment) may be injected into the re-provisioned virtualization environment at or close to the time of divergence to enable further execution of the suspect data and assessment of related data.

[0043] In step 210, a report is generated that may identify threats and vulnerabilities based on the monitored behaviors of the suspect data and the suspect processes within the testing environments. In various embodiments, the report may include a description of exploits, vulnerabilities of applications or operating systems, behaviors of the suspect data, payloads associated with the suspect data, command and control protocols, and probable targets of the suspect data (e.g., what valuable information the suspicious data was attempting to steal). Further, the report may include heuristics, additions to whitelists, additions to blacklists, statistics, or signatures designed to detect the suspect data.

[0044] In various embodiments, the exemplary process 200 may be used to detect distributed attacks characteristic of advanced persistent threats. One exemplary scenario of a distributed attack is that an attacker may send a package to be stored in a specific location in the target computer. The package and the act of storing the package may be benign. The attacker may, over time, subsequently send an attack program. Without the previously stored package, the attack program may also appear benign and may not be detectable as malware by preexisting security solutions. Once the attack program retrieves the previously stored package, however,

the attack program may attack the target system (e.g., exploit a vulnerability in the operating system to take over the target computer or copy valuable data).

[0045] In various embodiments, the security server 108 may first receive and test a package in at least one of the different environments. A report or other characteristic of the storage (e.g., the location of the stored data and the stored data) may be logged and stored for later testing within the environments. For example, an object that stores a package in memory but does not refer to the package after storage may be deemed to be suspicious. As such, the object may be tested in a variety of different environments and/or the package may be stored (e.g., in a protected long term storage memory such as a hard drive). When the security server 108 subsequently receives the attack program and, during testing, notes that the attack program is suspiciously checking a particular location in memory for data, the security server 108 may recognize that the previously stored package was stored in that particular location of memory. The security server 108 may retrieve the previously received package and store the package within the location in memory in one of the environments and retest the attack program. If the attack program acts maliciously after receiving the package, the security server 108 may generate a report (e.g., information, signature file, heuristic, and/or the like) to identify the package as well as the attack program in order to protect against similar attacks. Moreover, the security server 108 may generate a report identifying the exploited vulnerability so that the vulnerability may be corrected (e.g., the operating system patched or upgraded to correct the exploit). The security server 108 may also generate a report identifying the targeted information (e.g., a password file or file of credit card numbers) so that corrective action may be taken (e.g., move the file or encrypt the information).

[0046] FIG. 3 is a block diagram of an exemplary security server 108 in some embodiments. In various embodiments, the security server 108 leverages both virtualization and emulation systems and methods to detect malware anti-virtualization protections and accelerate “on-demand” virtualized environments for faster prediction. The security server 108 comprises a collection module 302, a data flagging module 304, a virtualization module 306, an emulation module 308, a control module 310, a reporting module 312, a signature module 314, and a quarantine module 316.

[0047] The collection module 302 is configured to receive network data (e.g., potentially suspicious data) from one or more sources. Network data is data that is provided on a network

from one digital device to another. The collection module 302 may flag the network data as suspicious data based on, for example, whitelists, blacklists, heuristic analysis, statistical analysis, rules, and/or atypical behavior. In some embodiments, the sources comprise data collectors configured to receive network data. For example, firewalls, IPS, servers, routers, switches, access points and the like may, either individually or collectively, function as or include a data collector. The data collector may forward network data to the collection module 302.

[0048] In some embodiments, the data collectors filter the data before providing the data to the collection module 302. For example, the data collector may be configured to collect or intercept data that includes executables and batch files. In some embodiments, the data collector may be configured to follow configured rules. For example, if data is directed between two known and trustworthy sources (e.g., the data is communicated between two device on a whitelist), the data collector may not collect the data. In various embodiments, a rule may be configured to intercept a class of data (e.g., all MS Word documents that may include macros or data that may comprise a script). In some embodiments, rules may be configured to target a class of attack or payload based on the type of malware attacks on the target network in the past. In some embodiments, the security server 108 may make recommendations (e.g., via the reporting module 312) and/or configure rules for the collection module 302 and/or the data collectors. Those skilled in the art will appreciate that the data collectors may comprise any number of rules regarding when data is collected or what data is collected.

[0049] In some embodiments, the data collectors located at various positions in the network may not perform any assessment or determination regarding whether the collected data is suspicious or trustworthy. For example, the data collector may collect all or a portion of the network data and provide the collected network data to the collection module 302 which may perform filtering.

[0050] The data flagging module 304 may perform one or more assessments to the collected data received by the collection module 302 and/or the data collector to determine if the intercepted network data is suspicious. The data flagging module 304 may apply rules as discussed herein to determine if the collected data should be flagged as suspicious. In various embodiments, the data flagging module 304 may hash the data and/or compare the data to a

whitelist to identify the data as acceptable. If the data is not associated with the whitelist, the data flagging module 304 may flag the data as suspicious.

[0051] In various embodiments, collected network data may be initially identified as suspicious until determined otherwise (e.g., associated with a whitelist) or heuristics find no reason that the network data should be flagged as suspicious. In some embodiments, the data flagging module 304 may perform packet analysis to look for suspicious characteristics in the header, footer, destination IP, origin IP, payload, and the like. Those skilled in the art will appreciate that the data flagging module 304 may perform a heuristic analysis, a statistical analysis, and/or signature identification (e.g., signature-based detection involves searching for known patterns of suspicious data within the collected data's code) to determine if the collected network data is suspicious.

[0052] The data flagging module 304 may be resident at the data collector, at the security server 108, partially at the data collector, partially at the security server 108, or on a network device. For example, a router may comprise a data collector and a data flagging module 304 configured to perform one or more heuristic assessments on the collected network data. If the collected network data is determined to be suspicious, the router may direct the collected data to the security server 108.

[0053] In various embodiments, the data flagging module 304 may be updated. In one example, the security server 108 may provide new entries for a whitelist, entries for a blacklist, heuristic algorithms, statistical algorithms, updated rules, and/or new signatures to assist the data flagging module 304 to determine if network data is suspicious. The whitelists, entries for whitelists, blacklists, entries for blacklists, heuristic algorithms, statistical algorithms, and/or new signatures may be generated by one or more security servers 108 (e.g., via the reporting module 312).

[0054] The virtualization module 306 and emulation module 308 may analyze suspicious data for untrusted behavior (e.g., malware or distributed attacks). The virtualization module 306 is configured to instantiate one or more virtualized environments to process and monitor suspicious data. Within the virtualization environment, the suspicious data may operate as if within a target digital device. The virtualization module 306 may monitor the operations of the suspicious data within the virtualization environment to determine that the suspicious data is probably trustworthy, malware, or requiring further action (e.g., further monitoring in one or

more other virtualization environments and/or monitoring within one or more emulation environments). In various embodiments, the virtualization module 306 monitors modifications to a system, checks outbound calls, and checks tainted data interactions.

[0055] In some embodiments, the virtualization module 306 may determine that suspicious data is malware but continue to process the suspicious data to generate a full picture of the malware, identify the vector of attack, determine the type, extent, and scope of the malware's payload, determine the target of the attack, and detect if the malware is to work with any other malware. In this way, the security server 108 may extend predictive analysis to actual applications for complete validation. A report may be generated (e.g., by the reporting module 312) describing the malware, identify vulnerabilities, generate or update signatures for the malware, generate or update heuristics or statistics for malware detection, and/or generate a report identifying the targeted information (e.g., credit card numbers, passwords, or personal information).

[0056] In some embodiments, the virtualization module 306 may flag suspicious data as requiring further emulation if the data has suspicious behavior such as, but not limited to, preparing an executable that is not executed, performing functions without result, processing that suddenly terminates, loading data into memory that is not accessed or otherwise executed, scanning ports, or checking in specific portions of memory when those locations in memory may be empty. The virtualization module 306 may monitor the operations performed by or for the suspicious data and perform a variety of checks to determine if the suspicious data is behaving in a suspicious manner.

[0057] The emulation module 308 is configured to process suspicious data in an emulated environment. Those skilled in the art will appreciate that malware may require resources that are not available or may detect a virtualized environment. When malware requires unavailable resources, the malware may "go benign" or act in a non-harmful manner. In another example, malware may detect a virtualized environment by scanning for specific files and/or memory necessary for hypervisor, kernel, or other virtualization data to execute. If malware scans portions of its environment and determines that a virtualization environment may be running, the malware may "go benign" and either terminate or perform nonthreatening functions.

[0058] In some embodiments, the emulation module 308 processes data flagged as behaving suspiciously by the virtualization environment. The emulation module 308 may process the

suspicious data in a bare metal environment where the suspicious data may have direct memory access. The behavior of the suspicious data as well as the behavior of the emulation environment may be monitored and/or logged to track the suspicious data's operations. For example, the emulation module 308 may track what resources (e.g., applications and/or operating system files) are called in processing the suspicious data.

[0059] In various embodiments, the emulation module 308 records responses to the suspicious data in the emulation environment. If a divergence in the operations of the suspicious data between the virtualization environment and the emulation environment is detected, the virtualization environment may be configured to inject the response from the emulation environment. The suspicious data may receive the expected response within the virtualization environment and continue to operate as if the suspicious data was within the targeted digital device. This process is further described herein.

[0060] The control module 310 synchronizes the virtualization module 306 and the emulation module 308. In some embodiments, the control module 310 synchronizes the virtualization and emulation environments. For example, the control module 310 may direct the virtualization module 306 to instantiate a plurality of different virtualization environments with different resources. The control module 310 may compare the operations of different virtualization environments to each other in order to track points of divergence. For example, the control module 310 may identify suspicious data as operating in one manner when the virtualization environment includes Internet Explorer v. 7.0 or v. 8.0, but operating in a different manner when interacting with Internet Explorer v. 6.0 (e.g., when the suspicious data exploits a vulnerability that may be present in one version of an application but not present in another version).

[0061] The control module 310 may track operations in one or more virtualization environments and one or more emulation environments. For example, the control module 310 may identify when the suspicious data behaves differently in a virtualization environment in comparison with an emulation environment. Divergence analysis is when operations performed by or for suspicious data in a virtual environment is compared to operations performed by or for suspicious data in a different virtual environment or emulation environment. For example, the control module 310 may compare monitored steps of suspicious data in a virtual environment to monitored steps of the same suspicious data in an

emulation environment. The functions or steps of or for the suspicious data may be similar but suddenly diverge. In one example, the suspicious data may have not detected evidence of a virtual environment in the emulation environment and, unlike the virtualized environment where the suspicious data went benign, the suspicious data undertakes actions characteristic of malware (e.g., hijacks a formerly trusted data or processes).

[0062] When divergence is detected, the control module 310 may re-provision or instantiate a virtualization environment with information from the emulation environment (e.g., a page table including state information and/or response information further described herein) that may not be previously present in the originally instantiation of the virtualization environment. The suspicious data may then be monitored in the new virtualization environment to further detect suspicious behavior or untrusted behavior. Those skilled in the art will appreciate that suspicious behavior of an object is behavior that may be untrusted or malicious. Untrusted behavior is behavior that indicates a significant threat.

[0063] In some embodiments, the control module 310 is configured to compare the operations of each virtualized environment in order to identify suspicious or untrusted behavior. For example, if the suspicious data takes different operations depending on the version of a browser or other specific resource when compared to other virtualized environments, the control module 310 may identify the suspicious data as malware. Once the control module 310 identifies the suspicious data as malware or otherwise untrusted, the control module 310 may continue to monitor the virtualized environment to determine the vector of attack of the malware, the payload of the malware, and the target (e.g., control of the digital device, password access, credit card information access, and/or ability to install a bot, keylogger, and/or rootkit). For example, the operations performed by and/or for the suspicious data may be monitored in order to further identify the malware, determine untrusted acts, and log the effect or probable effect.

[0064] The reporting module 312 is configured to generate reports based on the processing of the suspicious data of the virtualization module 306 and/or the emulation module 308. In various embodiments, the reporting module 312 generates a report to identify malware, one or more vectors of attack, one or more payloads, target of valuable data, vulnerabilities, command and control protocols, and/or behaviors that are characteristics of the malware. The reporting module 312 may also make recommendations to safeguard information based on the attack

(e.g., move credit card information to a different digital device, require additional security such as VPN access only, or the like).

[0065] In some embodiments, the reporting module 312 generates malware information that may be used to identify malware or suspicious behavior. For example, the reporting module 312 may generate malware information based on the monitored information of the virtualization environment. The malware information may include a hash of the suspicious data or a characteristic of the operations of or for the suspicious data. In one example, the malware information may identify a class of suspicious behavior as being one or more steps being performed by or for suspicious data at specific times. As a result, suspicious data and/or malware may be identified based on the malware information without virtualizing or emulating an entire attack.

[0066] The optional signature module 314 is configured to store signature files that may be used to identify malware. The signature files may be generated by the reporting module 312 and/or the signature module 314. In various embodiments, the security server 108 may generate signatures, malware information, whitelist entries, and/or blacklist entries to share with other security servers. As a result, the signature module 314 may include signatures generated by other security servers or other digital devices. Those skilled in the art will appreciate that the signature module 314 may include signatures generated from a variety of different sources including, but not limited to, other security firms, antivirus companies, and/or other third-parties.

[0067] In various embodiments, the signature module 314 may provide signatures which are used to determine if network data is suspicious or is malware. For example, if network data matches the signature of known malware, then the network data may be classified as malware. If network data matches a signature that is suspicious, then the network data may be flagged as suspicious data. The malware and/or the suspicious data may be processed within a virtualization environment and/or the emulation environment as discussed herein.

[0068] The quarantine module 316 is configured to quarantine suspicious data and/or network data. In various embodiments, when the security server 108 identifies malware or probable malware, the quarantine module 316 may quarantine the suspicious data, network data, and/or any data associated with the suspicious data and/or network data. For example,

the quarantine module 316 may quarantine all data from a particular digital device that has been identified as being infected or possibly infected.

[0069] In some embodiments, the quarantine module 316 is configured to alert a security administrator or the like (e.g., via email, call, voicemail, or SMS text message) when malware or possible malware has been found.

[0070] In various embodiments, the security server 108 allows an administrator or other personnel to log into the security server 108. In one example, the security server 108 provides a graphical user interface or other user interface that authenticates a user (e.g., via digital signature, password, username, and the like). After the user is authenticated, the security server 108 may allow the user to view the processing of the virtualization module 306 and the emulation module 306 including infection vectors, and vulnerability vectors. The security server 108 may also provide the user with threshold reasoning which is further described regarding FIG. 4.

[0071] FIG. 4 is a conceptual block diagram 400 of a virtualization module in some embodiments. In various embodiments, different processes 402 may be virtualized within one or more virtualization environments 404. The virtualization environments execute on a host 406 that runs over hardware 408 that is isolated from the suspicious data and/or processes. The control module 310 may identify various results to identify when suspicious behavior is present (e.g., value X), in what sequence the suspicious behavior occurs (e.g., value Y) and what process (e.g., value Z).

[0072] For example, a particular process 402 may be intercepted and tested in a variety of different virtualization environments 404. Each virtualization environment 404 may operate on a host 406 (e.g., operating system and/or virtual machine software) that executes over a digital device's hardware 408. The functions of the tested process may be isolated from the host 406 and hardware 408. Suspicious or untrusted behavior may be identified within the virtualization. A time of exploitation may be identified as value X, an exploited sequence may be identified as value Y, and a process of exploitation may be identified as value Z.

[0073] The X, Y, Z values may form a description of suspicious data or the process which may be used to measure the threat against a threat matrix. In some embodiments, an administrator may store a threat threshold, based on the threat matrix depending upon the level

of risk that is acceptable. The threat matrix may be based on interactions with the operating system, time sequence, resources, or events. In some embodiments, the degree of malicious behavior may be determined based on a threat value (e.g., comprising a function including the X, Y, and Z values). In one example, the interactions with the OS, time sequences, types of interactions, and resources requested, may all be elements of the threat matrix. Once a threat value is determined, the threat value may be compared to a threat threshold to determine the degree of maliciousness and/or what actions will be taken. Those skilled in the art will appreciate that the threat threshold may be determined and/or generated based on an administrator's acceptable level of risk.

[0074] Time, sequence, and process values may be generated for each tested process or data. The time, sequence, and process values may be measured against the threshold using the threat matrix to determine a possible course of action (e.g., quarantine, generate a report, alert an administrator, or allow the process to continue unobstructed).

[0075] The X, Y, Z values may be compared to X, Y, Z values associated with the same suspicious data from the emulation environment. If the emulation environment values are different or divergent, further testing within the virtualization environment and/or the emulation environment may be required.

[0076] FIG. 5 is a block diagram of an exemplary virtualization module 306 in some embodiments. The virtualization module 306 may comprise a virtual machine module 502, a resource module 504, a monitor module 506, a taint module 508, a time module 510, a state module 512, and a state database 514.

[0077] The virtual machine module 502 is configured to generate one or more virtualization environments to process and monitor suspicious data. Those skilled in the art will appreciate that many different virtual machines may be used (e.g., virtual machines from VMWare or custom virtual machines).

[0078] The resource module 504 is configured to provision one or more virtualization environments with plug-ins or other resources. In various embodiments, plug-ins are modules build in the virtual and emulation environments that collect specific data sets from certain system components. This process may be chained to follow an execution through the system or may run in parallel if there is a threaded malicious or clean object.

[0079] In some embodiments, the resource module 504 provisions a virtualization environment with an initial set of resources (e.g., operating system, OS updates, applications, and drivers). In some embodiments, the resource module 504 provisions virtualization environments to include resources based on the destination of the suspicious data (e.g., the digital device targeted to receive the suspicious data), device images provisioned by information technology management, or metadata associated with the suspicious data. In some embodiments, the resource module 504 comprises a pre-processing module that determines specific requirements based on network meta-data to determine which plug-ins should be implemented within the virtualization environment and in what combination the plug-ins may be launched.

[0080] In some embodiments, the resource module 504 provisions a virtualization environment based on the suspicious data's similarity to malware or other suspicious data. In one example, the virtualization module 306 may scan and find that the suspicious data appears to be similar to previously tested suspicious data or malware. Subsequently, the resource module 504 may provision one or more virtualization environments to include resources with known vulnerabilities to monitor whether the suspicious data acts in a similarly untrusted manner.

[0081] In various embodiments, the resource module 504 provisions a virtualization environment based in part on metadata associated with the suspicious data. For example, the virtualization module 306 may receive or retrieve metadata associated with the suspicious data. The resource module 504 may determine, based on the metadata, that one or more applications are required for the suspicious data to function. Subsequently, the resource module 504 may provision one or more virtualization environments with the necessary applications and related support file (e.g., operating system, shared resources, or drivers).

[0082] Those skilled in the art will appreciate that multiple virtualized environments may be instantiated. Each of the virtualized environments may have one or more different resources. In one example, one virtualized environment may include Internet Explorer v. 6 while another virtualized environment may include Internet Explorer v. 7. Different virtualized environments may include, in some embodiments, different browser programs (e.g., Mozilla Firefox), different operating systems (e.g., Unix), and/or different drivers. The different virtualization environments may have similar applications or operating systems but different

versions or different patches or updates. In this way, the same suspicious data may be processed using different resources. If the suspect data behaves differently with one browser than with another, then there is evidence that the suspicious data may be malware.

[0083] In various embodiments, suspicious data is processed in a plurality of different virtualized environments where each of the different virtualized environments includes a limited number of differences. As a result, if malware is only effective in the presence of Internet Explorer v. 6.0 (i.e., there is a vulnerability in Internet Explorer v. 6.0 that the malware is programmed to exploit), then the malware's behavior as well as the exploit may be identified.

[0084] The control module 310 may provision the virtualization module 306. In some embodiments, the control module 310 may review metadata associated with the suspicious data to determine resources to be available in one or more virtualization environments. Those skilled in the art will appreciate that the metadata may come from a variety of sources. For example, some metadata may be apparent from the suspicious data such as a file extension or calls associated with the suspicious data. In some embodiments, the control module 310 may retrieve information regarding the suspicious data in order to provision the virtualization environment. For example, the control module 310 may determine that the suspicious data may be similar to other malware or suspicious data and provision one or more virtualized environments in a manner to see if the newly acquired suspicious data behaves in an untrusted manner.

[0085] The control module 310 may also provision the emulation module 308. In some embodiments, the control module 310 may review metadata associated with the suspicious data to determine resources to be available in one or more emulation environments. The control module 310 may also provision an emulation environment based on the provisioning of one or more virtualized environments. For example, the control module 310 may provision the emulation environment based on a virtualized environment where the suspicious data may have behaved abnormally (e.g., in an environment with a specific version of an operating system, the suspicious data scanned one or more areas of memory and then terminated further operations). The emulation environment may, in some embodiments, share similar resources as what was provided in a virtualization environment.

[0086] The virtualization module 306 and/or the collection module 302 may determine resource requirements of or for the suspicious data. In various embodiments, the virtualization module 306 receives metadata associated with the suspicious data to determine resources as described herein. For example, the metadata may indicate that the network data is an executable to be run in a Windows environment or the metadata may indicate that the network data is an executable file to be operated by a browser (e.g., a web application). The virtualization module 306 and/or the control module 310 may dynamically select a variety of resources to provision and instantiate a virtualization environment in order to process the network data and monitor actions.

[0087] In various embodiments, a resource may be missing from one, some, or all of the virtualized environments. For example, the suspicious data may require a different application to be able to execute. In some embodiments, the virtualization module 306 may halt a virtualization environment, dynamically provision the virtualization environment with the necessary resources, and re-instantiate the virtualized environment to monitor for changes in behavior of the suspicious data.

[0088] The monitor module 506 is configured to monitor the virtualization environments instantiated by the virtual machine module 502. In various embodiments, the monitor module 506 logs each step or function performed by or for the suspicious data within each virtualization environment. In various embodiments, the monitor module 506 logs each operation of the suspicious data, logs changes caused by the operation (e.g., what information is stored in memory and where in memory the information is stored), and logs at what time the operation occurred.

[0089] The monitor module 506 may compare the operations of the suspicious data in various virtualization environments during or after virtualization. When a divergence is identified between a virtualization environment and an emulation environment or between two virtualization environments, the monitor module 506 may generate a flag or track the results to identify if different operations perform untrusted actions.

[0090] The taint module 508 is configured to perform taint analysis and/or other techniques to identify and track operations provided by and for the suspect data. As a result, acts associated with the suspicious data, including executions by the suspect data and executions performed by an application or operating system for the suspect data are tracked and logged.

By using dynamic taint analysis, the taint module 508 and/or the monitor module 506 may monitor actions to detect whether a value that is normally derived from a trusted source is instead derived by some operation associated with the suspect data.

[0091] For example, values such as jump addresses and format strings should usually be supplied by the code itself, not from external untrusted inputs. However, an attacker may attempt to exploit a program by overwriting these values with their own data. In various embodiments, the taint module 508 may initially mark input data from untrusted sources tainted, then monitor program execution to track how the tainted attribute propagates (i.e., what other data becomes tainted) and to check when tainted data is used in dangerous ways (e.g., use of tainted data as jump addresses or format strings which may indicate an exploit of a vulnerability such as a buffer overrun or format string vulnerability). In various embodiments, based on the taint analysis, the monitor module 506 may look for variable, string, particular component and feedback that causes a jump in the code.

[0092] In various embodiments, the monitor module 506 and/or the taint module 508 may be plug-ins within the virtualization environment. In one example, the resource module 504 may provision a monitoring plug-in and a taint analysis plug-in with one or more virtualization environments.

[0093] Those skilled in the art will appreciate that the virtualization module 306 (e.g., via the monitor module 506) may detect attacks at time of use in the virtualized environment as well as at the time of writing to memory. In some embodiments, the virtualization module 306 detects when a certain part of memory is illegitimately overwritten by the suspicious data at the time of writing to the memory.

[0094] The time module 510 may speed up the perceived time within the virtualization and/or emulation environment. By increasing or slowing clock signals and processing, the suspicious data may be analyzed in a more detailed manner and/or in a faster time than if the clock signal was allowed to operate in real time.

[0095] In some embodiments, malware requires a passage of time. For example, some malware requires seconds, minutes, days, or weeks to pass before becoming active. The time module 510 may increase the clock time in the virtualization or emulation environments in order to trigger suspicious behavior.

[0096] Further, the time module 510 can slow clock time within the virtualization and/or emulation environments. For example, the time module 510 may take time slices to specifically identify and characterize processes that are taken by or for the suspicious data. In some embodiments, time slice information may be used to isolate an attack vector, describe the suspicious data, or determine the target of the attack. For example, time slice information may indicate that at a certain time and associated step, the suspicious data takes over a formerly trusted process. This information may be used to characterize malware such that when other suspicious data take similar action at the same time and associated step, the suspicious data may be classified as a similar type of malware. The time module 510 may also segment operations by or for the object in the virtualization environment and the emulation environment to simplify comparisons of operations between the virtualization environment and the emulation environment.

[0097] In various embodiments, the state module 512 tracks the various states of the virtualization environment (e.g., the time, date, process, as well as what was stored in memory where it was stored and when). In some embodiments, the virtual machine module 502 may halt a virtualization environment or instantiate a new virtualization environment utilizing the states of a previous virtualization. For example, the state module 512 may monitor the behavior of suspicious data which suspiciously terminates at time T. The virtual machine module 502 may instantiate a new virtualization environment. The state module 512 may perform dynamic state modification to change the new virtualization environment to include the logged states of the previous virtualization environment at time T. In some embodiments, the state module 512 and/or the time module 510 may increase the clock signal, decrease the clock signal, or simply change the clock signal depending on the processing of the suspicious data that needs to occur. As a result, the suspicious data may be allowed to execute in a similar environment at the desired time. Those skilled in the art will appreciate that the new virtualization environment may be slightly different (e.g., include and/or not include one or more resources) from the previous virtualization environment. In some embodiments, the virtual machine module 502 does not instantiate a new virtualization environment but rather halts the previous virtualization environment and re-instantiates the previous virtualization environment at a previously logged state with one or more resources.

[0098] The state database 514 is a database configured to store the state of one or more virtualization environments and/or one or more emulation environments. Those skilled in the

art will appreciate that the state database 514 is not limited to databases but may include any data structure.

[0099] Once the control module 310 identifies the suspicious data as malware or otherwise untrusted, the control module 310 may continue to monitor the virtualized environment to determine the vector of attack of the malware, the payload of the malware, and the target (e.g., control of the digital device, password access, credit card information access, and/or ability to install a bot, keylogger, and/or rootkit). For example, the operations performed by and/or for the suspicious data may be monitored in order to further identify the malware, determine untrusted acts, and log the effect or probable effect.

[00100] If the behavior of the suspicious data is also suspicious, the virtualization module 306 may halt the virtualization environment and provide new resources. For example, if the suspicious data begins to execute a program but abruptly halts, prepares to run an executable but does not actually run the executable, or constantly checks a section in memory that should typically be empty, then the virtualization module 306 may instantiate new virtualization environments and/or re-provision existing virtualization environments with different resources to see if the suspicious data acts differently. In various embodiments, the emulation module 308 may instantiate an emulation environment to test the suspicious data.

[00101] In various embodiments, the virtualization module 306 tracks different behaviors by different suspicious data in order to identify complex attacks, distributed attacks and/or advanced persistent threats (APT). For example, one type of malware may store an executable in a specific place in memory and then, possibly much later, a second type of malware may access the stored executable and attack a computerized system. The virtualization module 306 may identify and record the behavior of suspicious data which, when executed in a virtualization environment, only stores an executable in a specific place in memory but performs no other functions. If other data is executed in the virtualization environment which checks that specific place in memory, the virtualization module 306 may halt the virtualization, provision the executable from the previous data in the specific location in memory, and re-run the virtualization environment to monitor changes.

[00102] FIG. 6 is an exemplary virtualization environment 600 for detection of malware in some embodiments. The virtualization environment 600 comprises objects 602, a network 604, applications 606, operating system 608, a virtual machine 610, a hypervisor 612, a

manager 614, a dynamic state manager 616, and a page table manager 618. Objects include, but are not limited to, suspicious data and/or processes that are tested in the virtualization environment 600. The network 604 comprises resources to allow the objects 602 to function and/or operate with access to network resources (e.g., network drivers and ports).

[00103] The applications 606 include one or more applications or other resources that function with the objects 602 to operate in the virtualization. The applications may include word processing applications, web browsers, applets, scripting engines, and the like. Different virtualization environments may include different applications and/or different versions. For example, one virtualization environment may comprise Internet Explorer v. 9 while another virtualization environment may comprise Mozilla Firefox v. 5.0. In another example, one virtualization environment may comprise Internet Explorer v. 9 while three other virtualization environments may comprise Internet Explorer v. 8, Internet Explorer v. 7, and Internet Explorer v. 6, respectively.

[00104] The operating system 608 includes all or part of the operating system necessary for the objects 602 to function within the virtualization. The operating system may include, for example, Ubuntu Linux, Windows 7.0, or OS X Lion. Different virtualization environments may include different operating systems 608, and/or include different versions of operating systems 608 (e.g., Windows XP and Windows 7.0). Further, different virtualization environments may include different applied patches and upgrades.

[00105] The virtual machine 610 may include any number of virtual machines configured to generate one or more virtualization environments to process the objects 602. The hypervisor 612, or virtual machine manager, manages resources for the virtualizations and may allow multiple operating systems (e.g., guests) to run concurrently on the host computer. The hypervisor 612 may manage execution of the guest operating systems. In various embodiments, a kernel manages resources for the virtualizations and may allow multiple operating systems (e.g., guests) to run concurrently on the host computer. In some embodiments, the kernel performs all or some of the functions of the hypervisor 612 such as, for example, managing execution of the guest operating systems.

[00106] The manager 614 is configured to manage monitoring and control the virtualization environment 600. In various embodiments, the control module 310 controls the virtualization

environment 600, including the provisioning, time acceleration, and logging through the manager 614.

[00107] The dynamic state manager 616 (i.e., DSM) tracks and logs the state of the machine. The DSM may also store the state for later use within the same or different virtualization environments (e.g., for dynamic state modification). The state may include, for example, the object or object identifier, resources available, time slices when events occurred, and logged events. The DSM 616 may also comprise contents in memory, and locations of contents in memory over time.

[00108] The page table manager 618 may receive one or more page tables from the emulation environment. In various embodiments, the object may be tested within both the virtualization environment and the emulation environment. Upon detection of a divergence of operations between the operations of the virtualization environment and the operations of the emulation environment, the emulation module 308 may log the state of the emulation environment and pass the state information to the virtualization environment 600 as a page table for dynamic state modification of the virtualization environment. In some embodiments, the virtualization module 306 re-instantiates the original virtualization environment and dynamically modifies the state of the virtualization environment using the page table(s) from the emulation environment or the virtualization module 306 may instantiate a new virtualization environment and load the information from the page table.

[00109] FIG. 7 is a flow diagram of an exemplary malware detection method. In step 702, an object is intercepted by a data collector. The data collector may be placed on any digital device and/or network device. In step 704, the resource module 504 inspects what resources the object may require for processing (e.g., dynamic libraries and/or registries the object may affect). In some embodiments, the collector includes metadata including where the object came from, where the object was to be received, and/or what application created the request. The resource module 504 may perform preprocessing by determining what resources are required based on the metadata.

[00110] In step 706, the virtual machine monitor 502 instantiates a first instance of a virtualization environment with one or more resources identified by the resource module 504. In one example, the virtual machine monitor 502 selects and initiates plug-ins within the virtualization environment for memory allocation, forensics, mutex, filesystem, monitoring,

taint analysis, and the like. In step 708, the object is executed and/or processed within the virtualization environment.

[00111] In step 710, the taint module 508 taints operations of the object within the virtualization environment. The taint module 508 may be a plug-in. In some embodiments, the taint module 508 taints the object, bit by bit, with trace capture information. In step 712, as data propagates through the application, the monitor module 506 monitors the operations assessing what resources were previously allocated and what resources are actually allocated and called within the virtualization environment.

[00112] Resources that are required and/or called by the object which were not initially provisioned may be assessed as further evidence of malware. In some embodiments, sets of newly requested resources may be assessed to determine the likelihood of malware. For example, a particular set of resources may be determined to be malicious. If an object calls that particular set of resources (e.g., by calling resources that have not been initially provisioned, calling resources that were initially provisioned, or calling a combination of resources of which only a few were initially provisioned), the object may be determined to be malicious.

[00113] In step 714, the monitor module 506 may identify untrusted actions from monitored operations. The monitor module 506 may be a plug-in. In various embodiments, the virtual machine module 502 may load only those resources called by the resource module 504 within the virtualization environment. If the object calls a driver that is not originally provided in the virtualization environment (e.g., the object went outside of the original boundaries or the initially accepted criteria), the object's operations may terminate. In some embodiments, the virtualization environment is re-instantiated or a new virtualization environment may be instantiated that includes the additionally called resource to further process and monitor the operations of the object.

[00114] In some embodiments, the object runs in a plurality of virtualization environments until all operations called on by or for the object are completed. The control module 310 may compare the operations performed by or for the object in one virtualization to actions performed in another virtualization to analyze for divergence. If the actions taken were similar between the two virtualization environments, then no divergence was found. If the actions

taken were different, divergence is found and the differences may be further assessed (e.g., found untrusted actions taken when an unpatched operating system was present).

[00115] Divergence may be evidence of malware. For example, if the object ceases to perform any operations at time T in one virtualization environment but continues to perform many additional operations after time T in another virtualization environment (e.g., use of different resources, point to different points in memory, open a socket, or open up output ports), the difference in the environment (e.g., an available exploit) likely influenced the actions of the object and, as such, vulnerabilities may be identified.

[00116] In some embodiments, the operations taken for or by the object within the virtualization environment may be measured to determine a threat value. The threat value may be compared to a customizable threshold to determine if the behavior of the object is untrustworthy. In some embodiments, the threat value is determined based on X values and Y values. The X values may include those operations taken by a plug-in while the Y value correlates to the plug-in and the virtualization environment (e.g., operating system or hypervisor). These two values may be part of a function to determine the threat value of each operation by or for the object, an entire execution path of the object, or a part of the execution path of the object. In one example, operations taken by or for an object may be weighted based on a matrix of actions regarding an operation system, application, network environment, or object. The threat value may be compared to a threat threshold to determine if the effect of the object within the virtualization environment is sufficiently trustworthy or if the object is behaving in a suspiciously sufficient to warrant running the object through the emulation environment. Further, the threat value may be compared to the threat threshold to determine that the operations are such that they may be characterized as untrusted and, therefore, the object may be quarantined and further corrective action may be taken.

[00117] In various embodiments, the threat value associated with one or more objects may be increased (e.g., determined to be more threatening and, therefore, indicative of an increasingly likelihood of maliciousness) based on the resources called by the object. As discussed herein, for example, a particular set of resources may be determined to be malicious. If an object calls that particular set of resources, a threat value associated with object may signify a significantly increased likelihood of maliciousness.

[00118] In step 716, the reporting module 312 generates a report identifying operations and untrusted actions of the object. The reporting module 312 may generate a report identifying the object, the payload, the vulnerability, the object of the attack, recommendations for future security, and so on.

[00119] Those skilled in the art will appreciate that using signatures to identify suspicious data or malware may be optional. For example, suspicious data may be provided to the virtualization environment. If the suspicious data behaves in a manner similar to known malware, a class of malware, or a class of data with suspicious behavior, then the object may be quarantined and remedial action taken (e.g., the user of the target digital device may be notified). In some embodiments, the process of testing the suspicious data within a virtualization environment to determine a potential threat may be faster than utilizing signatures in the prior art.

[00120] FIG. 8 is a flow diagram of an exemplary method of controlling a virtualization environment to detect malware. In step 802, the state module 512 may log a first instance of the virtualization environment. For example, the state module 512 may log or track the state of the virtualization environment (e.g., time, memory values, location of data within memory, and/or ports called). The state module 512 may log the state of a plurality of virtualization environments operating in parallel.

[00121] In step 804, the virtual machine module 502 may halt the first instance of the virtualization environment. For example, the object may have terminated functions after requesting a resource not originally provided in the first instance of the virtualization environment. In some embodiments, the request for a resource not originally provisioned is evidence of malware (e.g., requesting access to a resource that the object should not have reason to access). In various embodiments, the virtual machine module 502 may permit the first instance of the virtualization environment to continue running and the virtual machine module 502 may instantiate a new instance of the virtualization environment.

[00122] In step 806, the resource module 504 determines additional resources for the object. For example, if the object requests a resource not originally provided in the first instance of the virtualization environment, the resource module 504 may identify the desired additional resource. In various embodiments, if a divergence is also detected with another virtualization

environment, the resource module 504 may also identify differences in resources between the first and other virtualization environments.

[00123] In step 808, the virtual machine module 502 re-instantiates the first instance of the virtualization environment including the previously identified resources at the previously logged state. As a result, the object may be presented with an environment that may appear to be unprotected. Further, in step 810, the time module 510 may accelerate the clock signal to the time the object requested the unavailable resource.

[00124] In step 812, the monitor module 506 may monitor operations by or for the object within the re-instantiated virtualization environment. In some embodiments, the monitor module 506 monitors the operations by or for the object as if the virtualization environment had not changed. In some embodiments, a plug-in monitors the operations by or for the object and provides information to the monitor module 506. In step 814, the monitor module 506 may identify untrusted actions from monitored operations. As discussed herein, the operations, either taken alone or in combination, may be used to determine a threat value. The threat value may be compared to a threat threshold to determine if the object is behaving suspicious, not behaving suspiciously, or behaving in an untrustworthy manner.

[00125] In step 816, the reporting module 312 may generate a report identifying suspicious or untrusted operations as well as any untrusted actions (e.g., vulnerability exploits, target of payload, defenses of the object and so on).

[00126] Those skilled in the art will appreciate that the first instance of the virtualization environment may not be halted. In some embodiments, a new instance of the virtualization environment is instantiated (without halting the previous instance) including the state information and the like. In various embodiments, the first instance of the virtualization environment is halted and then re-instantiated including the state information.

[00127] FIG. 9 is a flow diagram of an exemplary model to detect malware through multiple virtualization environments. In step 902, the collection module 302 collects the object and the resource module 504 determines one or more required resources.

[00128] In step 904, the virtual machine module 502 may instantiate the first instance of the virtualization environment (e.g., the virtual machine module 502 may instantiate a modified image of the virtualization environment) with the determined resources. Further, in step 906,

the virtual machine module 502 may instantiate a second instance of the virtualization environment but with resources that are different from that provided in the first instance of the virtualization environment. For example, versions of applications may be different, operating system patches, may be different, or the like.

[00129] In step 908, the virtual machine module 502 executes the object within the first and second instances of the virtualization environment. In step 910, the monitor module 506 may monitor operations of the object within the first and second virtualization environments. In various embodiments, the monitor module 506 traces the operations of the object in both virtualization environments. As discussed herein, a trace may be based on X values (e.g., operations by or on a plug-in of the virtualization environment) and Y values (e.g., operations between an operating system of the plug-in which may be coordinated with the X values). In some embodiments, not all operations are relevant. In some embodiments, one or more actions or operations by the host during processing may be compared against a check system to determine if the action or operation is relevant. If the action or operation is relevant, then the action or operation may be given weight and may affect the trace. In various embodiments, the one or more actions or operations by the host during processing may be compared against a check system to determine if the action or operation is not relevant. If the action or operation is not relevant, then the action or operation may be given no weight and may not affect the trace.

[00130] In step 912, the control module 310 or the monitor module 506 compares the operations of the first instance and the operations of the second instance to determine divergence. In one example, the traces of the object in the respective virtualization environments may form an execution tree which may be compared to other execution trees associated with other virtualization environments.

[00131] In one example, divergence between the traces of the two virtualization environment may be found. In various embodiments, the control module 310 may halt one or both of the virtualization environments and may notify an administrator of malware. In some embodiments, the control module 310 continues processing the object within one or both virtualization environments to further identify characteristics of the suspicious data, targeted vulnerabilities, payload, goal, or the like.

[00132] In step 914, the reporting module 312 generates a report identifying operations suspicious behavior, and/or untrusted actions of the object based, in part, on the comparison. For example, the reporting module 312 may identify the exploit that is present in some digital devices but not others. Further, the report may include recommendations to improve security (e.g., moving valuable information to a more secure location).

[00133] FIG. 10 is a block diagram of an exemplary digital device 1000. The digital device 1000 comprises a processor 1002, a memory system 1004, a storage system 1006, a communication network interface 1008, an I/O interface 1010, and a display interface 1012 communicatively coupled to a bus 1014. The processor 1002 is configured to execute executable instructions (e.g., programs). In some embodiments, the processor 1002 comprises circuitry or any processor capable of processing the executable instructions.

[00134] The memory system 1004 is any memory configured to store data. Some examples of the memory system 1004 are storage devices, such as RAM or ROM. The memory system 1004 can comprise the ram cache. In various embodiments, data is stored within the memory system 1004. The data within the memory system 1004 may be cleared or ultimately transferred to the storage system 1006.

[00135] The storage system 1006 is any storage configured to retrieve and store data. Some examples of the storage system 1006 are flash drives, hard drives, optical drives, and/or magnetic tape. In some embodiments, the digital device 1000 includes a memory system 1004 in the form of RAM and a storage system 1006 in the form of flash data. Both the memory system 1004 and the storage system 1006 comprise computer readable media which may store instructions or programs that are executable by a computer processor including the processor 1002.

[00136] The communication network interface (com. network interface) 1008 can be coupled to a network (e.g., communication network 114) via the link 1016. The communication network interface 1008 may support communication over an Ethernet connection, a serial connection, a parallel connection, or an ATA connection, for example. The communication network interface 1008 may also support wireless communication (e.g., 802.11 a/b/g/n, WiMax). It will be apparent to those skilled in the art that the communication network interface 1008 can support many wired and wireless standards.

[00137] The optional input/output (I/O) interface 1010 is any device that receives input from the user and output data. The optional display interface 1012 is any device that is configured to output graphics and data to a display. In one example, the display interface 1012 is a graphics adapter. It will be appreciated that not all digital devices 1000 comprise either the I/O interface 1010 or the display interface 1012.

[00138] It will be appreciated by those skilled in the art that the hardware elements of the digital device 1000 are not limited to those depicted in FIG. 10. A digital device 1000 may comprise more or less hardware elements than those depicted. Further, hardware elements may share functionality and still be within various embodiments described herein. In one example, encoding and/or decoding may be performed by the processor 1002 and/or a co-processor located on a GPU (i.e., Nvidia).

[00139] The above-described functions and components can be comprised of instructions that are stored on a storage medium such as a computer readable medium. The instructions can be retrieved and executed by a processor. Some examples of instructions are software, program code, and firmware. Some examples of storage medium are memory devices, tape, disks, integrated circuits, and servers. The instructions are operational when executed by the processor to direct the processor to operate in accord with embodiments of the present invention. Those skilled in the art are familiar with instructions, processor(s), and storage medium.

[00140] The present invention is described above with reference to exemplary embodiments. It will be apparent to those skilled in the art that various modifications may be made and other embodiments can be used without departing from the broader scope of the present invention. Therefore, these and other variations upon the exemplary embodiments are intended to be covered by the present invention.

Claims

1. A method comprising:
 - intercepting an object provided from a first digital device to a second digital device;
 - determining one or more resources the object requires when the object is executed;
 - instantiating a virtual environment with the one or more resources;
 - processing the object within the virtual environment;
 - tainting operations of the object within the virtual environment;
 - monitoring the operations of the object while processing within the virtual environment;
 - identifying an additional resource of the object while processing that is not provided in the virtual environment;
 - re-instantiating the virtual environment with the additional resource as well as the one or more resources;
 - monitoring the operations of the object while processing within the re-instantiated virtual environment;
 - identifying untrusted actions from the monitored operations; and
 - generating a report identifying the operations and the untrusted actions of the object.
2. The method of claim 1, wherein the object comprises an executable file, a batch file, or a data file.
3. The method of claim 1, further comprising performing a heuristic process on the object and determining the one or more resources the object requires based on the result of the heuristic process.
4. The method of claim 1, wherein determining the one or more resources the object requires based on metadata associated with the object.
5. The method of claim 1, wherein the one or more resources may include one or more applications.
6. The method of claim 1, wherein generating the report identifying the operations and the untrusted actions of the object comprises generating a signature to be used to detect malware.

7. The method of claim 1, wherein generating the report identifying the operations and the untrusted actions of the object comprises identifying a vulnerability in an application based on the operations and the untrusted actions of the object.
8. The method of claim 1, wherein re-instantiating the virtual environment with the additional resource as well as the one or more resources comprises instantiating a second instance of a virtual environment with at least one resource that is different than a resource available in the prior virtual environment.
9. The method of claim 8, further comprising comparing identified monitored operations of the prior virtual environment to operations monitored in the second instance of the virtual environment.
10. The method of claim 9, wherein generating the report comprises generating the report based, at least in part, on the comparison.
11. The method of claim 1, further comprising increasing or decreasing a clock signal within the virtual environment.
12. The method of claim 1, further comprising logging a state of the virtual environment while monitoring the operations of the object.
13. The method of claim 12, wherein re-instantiating the virtual environment with the additional resource as well as the one or more resources comprises halting the virtual environment and re-instantiating the virtual environment with the logged state.
14. A system comprising:
 - a collection module configured to receive an object provided from a first digital device to a second digital device;
 - a virtualization module configured to instantiate a virtual environment with the one or more resources, to process the object within the virtual environment, to identify an additional resource of the object while processing that is not provided in the virtual environment, re-instantiate the virtual environment with the additional resource as well as the one or more resources, and to taint operations of the object within the virtual environment;
 - a control module configured to determine one or more resources the object requires when the object is processed, to monitor the operations of the object while processing within the

virtual environment, to monitor the operations of the object while processing within the re-instantiated virtual environment, and to identify untrusted actions from the monitored operations; and

a report module configured to generate a report identifying the operations and the untrusted actions of the object.

15. The system of claim 14, wherein the object comprises an executable file, a batch file, or a data file.

16. The system of claim 14, further comprising a heuristic module configured to perform a heuristic process on the object and determining the one or more resources the object requires based on the result of the heuristic process.

17. The system of claim 14, wherein the control module determines the one or more resources of the object based on metadata associated with the object.

18. The system of claim 14, wherein the one or more resources may include one or more applications.

19. The system of claim 14, wherein the report module configured to generate the report identifying the operations and the untrusted actions of the object comprises the report module configured to generate a signature to be used to detect malware.

20. The system of claim 14, wherein the report module configured to generate the report identifying the operations and the untrusted actions of the object comprises the report module configured to identify a vulnerability in an application based on the operations and the untrusted actions of the object.

21. The system of claim 14, wherein the virtualization module configured to re-instantiate the virtual environment with the additional resource as well as the one or more resources comprises the virtualization module configured to instantiate a second instance of a virtual environment with at least one resource that is different than a resource available in the prior virtual environment.

22. The system of claim 21, wherein the virtualization module is further configured to compare identified monitored operations of the prior virtual environment to operations monitored in the second instance of the virtual environment.

23. The system of claim 22, wherein the reporting module is configured to generate the report based, at least in part, on the comparison.
24. The system of claim 14, wherein the control module is further configured to increase or decrease a clock signal within the virtual environment.
25. The system of claim 14, wherein the control module is further configured to log a state of the virtual environment while monitoring the operations of the object.
26. The system of claim 25, wherein the virtualization module configured to re-instantiate the virtual environment with the additional resource as well as the one or more resources comprises the virtualization module configured to modify the re-instantiated virtual environment based on the logged state.
27. A computer readable medium comprising instructions, the instructions being executable by a processor for performing a method, the method comprising:
- intercepting an object provided from a first digital device to a second digital device;
 - determining one or more resources the object requires when the object is executed;
 - instantiating a virtual environment with the one or more resources;
 - processing the object within the virtual environment;
 - tainting operations of the object within the virtual environment;
 - monitoring the operations of the object while processing within the virtual environment;
 - identifying an additional resource of the object while processing that is not provided in the virtual environment;
 - re-instantiating the virtual environment with the additional resource as well as the one or more resources;
 - monitoring the operations of the object while processing within the re-instantiated virtual environment;
 - identifying untrusted actions from the monitored operations; and
 - generating a report identifying the operations and the untrusted actions of the object.

1/10




FIG. 1

2/10

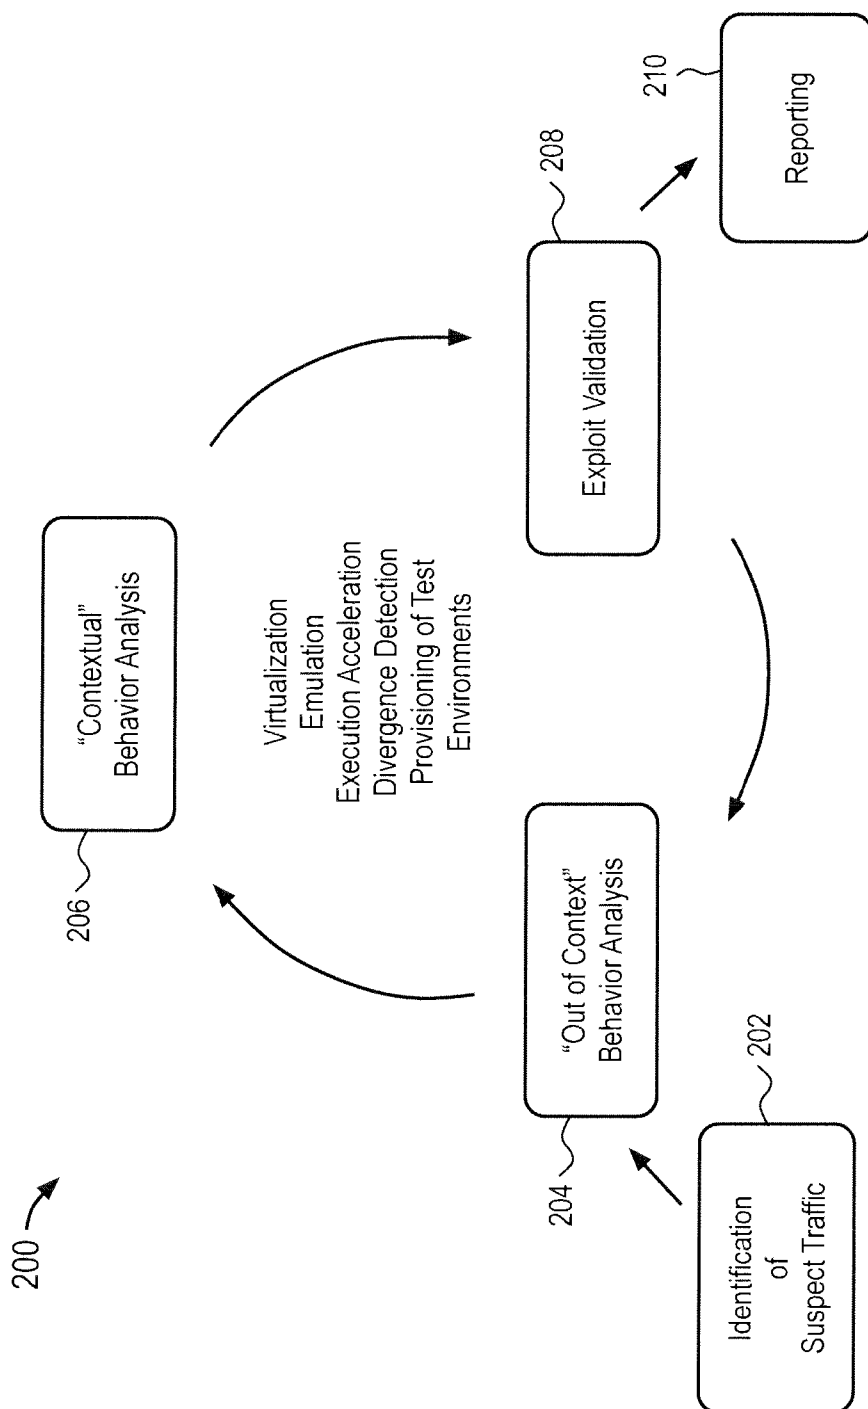


FIG. 2

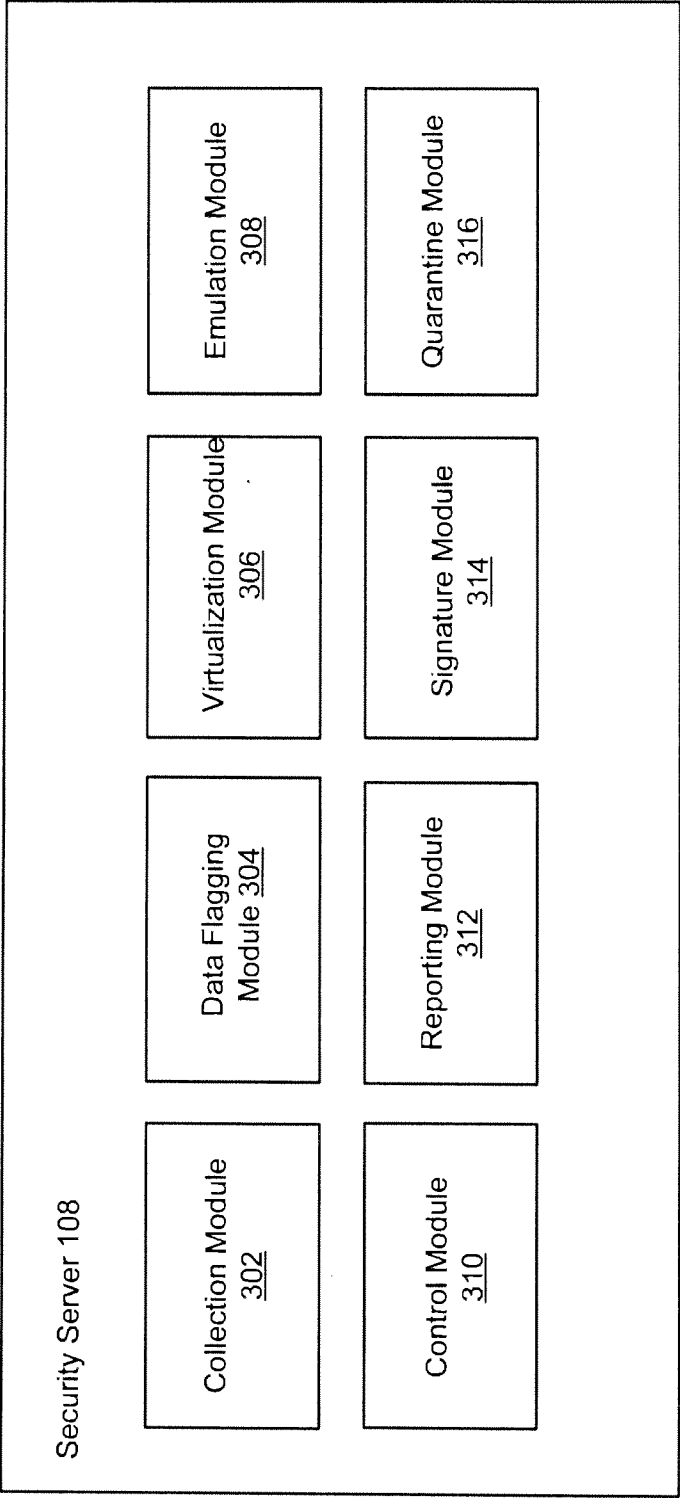


FIG. 3

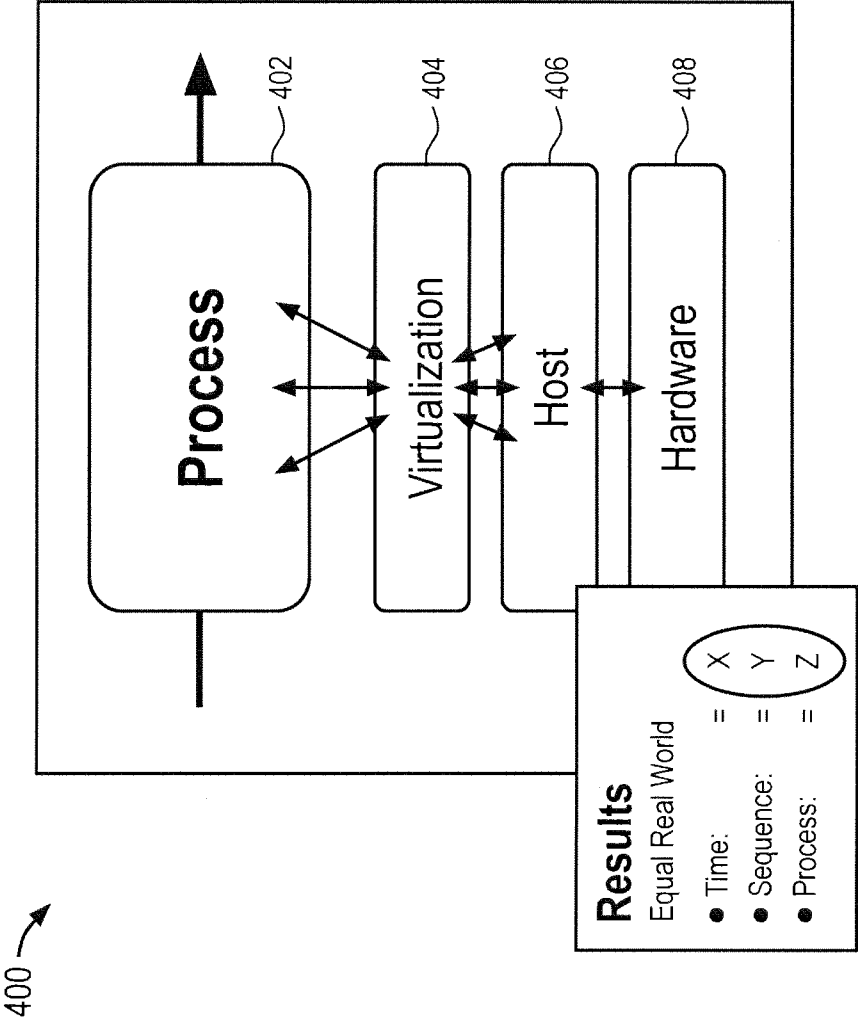


FIG. 4




FIG. 5



FIG. 6

7/10

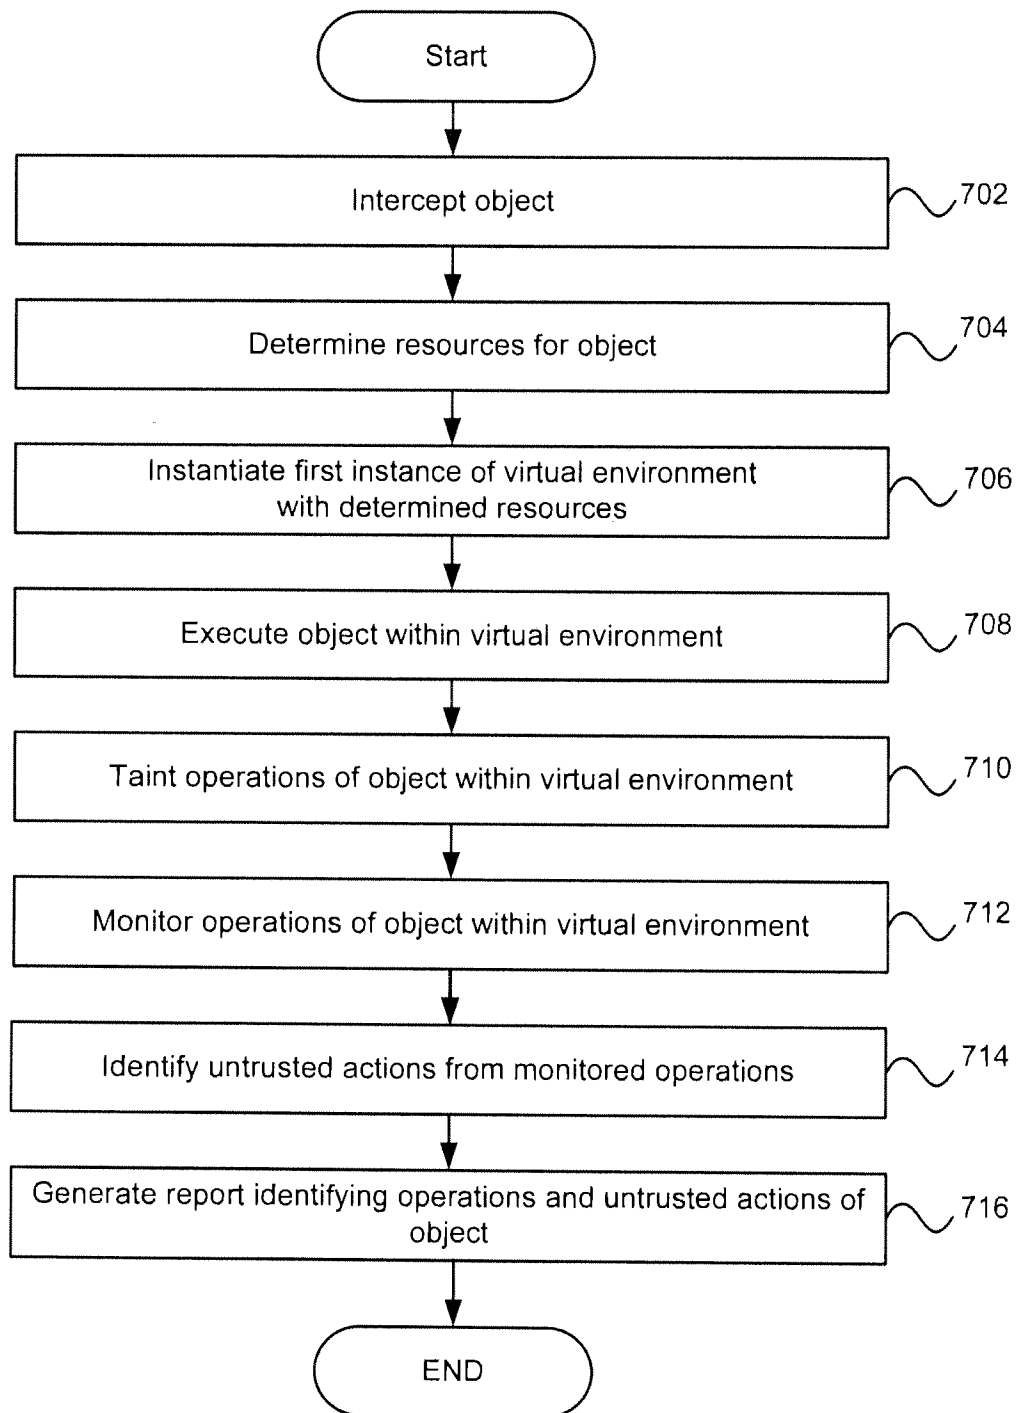


FIG. 7

8/10

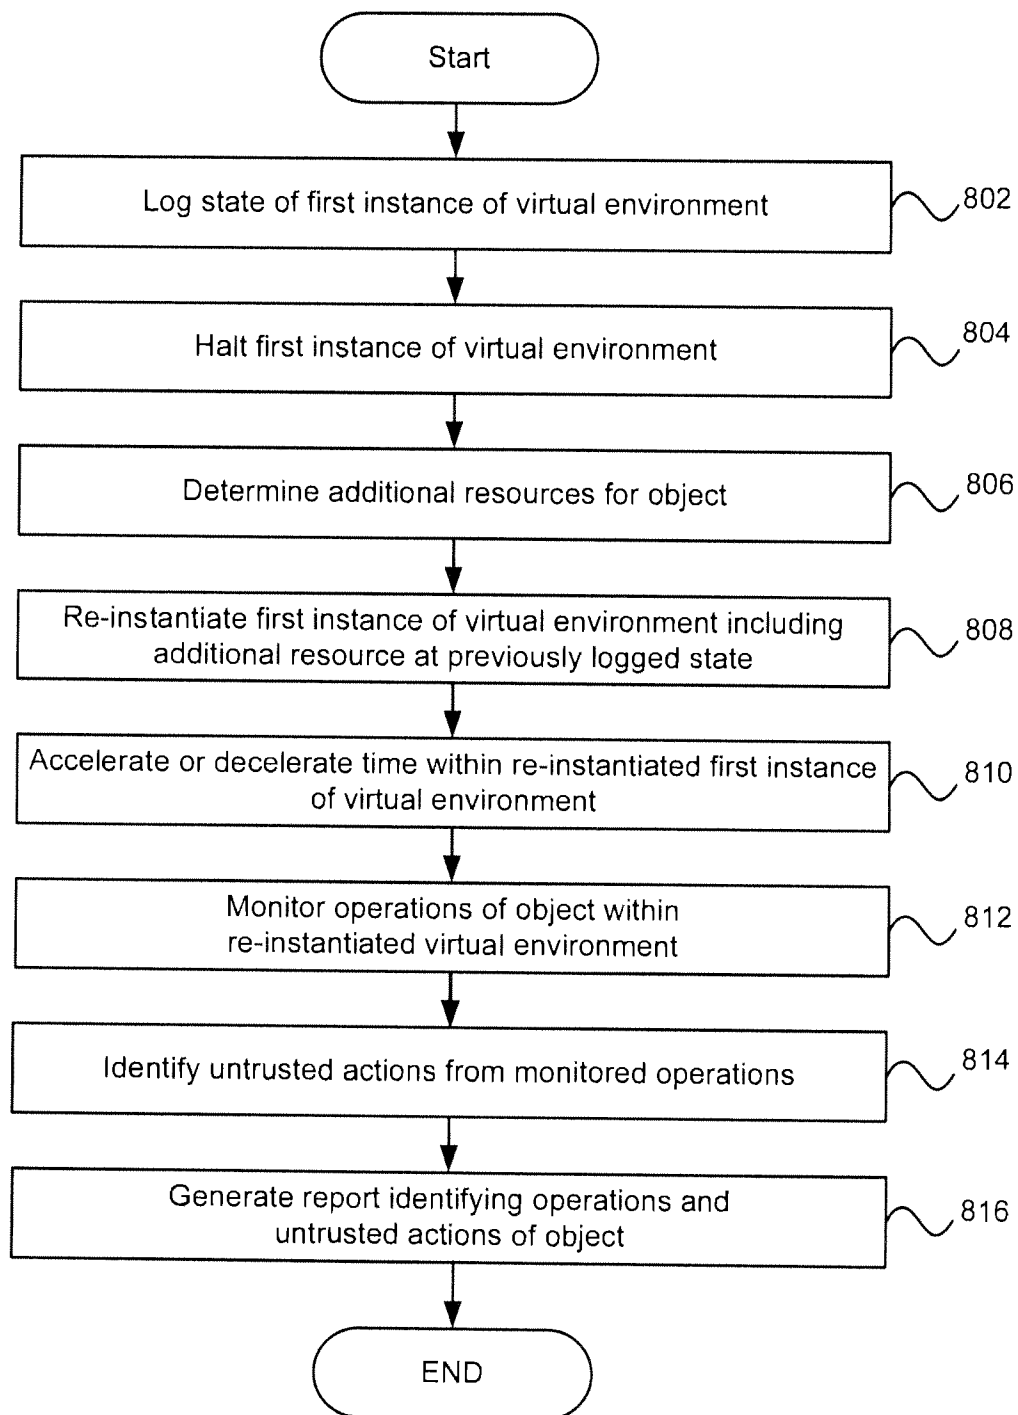


FIG. 8

9/10

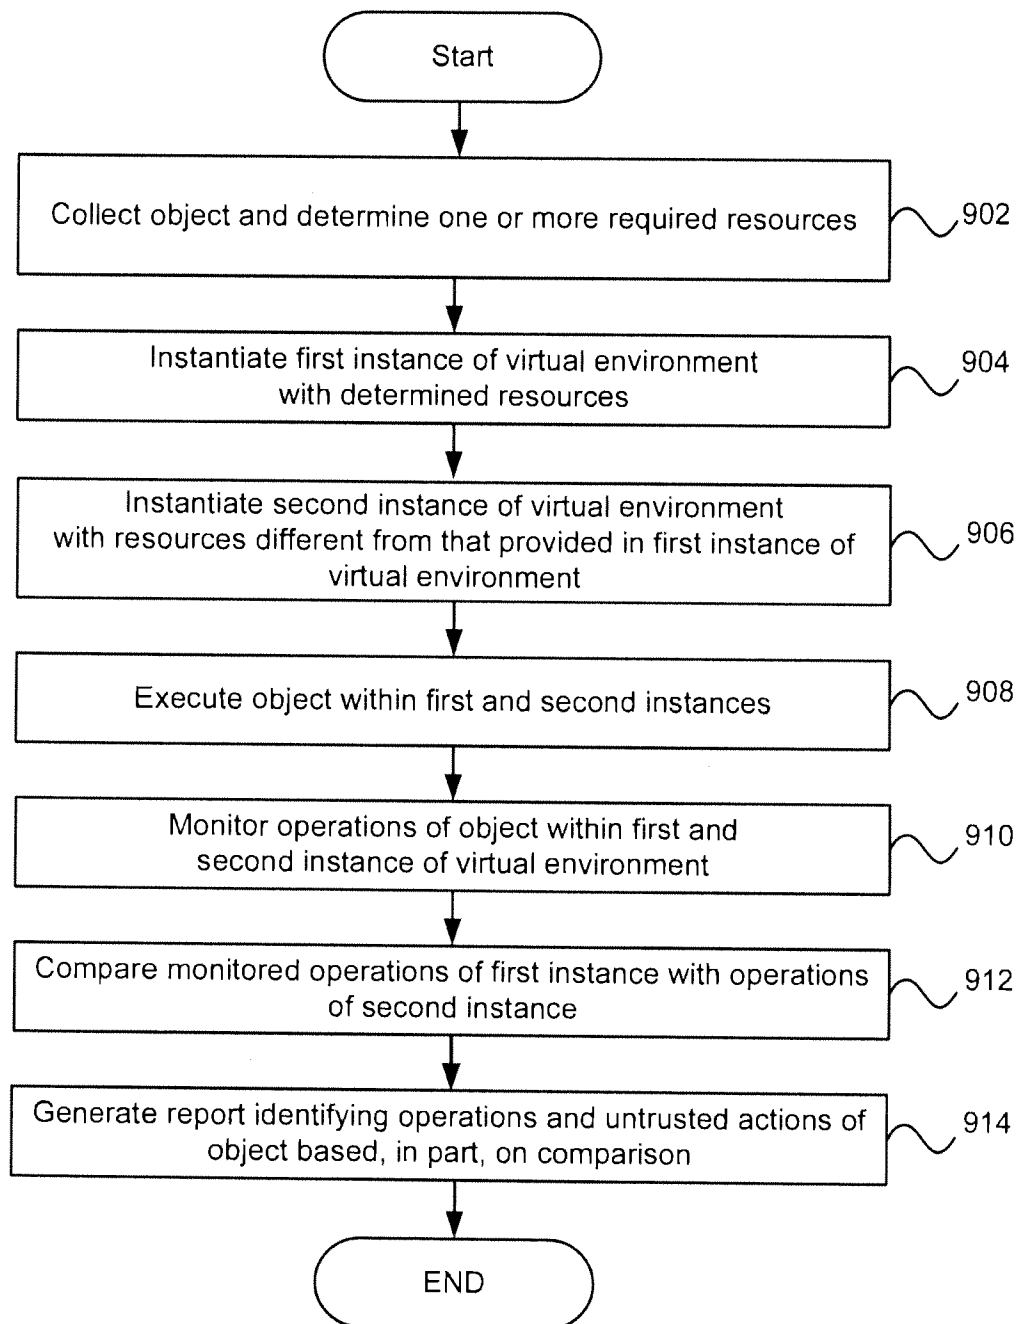


FIG. 9

10/10

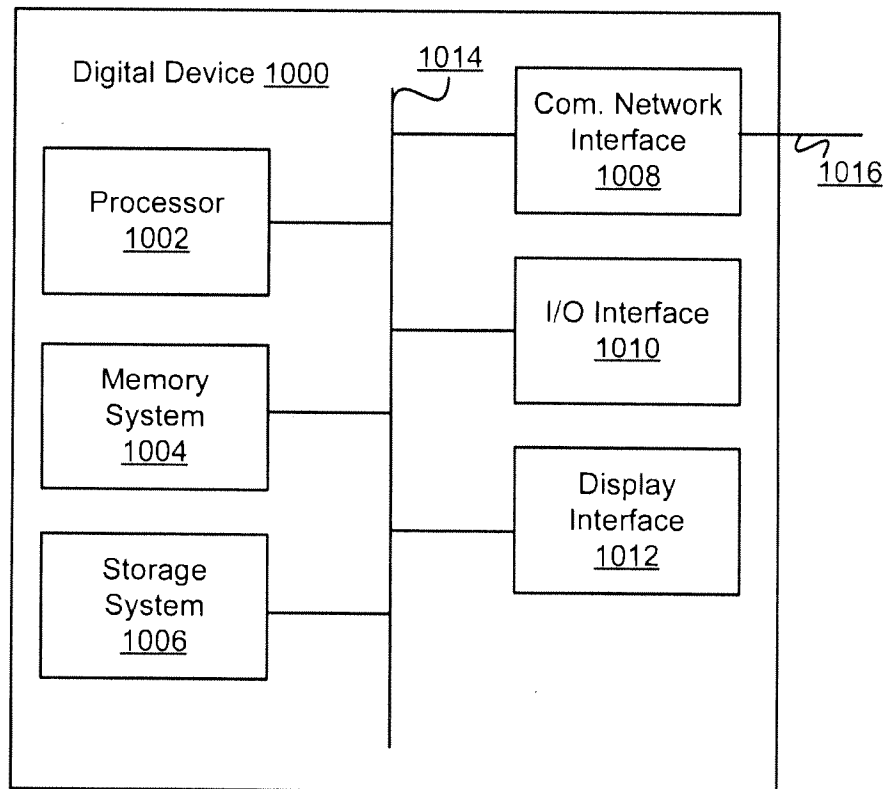


FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 12/63569

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 11/00 (2012.01)

USPC - 726/25

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 11/00 (2012.01)

USPC: 726/25

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

IPC(8): G06F 11/00 (2012.01) (search term limited; see terms)

USPC: 726/25; 726/22; 716/23; 711/6; 711/E12.001; 714/15; 714/E11.03 (search term limited; see terms)

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

Patbase, Google Search and Scholar, PubWEST(PGPB,USPT,EPAB,JPAB); Search Terms: malware, trojan, virus, worm, malicious, software, virtualization, virtual machine, virtual environment, emulate, antivirtualization, divergence, difference, trace, detecting, components, multi-stage, armoring

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2007/0250930 A1 (AZIZ et al.) 25 October 2007 (25.10.2007), entire document, especially Abstract; para [0037], [0046], [0075], [0091], [0096], [0148], [0150], [0151], [0169], [0170], [0171], [0190], [0197], [0199]	1 - 27
Y	US 2010/0064299 A1 (KACIN et al.) 11 March 2010 (11.03.2010), entire document, especially Abstract; para [0015], [0026][0027]	1 - 27
A	US 2009/0077544 A1 (WU) 19 March 2009 (19.03.2009) para [0014]-[0022]	1 - 27
A	US 2011/0167494 A1 (BOWEN et al.) 07 July 2011 (07.07.2011) para [0029]-[0032], [0049], [0084]-[0085]	1 - 27
A	US 2010/0192223 A1 (ISMAEL et al.) 29 July 2010 (29.07.2010) para [0021]-[0022], [0080]	1 - 27

☐ Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

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

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

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

04 January 2013 (04.01.2013)

Date of mailing of the international search report

18 JAN 2013

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774