



- (51) International Patent Classification:
G06F 21/22 (2006.01) G06F 11/36 (2006.01)
G06F 15/16 (2006.01)
- (21) International Application Number:
PCT/US2011/038609
- (22) International Filing Date:
31 May 2011 (31.05.2011)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant (for all designated States except US): **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **CHESS, Brian V.** [US/US]; Hewlett-Packard Company, 2215 Bridgepointe Pkwy, Suite 400, San Mateo, California 94404 (US). **RAGOLER, Iftach** [IL/US]; Hewlett-Packard Company, 5555 Windward Pkwy, Alpharetta, Georgia 30004 (US). **HAMER, Philip Edward** [US/US]; Hewlett-Packard Company, 5555 Windward Pkwy, Alpharetta, Georgia 30004 (US). **SPITLER, Russell Andrew** [US/US]; Hewlett-Packard Company, 2215 Bridgepointe Pkwy, Suite 400, San Mateo, California 94404 (US). **FAY, Sean**

Patrick [US/US]; Hewlett-Packard Company, 2215 Bridgepointe Pkwy, Suite 400, San Mateo, California 94404 (US). **JAGDALE, Prajakta Subbash** [IN/US]; Hewlett-Packard Company, 5555 Windward Pkwy, Alpharetta, Georgia 30004 (US).


(74) Agents: **MC KINNEY, Jack H.** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, Colorado 80528 (US).

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

[Continued on next page]

(54) Title: APPLICATION SECURITY TESTING



200
FIG. 2

(57) Abstract: The present disclosure provides a system that includes a server hosting an application under test (AUT), an observer configured to monitor instructions executed by the AUT, and a computing device communicatively coupled to the AUT and the observer through a common communication channel. The computing device may be configured to send an application request to the AUT, wherein the application request is configured to expose a potential vulnerability of the AUT. The computing device may receive an application response from the AUT in accordance with the AUT's programming. The computing device may send a service request to the observer, and receive a service response from the observer that contains information corresponding to the instructions executed by the AUT due to the application request, information about the AUT, or information about a server hosting the AUT.

WO 2012/166120 A1

EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, **Published:**
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, — *with international search report (Art. 21(3))*
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *as to the identity of the inventor (Rule 4.17(i))*

APPLICATION SECURITY TESTING

BACKGROUND

[0001] Software security testing is used to identify vulnerabilities in an application such as a Web application. Traditional black-box security testing for Web-based software works by using a security testing application, often referred to as a scanner, which poses as an attacker. In a black-box approach, the scanner explores an Application Under Test (AUT) by making HTTP requests and evaluating the HTTP responses in order to find all of the URLs where the AUT accepts input. The URLs where the AUT accepts input may be referred to the attack surface of the AUT. The scanner then creates attacks based on the attack surface and likely categories of vulnerabilities. The scanner applies the attacks to diagnose the presence or absence of vulnerabilities by evaluating the program's HTTP responses. In a black-box approach, the scanner does not have any insight into the internal workings of the AUT.

[0002] Black-box vulnerability testing is straightforward in concept, but in practice it presents a number of challenges. For example, exploring the AUT might not reveal all of the attack surface, so the scanner might not launch attacks against all of the places where the AUT is vulnerable. Additionally, some vulnerabilities cannot be accurately identified through the information returned in the HTTP response. If the scanner does discover a vulnerability, the scanner cannot provide information about where the vulnerability is inside the code of the AUT. Furthermore, the scanner may report several vulnerabilities that are all related to the same underlying problem in the AUT, causing a programmer trying to fix the vulnerabilities to carry out a great deal of repetitive work.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] Certain embodiments are described in the following detailed description and in reference to the drawings, in which:

[0004] Fig. 1 is a block diagram of a system that may be used to conduct gray-box security testing, in accordance with embodiments;

[0005] Fig. 2 is a block diagram showing a testing system configuration for conducting gray-box security testing, in accordance with embodiments;

[0006] Fig. 3 is a process flow diagram of a method of performing gray-box security testing, in accordance with embodiments; and

[0007] Fig. 4 is a block diagram showing a non-transitory, computer-readable medium that stores code configured to conduct gray-box security testing, in accordance with embodiments.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

[0008] Embodiments described herein provide techniques for performing gray-box security testing of a Web application. In gray-box security testing, a software program referred to herein as an observer is used to watch the internal operations performed by the AUT. The observer enables the scanner to determine the operation of the AUT and how it behaves in response to attacks. The observer also enables the scanner to determine the behavior of the AUT in response to normal application requests, which the scanner may use to determine what types of attacks to send. The scanner continues to deliver attacks to the AUT, and the scanner receives knowledge of the internal workings of the AUT from the observer. In this way, the scanner can find more vulnerabilities and produce better vulnerability reports, thereby providing a more comprehensive and detailed software security test of web-based applications.

[0009] In accordance with embodiments, a communication channel is provided between the observer and the scanner. The scanner uses this communication channel to get insight into the AUT during its scanning. The communication channel between the scanner and the observer may be implemented by using the communication channel already in use by the AUT. In this way, the person conducting the test does not have to perform additional configuration or setup work, and the communication channel does not interfere with the normal operation of the AUT or the computer system on which the AUT

is running. Additional benefits of the present invention may be better understood with reference to the description provided below.

[0010] Fig. 1 is a block diagram of a system that may be used to conduct gray-box security testing, in accordance with embodiments. The system is generally referred to by the reference number 100. Those of ordinary skill in the art will appreciate that the functional blocks and devices shown in Fig. 1 may comprise hardware elements including circuitry, software elements including computer code stored on a non-transitory, machine-readable medium, or a combination of both hardware and software elements. Further, the configuration is not limited to that shown in Fig. 1, as any number of functional blocks and devices may be used in embodiments of the present invention. Those of ordinary skill in the art would readily be able to define specific functional blocks based on design considerations for a particular electronic device.

[0011] As illustrated in Fig. 1, the system 100 may include a computing device 102, which will generally include a processor 104 connected through a bus 106 to a display 108, a keyboard 110, and one or more input devices 112, such as a mouse, touch screen, or keyboard. In an embodiment, the computing device 102 is a general-purpose computing device, for example, a desktop computer, laptop computer, server, and the like. The computing device 102 can also have one or more types of non-transitory, computer readable media, such as a memory 114 that may be used during the execution of various operating programs, including operating programs used in embodiments of the present invention. The memory 114 may include read-only memory (ROM), random access memory (RAM), and the like. The device 102 can also include other non-transitory, computer-readable media, such as a storage system 116 for the long-term storage of operating programs and data, including the operating programs and data used in embodiments of the present invention.

[0012] In an embodiment, the computing device 102 includes a network interface controller (NIC) 118, for connecting the device 102 to a server 120. The computing device 102 may be communicatively coupled to the server 120 through a network 122, such as the internet, a local area network (LAN), a wide-area network (WAN), or another network configuration. The server 120 may

have a non-transitory, computer-readable media, such as storage device, for storing data, buffering communications, and storing operating programs of the server 120. Communications between the device 102 and the server 120 may be conducted using a request-response protocol such as the Hyper-Text Transfer Protocol (HTTP).

[0013] The server 120 may be an application server that hosts the AUT 124. The server 120 also includes an observer 126 that monitors the AUT 124 during execution. The computing device 102 may include a scanner 128 that performs a security test against the AUT 124. For example, the scanner 128 may send HTTP requests to the AUT 124 through the network 122, wherein the HTTP request is configured to attempt to expose a vulnerability of the AUT 124. The HTTP requests may include HTTPS requests, which combine the Hypertext Transfer Protocol with the SSL (Secure Sockets Layer) and TLS (Transport Layer Security) protocols to provide encrypted communication and secure identification of a network Web server. During the processing of the HTTP request by the AUT 124, the observer 126 monitors the internal processes performed by the AUT 124. For example, the observer 126 can identify lines of code that are executed by the AUT 124, files that are accessed, database queries that are performed, and the like. Both the observer 126 and the AUT 124 may be configured to communicate with the scanner 128 over a same HTTP channel. As described further in reference to Fig. 2, some requests sent from the scanner 128 to the server 120 may target the AUT 124 to elicit a response from the AUT 124 in accordance with its programming. Other requests sent from the scanner 128 to the server 120 may target the observer 126 to obtain additional information regarding the effect that a particular request had on the operations performed by the AUT 124 or other information related to the AUT 124, the observer 126, or the server 120 hosting the AUT 124. Data received by the scanner 128 in responses to application requests and service requests may be used by the scanner 128 to generate vulnerability reports. Vulnerability reports may be displayed to the user through a user interface provided by the scanner 128.

[0014] Fig. 2 is a block diagram showing a testing system configuration for conducting gray-box security testing, in accordance with embodiments. The system 200 may include the scanner 128, the AUT 124, and the observer 126. The AUT 124 may be encoded in any suitable Web-based computer language, such as JAVA, or .NET, among others. The AUT 124 may operate within an suitable software framework, such as Struts, Struts 2, ASP.NET MVC, Oracle WebLogic, and Spring MVC, among others. The software framework includes a set of common code modules that provide generic functionality, which can be selectively overridden or specialized by user code to providing specific functionality. The AUT 124 may be configured to execute one or more instances of a Java Virtual Machine (JVM), Common Language Runtime (CLR), other runtime environment for processing requests from the scanner 128. The programming instructions provided by the common code modules of the software framework or runtime environment may be referred to as container code. The custom programming instructions specific to the AUT 124 may be referred to as user code.

[0015] The AUT 124 includes a network interface 202 for enabling communications between the scanner 128 and the AUT 124 through the network 122. The network interface 202 exposes the attack surface of the AUT 124 and is the same interface that would eventually be used to provide access to the AUT 124 when the AUT 124 is made available for general use. Communication between the scanner 128 and the AUT 124 over the network interface 202 may be conducted through HTTP requests issued from the scanner 128 to the AUT 124 and HTTP responses issued from the AUT 124 to the scanner 128. Requests targeting the AUT 124 may be referred to as application requests, and responses received from the AUT 124 may be referred to as application responses. The application requests generated by the scanner 128 may be configured to expose potential vulnerabilities of the AUT 124.

[0016] The AUT 124 may be coupled to a file system 204, a database 206, and other resources used by the AUT 124. The database 206 may include a variety of user information such as a table of usernames and passwords used to

grant access to various resources of the AUT 124, for example. The file system 204 may include data and programs used by the AUT 124, as well data which may be requested by a user such as HTTP pages, software programs, media files, and the like.

[0017] The observer 126 operates within the execution environment of the AUT 124 and has access to the internal operations performed by the AUT 124. For example, the observer may modify the bytecode of the AUT 124 by injecting additional code, such as a JAVA class, at various program points. The injected code acts as a monitor that observes the AUT 124. The injected monitor code may be located at strategic program points in the AUT 124, for example, application programming interface (API) calls that perform specific operations such as reading a URL parameter or writing to the file system 204. Whenever such a program point in the AUT 124 is executed, the monitor calls into services provided by the observer 126 to record operations performed by AUT 124. The observer 126 may be coupled to a buffer 210 for storing information that has been collected regarding the internal operations of the AUT 124. The buffer 210 may be used to store data that has been collected but has not yet been reported to the scanner 128. The buffer 210 may be stored in non-volatile storage medium such as a hard disk, solid state drive, and the like.

[0018] The observer 126 may also include an additional network interface 208 for enabling communications between the observer 126 and the scanner 128 through the network 122. As noted above, both network interfaces 202 and 208 may use the same communication channel, for example, the same HTTP channel. Communications between the scanner 128 and the observer 126 may be implemented through the use of custom request and response headers. Custom headers may be added to the application requests by the scanner 128, and custom headers may be added to the application responses by the observer 126. In this way, at least some of the communications between the scanner 128 and the observer 126 may be piggy-backed on normal communications with the AUT 124. Using a single channel of communication eliminates any problems with opening a dedicated, secondary channel, and

adding HTTP headers typically does not interfere with the normal operation of the AUT 124.

[0019] The scanner 128 may add one or more custom headers to each application request, wherein the custom headers include information that the observer 126 may use to diagnose a vulnerability related to an attack in progress. Information within the custom header can include the version of the scanner 128, or the payload that the scanner 128 is using in the attack. The payload information may be used by the observer 126 to determine whether the attack succeeded.

[0020] The scanner 128 may also use custom request headers to generate requests that target the observer 126 to obtain additional information regarding the internal processes performed by the AUT 124, or information about AUT, the server 120 (Fig. 1), or the observer 126. Requests targeting the observer 126 may be referred to as service requests, and responses received from the observer 126 may be referred to as service responses. Service responses issued by the observer 126 may include supplemental information in the body of the service response, as described further below.

[0021] In embodiments, the observer 126 is configured to receive the application requests and the service requests sent from the scanner 128 to the AUT 124. The observer 126 may then analyze the header information to determine whether the request is an application request or a service request. Upon receiving an application request, the observer 126 may analyze the header information to acquire data used by the observer 126 regarding the specific application request. Application requests may then be delivered by the observer 126 to AUT 124 for processing by the AUT 124 according to the AUT's programming. When the AUT 124 generates the application response, the observer 126 may add one or more custom headers to the application response to send additional information back to the scanner 128. The custom headers added to application requests and application responses may be referred to as per-request headers and are further described in the section entitled "Per-request Headers."

[0022] Upon receiving a service request, the observer 126 may process the request without delivering the service request to the AUT 124. The service request may include one or more custom headers that include information configured to request a particular service of the observer 126, such as name of the service being requested. The observer 126 may respond with the requested information in the body of an HTTP response, referred to herein as a “service response.” In embodiments, the information provided by the observer 126 in the body of the service response may be formatted using Java Script Object Notation (JSON) and may be a self-identifying JSON object. If the observer 126 has no information to send, the response body may be empty. Service requests are further described in the sections entitled “Service Requests - .”

[0023] *PER-REQUEST HEADERS*

[0024] The per-request headers may be custom HTTP headers, which may include a custom field name followed by one or more field values that are understood by the observer 126 and the scanner 128. The custom HTTP headers are ignored by the AUT 124. It will be appreciated that the field names described herein are merely used as examples of field names that could be used in a particular implementation and are not intended to limit the scope of the claims.

[0025] The per-request headers may include a version header used to coordinate the interaction between the scanner 128 and the observer 126. The observer 126 may add a version header to every application response. The scanner 128 may use the version header to verify that the observer 126 is installed. As an example, the version header may be formatted as follows:

X-WIPP-Version: <language> / <version> / <vm_id>

[0026] In the example version header, the prefix “X-WIPP” identifies the header as a custom header that is used by the observer 126. The field name “X-WIPP-VERSION” is a string of characters that uniquely identifies the custom header as a version header. The <language> field value may be the name of a runtime environment used by the AUT 124 to process the application request,

such as Java or .NET, among others. In some cases, the AUT 124 may execute two or more processes, which may be handled by different runtime instances. For example, the AUT 124 may use a load balancer or other work distribution arrangement for processing application requests. The scanner 128 may use the <vm_id> field value to identify the process within the AUT 124 handling the application request. For example, the <vm_id> field value may be a name that uniquely identifies the runtime instance that processed the application request, such as a specific JVM instance (in the case of JAVA), CLR instance (in the case of .NET), or other type of runtime instance. The <version> field value may be a number or character string that identifies the version of the observer 126. The scanner 128 may use the identified version of the observer 126 to properly coordinate interactions with the observer 126 in the event that the observer's interface 208 changes between versions of the observer software.

[0027] The per-request headers may also include a File-Not-Found (FNF) header used by the scanner 128 to identify a file-not-found condition. In HTTP, if a client requests a resource that doesn't exist or cannot be found, the Web application may generate a standard error code referred to as HTTP code 404. Simple Web applications often indicate the file-not-found condition by returning the HTTP code 404 in the HTTP response. A more complex Web application may "swallow" the 404 code. In other words, rather than simply returning the HTTP code 404 error in the HTTP response, the code 404 may trigger the Web application to redirect the client to an error page, landing page, or any other part of the Web application. In traditional black-box testing, redirecting the scanner 128 to a different part of the Web application rather than simply reporting the error may cause the scanner 128 to incorrectly report a false positive. The FNF header can be used to avoid this result. Because the observer 126 is operating inside the application, the observer 126 can detect the file-not-found error and report the file-not-found error by adding the FNF header to the application response. For example, if the application request provokes an file-not-found response from the AUT 124, the observer 126 may add the following header to the application response:

X-WIPP-FNF: 404

In this way, the file-not-found error can be reported to the scanner 128 regardless of the HTTP response provided by the AUT 124.

[0028] The per-request headers may also include a Request ID header. The Request ID header may be formatted as follows:

X-WIPP-RequestID: <request_id>

[0029] As described further below, under the section entitled "Service Requests," the scanner 128 may request additional information that was gathered by the observer 126 in response to an application request but not reported in the application response. The additional information may be included in a data structure referred to herein as a "trace." The information included in the trace describes the operations of the AUT 124 that were triggered by a specific application request. To support the trace request service, the observer 126 may add the Request ID header to each application response to enable the scanner 128 to associate the requested trace with the specific application response corresponding to the requested trace. In embodiments, the <request_id> field value is assigned by the scanner 128 and included in the application request. The observer 126 may then use the same request_id value in the Request ID header that is added to the corresponding application response. In embodiments, the scanner 128 does not add a request ID header to application requests, in which case the observer 126 may generate a unique value for the request_id and include the request_id in a Request ID header added to the application response. In either case, the same request_id value may be used by the scanner 128 to request the corresponding trace from the observer 126.

[0030] The per-request headers may also include an update header, used by the observer 126 to notify the scanner 128 of various types of changes regarding the AUT 124. The update header may be formatted as follows:

X-WIPP-Update: <service_name_list>

[0031] The <service_name_list> value may be a comma-separated list of service names, each of which refers to a services provided by the observer 126 that may provide new information as a result of a change in the AUT 124. When new information about the AUT 124 becomes available, the observer 126 may notify the scanner 128 by adding the update header to an application response. In this way, the update header informs the scanner 128 that new information about the AUT 124 is available and what service to request to obtain the information. For example, if the observer 126 detects additional URLs that have been generated during the security test of the AUT 124, the observer 126 may send an update header to the scanner 128 wherein <service_name_list> equals "AttackSurface." Upon receiving the update header, the scanner 128 may send a service request to the observer 126 requesting the identified service or services. The observer 126 may continue to send the update header in every application response until the scanner 128 issues a service request for the named service or services.

[0032] *SERVICE REQUESTS – TRACES*

[0033] In response to an application request, the observer 126 may determine the effect of the application request by determining, for example, particular lines of code that have been executed by the AUT 124, files that have been accessed by the AUT 124, database queries executed by the AUT 124, or other information. The data collected by the observer 126 can be stored to a data structure, referred to herein as a "trace." In embodiments, each trace may be stored to the buffer 210. Each trace may include the Request ID of the application request and the application response that corresponds with the trace. The scanner 128 can learn about the internal operations of the AUT 124 that were triggered by a particular application request by retrieving the corresponding trace from the observer 126. To retrieve a trace, the scanner 128 may issue a service request to the observer 126 that includes a header field name/value pair configured to indicate the request of the trace corresponding to

a particular application request or response. For example, the field name/value pair for requesting a trace may be formatted as follows:

Trace=<request_id>

[0034] The value <request_id> is the value assigned by the scanner 128 or the observer 126 that corresponds with the application request and/or application response associated with the requested trace, as described above in relation to the section entitled "PER-REQUEST HEADERS." Upon receiving the trace service request, the observer 126 may bypass the AUT 124 and generate a service response that includes the requested trace. In embodiments, the requested trace may be retrieved by the observer 126 from the buffer 210 and added to the body of the service response, which may then be sent to the scanner 128. The service response header includes the request_id value of the requested trace, and the body of the service response may be formatted as a JSON object.

[0035] The observer 126 may maintain a plurality of traces in the buffer 210 so that the scanner 128 can request a trace for any application request that has been made. The buffer 210 may be of any size suitable for a particular implementation. In an embodiment, traces stored to the buffer 210 may be removed from the buffer 210 in a first-in-first-out manner if the buffer 210 becomes full. If the scanner 128 requests an unknown request ID, the observer 126 may return an error. A request ID may be unknown if it is invalid, was never used, or has aged-out of the buffer 210 of traces maintained by the observer 126.

[0036] The scanner 128 may be configured to send a separate trace service request after the corresponding application request was made and a response received from the AUT 124. The request_id enables the observer 126 to receive trace requests out-of-sequence while still being able to associate the received trace with the appropriate application request and response. Trace requests may be received out-of-sequence due, in part, because the scanner 128 might have multiple threads of execution issuing application requests to the

AUT 124. The scanner 128 may also be configured to abort a timed-out application request, in which case the scanner 128 may retrieve an incomplete trace from the observer 126. To differentiate between complete and incomplete traces, the observer 126 may be configured to add a special node to each completed trace that indicates that the application request corresponding to this trace request was successfully completed. For example, the observer 126 may add a special node of type "request_complete" at the end of each completed trace. The absence of the "request_complete" node may indicate to the scanner 128 that the corresponding application request failed.

[0037] The observer 126 may monitor processes performed by the AUT 124 that occur outside the context of an application request, such as processes initiated by the additional monitor code injected by the observer 126. To avoid incurring an unacceptable level of performance overhead, the observer 126 may be configured to minimize the performance overhead of monitoring processes that are unrelated to application requests. For example, the performance overhead may be minimized by injecting the monitor code to selectively monitor specific API calls and relevant portions of the AUT's user code.

[0038] The trace returned to the scanner 128 may include one or more trace nodes of various types. Each trace node conveys some bit of information corresponding to the internal processes performed by the AUT 124. In embodiments, each trace node includes a type property, which may be any suitable string of characters that uniquely identifies the type of the trace node. Some trace node types may be based on the type of action that was performed by the AUT 124.

[0039] In embodiments, the observer 126 may record information about a call stack used by the AUT 124 and the container code. A call stack is a data structure that stores information about the active subroutines of a computer program. For example, the call stack may keep track of a line of code to which an active subroutine should return control when it finishes executing. The call stack may also be used to pass parameters to a subroutine, and allocate memory for variables local to the subroutine, among other functions. The call stack may include a top stack frame that represents a currently executing

subroutine called by the AUT 124 or the AUT's container code. The top stack frame may include a file name and line number that identifies a particular line of code.

[0040] Each trace may include one or more trace nodes, wherein each trace node describes details regarding a particular call stack generated by the AUT 124. Trace nodes may include a location property that identifies a file name and line number of the top stack frame outside the observer 126 in the AUT 124 container code. Trace nodes may also include a "user_context" property giving a file name and line number for the top stack frame in the AUT 124 user code if such a stack frame can be identified. The context property enables the scanner 128 to create a vulnerability report that includes a root-cause analysis for vulnerabilities and enables the scanner 128 to group together vulnerabilities that are associated with the same location in the code.

[0041] In embodiments, the observer 126 is configured to detect vulnerabilities. For example, the scanner 128 may send an attack to the AUT 124 in the form of an application request configured to create of an arbitrary file on the file system 204. The application request may include custom header information that informs the observer 126 regarding the nature of the attack. If the AUT124 is vulnerable, the observer 126 will encounter a file creation API call, thus informing the observer 126 that the File Upload attack was successful and a vulnerability has been detected.

[0042] If the observer 126 detects a vulnerability, the observer 126 may generate a trace node referred to herein as a "vulnerability trace node." A vulnerability trace node may include one or more stack traces which provide code location information to the scanner 128 such as sink program points and, when available, one or more potential source program points. A source program points is a code location where malicious input is consumed by the AUT 124, and a sink program point is a code location where the malicious input modifies the behavior of the AUT 124. For example, in the case of a Cross-Site Scripting vulnerability the source program point is where a user supplied value is read from a URL parameter, and the sink program point is where the tainted parameter value is written to an HTML page. Stack traces may be elided such

that they do not include stack frames from the observer 126. The vulnerability trace node may include stack frames relating to both the AUT 124 code and the container code, and each stack frame may include an indication of whether the stack frame relates to user code or container code. Due to the code location information provided by the observer 126, the scanner 128 can create vulnerability reports that pinpoint the location of the vulnerability in the AUT 124, and can group vulnerabilities that occur at the same location in the AUT 124, thus reducing duplication in the vulnerability reports generated by the scanner 128. The code location information also provides the user insight into the nature of the vulnerability and therefore reduces the amount of remediation effort used to repair the problem.

[0043] The vulnerability trace node may also include a vulnerability category, such as “Cross-Site Scripting” or “SQL Injection,” and standard vulnerability identifier such as a Common Weakness Enumeration (CWE) identifier corresponding to each vulnerability category. The vulnerability trace node may also include relevant details about the detection of the vulnerability. For example, if the vulnerability is an SQL injection vulnerability, the vulnerability trace node may include the SQL query involved in the detection of the vulnerability by the observer 126. In this way, the observer 126 is able to detect and report an SQL injection vulnerability to the scanner 128 even if the vulnerability does not manifest itself in the application response returned to the scanner 128.

[0044] The observer 126 may also generate a trace node referred to herein as a “database trace node” if the AUT 124 performs a query against a database, such as an SQL query. The database trace node may include the text of the database query and values for the bind parameters used by the AUT 124 in the database query. Other types of trace nodes may include begin nodes and end nodes for source code files invoked by the AUT 124 such as, JAVA servlets, JAVA Server Pages (JSPs), and the like. Begin nodes and end nodes refer to the nodes in the control flow structure representing the flow of execution through the AUT 124. Begin nodes and end nodes may include the file name of the source code file and parameters passed to the source code file. The

observer 126 may also generate other types of trace nodes to represent, for example, reads and writes of the file system 204 performed by the AUT 124, Web service calls performed by the AUT 124, and network service operations performed by the AUT 124, among others.

[0045] As stated above, the scanner 128 may use trace information to group together duplicate vulnerabilities. The grouping of duplicate vulnerabilities may be performed by a de-duplication process implemented by the scanner 128. In de-duplication, the scanner 128 may apply a hashing algorithm to portions of a vulnerability trace node, such as the user_context property and vulnerability category, in order to create an identifier for the vulnerability. Two vulnerabilities that have the same identifier can be considered duplicates from the perspective of the AUT 124. The vulnerability identifier may be used to notify the user that fixing one of these vulnerabilities will likely remedy other vulnerabilities with the same vulnerability identifier. The scanner's user interface may be configured to present duplicate vulnerabilities to the user in a group.

[0046] In embodiments, the scanner 128 is configured to optimize the attacks it sends to the AUT 124 based on the trace information. For example, if a trace shows that a particular application request does not access the file system 204, the scanner can be configured to omit similar attacks related to that application request that are directed at vulnerabilities related to the file system 204. Similarly, if a trace shows that a particular application request does not invoke a database query, the scanner can be configured to omit similar attacks related to that application request that are directed at vulnerabilities in the database, such as SQL injection.

[0047] In embodiments, database query information from database trace nodes may be used to identify more persistent cross-site scripting vulnerabilities. For example, application requests targeting the database 206 may be associated with the database tables and columns that are accessed by the AUT 124 due to the application request. The scanner 128 can use this information to send an attack that attempts to store data in the database 206 with an application request that writes the data to a particular location in the database 206. The scanner 128 may determine the effect of the attack by

sending an application request that reads from the same location in the database 206.

[0048] *SERVICE REQUESTS – SERVER INFORMATION*

[0049] The observer 126 may be configured to provide a service, referred to herein as a “server information service,” used to inform the scanner 128 about the server 120. To retrieve server information, the scanner 128 may issue a server information service request to the observer 126 that includes a header field name, such as “Server”, which is configured to indicate the request of server information. Upon receiving a server information service request, the observer 126 may bypass the AUT 124 and return the requested server information to the scanner 128.

[0050] The requested server information may be returned in the body of a service response generated by the observer 126 and formatted, for example, as a JSON object. Examples of server information included in the service response may include the name and version of the host operating system, the name and version of the application server, the amount of time that the application server has been running without any downtime, number of threads currently being processed, and the amount of memory currently in use, among other information. The scanner 128 may use the server information to generate attacks that are appropriate for the server 120 hosting the AUT 124. For example, the scanner 128 may be configured to avoid sending Microsoft® Windows-based attacks to the AUT 124 if the server 120 hosting the AUT 124 is running Linux.

[0051] *SERVICE REQUESTS – APPLICATION INFORMATION*

[0052] The observer 126 may be configured to provide a service, referred to herein as an “application information service,” that is used to inform the scanner 128 about the AUT 124. To retrieve application information, the scanner 128 may issue an application information service request to the observer 126 that includes a field name, such as “Application”, which is configured to indicate the request of application information. Upon receiving the application information service request, the observer 126 may bypass the AUT 124 and return the requested application information to the scanner 128.

[0053] The requested application information may be returned in the body of a service response generated by the observer 126 and formatted, for example, as a JSON object. If there is no application information available to be sent, the body of the service response may be empty. Examples of application information returned by the service response may include the name and version of all of the databases that the AUT 124 interacts with, file libraries used by the AUT 124, Web services subsystems, and other subsystems and software frameworks that the AUT 124 interacts with, among other information. The scanner 128 can use the application information to generate attacks more efficiently. For example, information about the databases being used by the AUT 124 may enable the scanner 128 to generate attacks that are appropriate for the identified databases and avoid generating attacks for databases that are not in use. Additionally, the scanner 128 can avoid sending attacks that target a Microsoft® SQL Server to an AUT 124 that uses only an Oracle database.

[0054] *SERVICE REQUESTS – ATTACK SURFACE*

[0055] The observer 126 may be configured to provide a service, referred to herein as an “attack surface service,” used to identify components of the attack surface that might not be detected by simple Web crawler. To retrieve attack surface information, the scanner 128 may issue an attack surface service request that includes a header field name, such as “AttackSurface”, which is configured to indicate the request of attack surface information. Upon receiving the application information service request, the observer 126 may bypass the AUT 124 and return the requested attack surface information to the scanner 128. The attack surface information may be returned in the body of a service response generated by the observer 126 and formatted, for example, as a JSON object. If there is no attack surface information available to be sent, the body of the service response may be empty.

[0056] The attack surface of the AUT 124 includes the resources that are accessible to the scanner 128, such as Web page links, for example. The scanner 128 or the observer 126 may be configured to analyze, or “crawl,” the AUT 124 to discover such Web page links. Some resources that are accessible to the scanner 128 may not be associated with a Web page link and are, in this

sense, hidden resources that are not discoverable by crawling the AUT 124. Hidden resources may exist as files in the file system 204 and may be discovered by the observer 126, which has access to the file system 204. Additionally, some resources may be generated by the AUT 124 dynamically at runtime, in other words, during execution of the AUT 124 in response to a received application request. Dynamic resources may be generated during runtime based on configuration files and mapping files, such as Web.xml files, which are files that map a requested URL to a resource on the file system 204 based on pre-defined rules. Dynamically generated resources may be discovered by the observer 126 by examining mapping and configuration files and by observing the execution of the AUT 124 to identify Uniform Resource Locators (URLs) that are bound dynamically at runtime.

[0057] Each resource accessible to the scanner 128 may be referred to as an attack surface component. Each attack surface component discovered by the observer 126 may be formatted as a URL and reported to the scanner 128 in the body of the attack surface service response. Each attack surface component may also be tagged within the body of the attack surface service response to identify the attack surface component as static or dynamic. Resources discovered by exploring the file system 204, including files located in the root directory of the file system 204 and below, may be tagged as static. Resources discovered by examining mapping and configuration files may be tagged as dynamic. In the case of containers such as WebLogic that do not expand a Web application ARchive (WAR) file as part of running the application, the observer 126 may let the container code handle the task of extracting the resources from the WAR files and then use the list of extracted resources to define the attack surface.

[0058] *Error Handling*

[0059] In some cases, the scanner 128 may issue a service request that cannot be fulfilled by the observer 126. For example, the scanner 128 may issue a service request that is not recognized by the observer 126 or a service request with a header field value that is not recognized by the observer 126, such as a trace request with an unknown request ID. If the observer 126

encounters a service requests that cannot be fulfilled, the observer 126 may return an error to the scanner 128 in the header of a service response or application response. For example, the observer 126 may issue a service response formatted as follows:

X-WIPP-Error: <error_text_string>

[0060] The header field value <error_text_string> may be any suitable text string that gives a brief description of the error encountered. The text string describing the error may be stored by the scanner 128 to an error log that is viewable by the user. Furthermore, the observer 126 may maintain an error log, wherein each error log entry includes a more detailed description of the problem. The error log maintained by the observer 126 may records any errors encountered during the operation of observer 126.

[0061] Errors may be generated by the observer 126 in response to service requests or the portion of an application request header that includes information intended to target the observer 126. If the error is generated in response to a service request, the error message may be returned in a service response, and the body of the service response may be empty. If the error is generated in response to information included in an application request, the error message may be returned in an application response, and the body of the application response may contain whatever data the AUT 124 provides in accordance with its programming.

[0062] Fig. 3 is a process flow diagram summarizing a method of performing gray-box security testing, in accordance with embodiments. The method 300 may be performed by the scanner 128 in communication with the AUT 124 and the 126, as described in reference to Fig. 2. The method 300 may begin at block 302, wherein an application request may be sent to AUT 124. The application request may be configured to expose a potential vulnerability of the AUT 124. As discussed above, the application request may include custom headers used to communicate information to the observer 126. For example,

the scanner 128 may add a request ID value to a header of the application request that uniquely identifies the application request to the observer 126.

[0063] At block 304, the scanner 128 may receive an application response from the AUT 124 in accordance with the AUT's programming. If the scanner 128 is configured to add request IDs to each application request, the observer 126 may add the same request ID to the header of the application response. Otherwise, the scanner 128 may generate a unique request ID and add the request ID to the header of the application response. As described above, the observer 126 may also add additional information to the header of the application response, such as a file-not-found header, an observer version header, an update header, and the like.

[0064] At block 306, the scanner 128 may send a service request to the observer 126. In embodiments, the service request may be included in the application request of block 302 by adding additional header information to the application request. In embodiments, the service request may be a separate request, in other words, not combined with an application request. Service requests are processed by the observer 126 and not passed to the AUT 124. Service requests may be configured to request information such as attack surface information, server or application information, and trace information that relates to the internal processes of the AUT 124, among others. The scanner 128 may add the request ID to the header of the service request to obtain trace information corresponding to a particular application response.

[0065] At block 308, the scanner 128 may receive a service response from the observer 126. The service response may include information about the processes performed by the AUT 124 due to the application request. For example, the service response may include stack trace information identifying processes performed by the AUT 124 as a result of the application request. The service response may include vulnerability trace nodes that contain code locations corresponding to a vulnerability detected by the observer 126. The service response may also include information corresponding to a database query performed by the AUT 124 as a result of the application request. In embodiments, the service response may include information about the AUT 124

such as the programming language of the AUT 124, the name and version of the AUT 124, and the attack surface of the AUT 124 including static and dynamic URLs, among other information. The service response may also include information about the observer 126, such as the operating system name and version, application server name and version, number of threads, memory usage, and amount that the server 120 has been running without any downtime, among other information.

[0066] It will be appreciated that method 300 is merely an example process flow used to explain embodiments of the present techniques, and that the actual process flow may vary depending on a particular implementation. For example, the scanner 128 may not issue a service request for every application request. Additionally, the scanner 128 may send and receive multiple application requests and responses before sending a service request related to a particular one of the application responses.

[0067] At block 310, the scanner 128 may generate a vulnerability report based on the information received from the AUT 124 and the observer 126. The vulnerability report may group detected vulnerabilities based on the code location information included in the trace nodes received at block 308. The vulnerability report may be presented to the user through a user interface provided by the scanner 128. The vulnerability report may also be stored to memory, printed, and the like.

[0068] Fig. 4 is a block diagram showing a non-transitory, computer-readable medium that stores code configured to conduct gray-box security testing, in accordance with embodiments. The non-transitory, machine-readable medium is referred to by the reference number 400. The non-transitory, machine-readable medium 400 can include RAM, a hard disk drive, an array of hard disk drives, an optical drive, an array of optical drives, a non-volatile memory, a universal serial bus (USB) drive, a digital versatile disk (DVD), a compact disk (CD), and the like. The non-transitory, machine-readable medium 900 may be accessed by a processor 902 over a communication path 904.

[0069] As shown in Fig. 4, the various components discussed herein can be stored on the non-transitory, machine-readable medium 400. A region 406 on

the non-transitory, machine-readable medium 400 can include an application interface configured to send application requests to the AUT 124, wherein the application requests are configured to expose a potential vulnerability of the AUT 124. The application interface can also receive application responses from the AUT 124, wherein the application responses are generated by the AUT 124 in accordance with the AUT's programming. A region 408 can include an observer interface configured to send service requests to the observer 126. The observer interface also receives the service responses, which contain information corresponding to processes performed by the AUT 124 due to the application request, information about the AUT 124, or information about the server 120 hosting the AUT 124, for example. A region 410 can include a vulnerability report generator configured to analyze the data received from the AUT 124 and the observer 126 and generate a vulnerability report based in the analysis.

CLAIMS

What is claimed is:

1. A system, comprising:
 - a server hosting an application under test (AUT);
 - an observer configured to monitor instructions executed by the AUT; and
 - a computing device communicatively coupled to the AUT and the observer through a common communication channel, the computing device comprising a processor and a memory device for storing computer-readable instructions configured to direct the processor to:
 - send an application request to the AUT, wherein the application request is configured to expose a potential vulnerability of the AUT;
 - receive an application response from the AUT in accordance with the AUT's programming;
 - send a service request to the observer; and
 - receive a service response from the observer that contains information corresponding to the instructions executed by the AUT due to the application request, information about the AUT, or information about a server hosting the AUT.
2. The system of claim 1, wherein the observer is configured to generate a trace identifying the instructions executed by the AUT as a result of the application request, and send the trace to the computing device in the body of the service response.
3. The system of claim 1, wherein the observer is configured to communicate with the computing device, at least in part, by adding a custom header to the application response.
4. The system of claim 1, the memory device comprising computer-readable instructions configured to direct the processor to receive trace information from the observer, the trace information comprising a plurality of

vulnerability trace nodes, each vulnerability trace node containing code locations corresponding to a vulnerability detected by the observer.

5. The system of claim 4, the memory device comprising computer-readable instructions configured to direct the processor to group the plurality of vulnerability trace nodes based on the vulnerability trace nodes containing the same code locations.

6. The system of claim 1, wherein the observer is configured to monitor the AUT to identify new uniform resource locators (URLs) that are generated dynamically during runtime of the AUT, and return an update field in a header of the application response, the update field configured to inform the computing device that an attack surface of the AUT has changed.

7. The system of claim 1, wherein the observer is configured to:

- receive a request from the computing device and analyze a header of the request;
- identify the request as the application request or the service request based on the analysis of the header;
- pass the application request to the AUT; and
- process the service request without passing the service request to the AUT.

8. A method, comprising:

- sending an application request to an application under test (AUT), wherein the application request is configured to expose a potential vulnerability of the AUT;
- receiving an application response from the AUT in accordance with the AUT's programming;
- sending a service request to an observer that monitors instructions executed by the AUT; and

receiving a service response from the observer that contains information corresponding to instructions executed by the AUT due to the application request, information about the AUT, or information about a server hosting the AUT;

wherein the application request, application response, service request, and service response are communicated over a same network channel.

9. The method of claim 8, comprising receiving a file-not-found header in the application response, the file-not-found header added to the application response by the observer to indicate a file-not-found error generated by the AUT.

10. The method of claim 8, wherein the service request is a trace service request, the method comprising receiving a stack trace in a body of the service response received from the observer.

11. The method of claim 8, comprising receiving a vulnerability trace node in the body of the service response, wherein the vulnerability trace node identifies a vulnerability detected by the observer.

12. The method of claim 8, wherein the service request is an attack surface service request, the method comprising receiving information about the attack surface of the AUT in a body of the service response, the attack surface comprising static URLs and dynamic URLs that are generated by the AUT during runtime.

13. A non-transitory, computer readable medium, comprising code configured to direct a processor to:

send an application request to an application under test (AUT), wherein the application request is configured to expose a potential vulnerability of the AUT;

receive an application response from the AUT in accordance with the AUT's programming;

send a service request to an observer that monitors instructions executed by the AUT; and


receive a service response from the observer that contains information corresponding to instructions executed by the AUT due to the application request, information about the AUT, or information about a server hosting the AUT;

wherein the application request, application response, service request, and service response are communicated over a same network channel.

14. The non-transitory, computer readable medium of claim 13, comprising code configured to direct the processor to add a request ID to a header of the application request that uniquely identifies the application request, wherein the service response received from the observer includes the request ID.


15. The non-transitory, computer readable medium of claim 13, wherein the service response includes a database trace node that includes information corresponding to a database query performed by the AUT as a result of the application request.

1/4



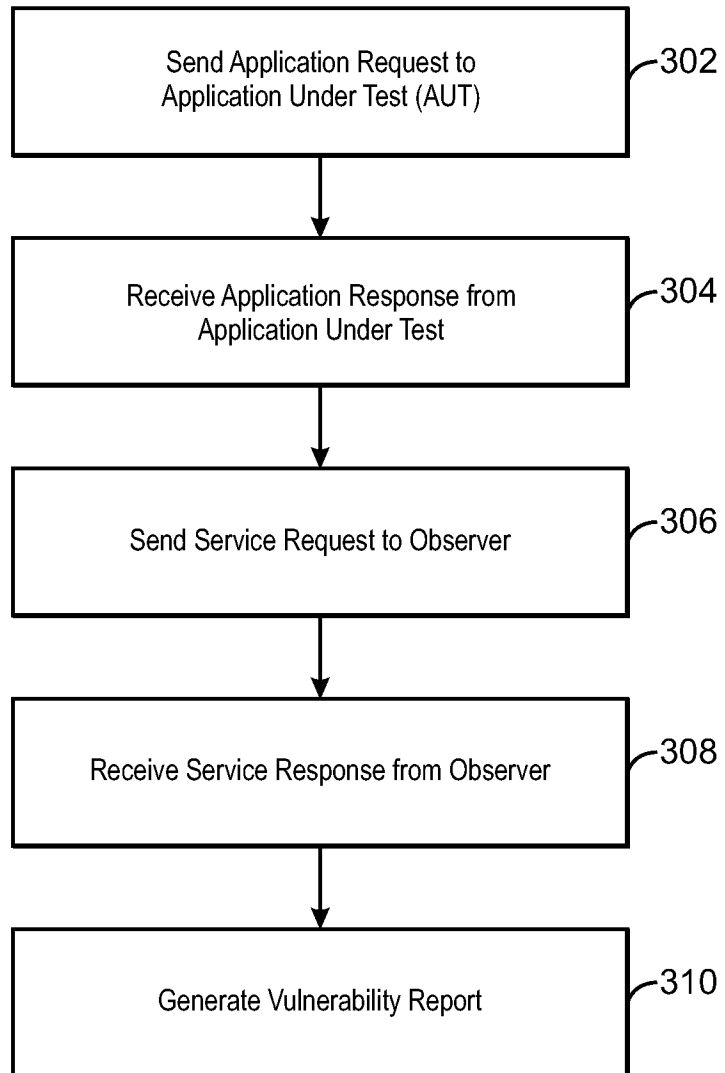
100
FIG. 1

2/4



200
FIG. 2

3/4



300
FIG. 3

4/4

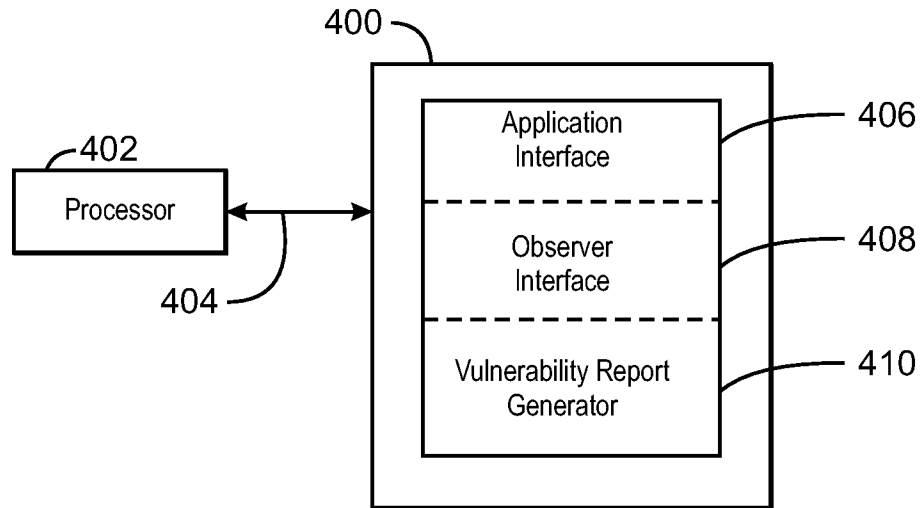


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2011/038609**A. CLASSIFICATION OF SUBJECT MATTER***G06F 21/22(2006.01)i, G06F 15/16(2006.01)i, G06F 11/36(2006.01)i*

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 21/22; G06F 11/30; G06F 12/19; G06F 12/14; G06F 7/04; G08B 23/00

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

Korean utility models and applications for utility models
Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: application, security, vulnerability, scanner, attacker, black, box, testing, request, response, observer, server, trace, node, header

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 6584569 B2 (RESHEF; ERAN et al.) 24 June 2003 See abstract; column 4, line 34 - column 11, line 46; figures 2, 3.	1-15
A	US 6311278 B1 (RAANAN; GIL et al.) 30 October 2001 See abstract; column 4, line 46 - column 6, line 13; figures 3, 4.	1-15
A	US 7000247 B2 (CARL E. BANZHOF) 14 February 2006 See abstract; column 7, line 32 - column 10, line 20; figures 3-5.	1-15
A	US 2007-0186285 A1 (DENNIS HURST et al.) 09 August 2007 See abstract; paragraphs [23] - [57]; figure 2.	1-15

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

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

31 JANUARY 2012 (31.01.2012)

Date of mailing of the international search report

09 FEBRUARY 2012 (09.02.2012)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
Government Complex-Daejeon, 189 Cheongsu-ro,
Seo-gu, Daejeon 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

Shin Sang Gil

Telephone No. 82-42-481-8480



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2011/038609

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6584569 B2	24.06.2003	US 2002-0010855 A1	24.01.2002
		US 2003-0233581 A1	18.12.2003
		US 7237265 B2	26.06.2007
US 6311278 B1	30.10.2001	AT 463793 T	15.04.2010
		AU 2000-55612 A1	22.01.2001
		AU 5561200 A	22.01.2001
		DE 60044133 D1	20.05.2010
		EP 1203297 A1	08.05.2002
		EP 1203297 A4	12.05.2004
		EP 1203297 B1	07.04.2010
		IL 147423 D0	14.08.2002
		JP 04-733886 B2	28.04.2011
		JP 2003-504723 A	04.02.2003
		JP 2003-504723 T	04.02.2003
		JP 2009-134725 A	18.06.2009
		JP 4733886 B2	27.07.2011
		US 2002-0116643 A1	22.08.2002
		US 2005-0044420 A1	24.02.2005
		US 6321337 B1	20.11.2001
US 7774835 B2	10.08.2010		
WO 01-02963 A1	11.01.2001		
US 7000247 B2	14.02.2006	US 2003-0126472 A1	03.07.2003
		US 2005-0091542 A1	28.04.2005
		US 2005-0229256 A2	13.10.2005
		US 7308712 B2	11.12.2007
US 2007-0186285 A1	09.08.2007	US 2006-0137014 A1	22.06.2006
		US 6996845 B1	07.02.2006
		US 7200867 B2	03.04.2007
		US 7444680 B2	28.10.2008