

(51) International Patent Classification:
G06F 21/55 (2013.01)

(21) International Application Number:
PCT/US2016/043040

(22) International Filing Date:
20 July 2016 (20.07.2016)

(25) Filing Language:
English

(26) Publication Language:
English

(30) Priority Data:
62/196,404 24 July 2015 (24.07.2015) US
15/213,896 19 July 2016 (19.07.2016) US

(71) Applicant: NEC LABORATORIES AMERICA, INC.
[US/US]; 4 Independence Way, Suite 200, Princeton, New Jersey 08540 (US).

(72) Inventors: CHEN, Zhengzhang; 44 York Road, Princeton Junction, New Jersey 08550 (US). TANG, LuAn; 10 Manley Road, Pennington, New Jersey 08534 (US). DONG, Boxiang; Lieb 223, Hoboken, New Jersey 07030 (US). JI-ANG, Guofei; 5 Danby Court, Princeton, New Jersey 08540 (US). CHEN, Haifeng; 68 Gate Lane, Old Bridge, New Jersey 08857 (US).

(74) Agent: KOLODKA, Joseph; NEC Laboratories America, Inc., 4 Independence Way, Suite 200, Princeton, New Jersey 08540 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,

[Continued on next page]

(54) Title: GRAPH-BASED INTRUSION DETECTION USING PROCESS TRACES




FIG. 1

(57) Abstract: Methods and systems for detecting malicious processes include modeling system data as a graph comprising vertices that represent system entities and edges that represent events between respective system entities. Each edge has one or more timestamps corresponding respective events between two system entities. A set of valid path patterns that relate to potential attacks is generated. One or more event sequences in the system are determined to be suspicious based on the graph and the valid path patterns using a random walk on the graph.



TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

GRAPH-BASED INTRUSION DETECTION USING PROCESS TRACES

RELATED APPLICATION INFORMATION

[0001] This application is a continuation in part of Application No. 15/098,861, filed on April 14, 2016, which in turn claims priority to Application Serial No. 62/148,232, filed on April 16, 2015. This application further claims priority to Application Serial No. 62/196,404, filed on July 24, 2015, and to Application Serial No. 62/360,572, filed on July 11, 2016, incorporated herein by reference in their entirety.

BACKGROUND

Technical Field

[0002] The present invention relates to computer and information security and, more particularly, to host-level intrusion detection on massive process traces.

Description of the Related Art

[0003] Enterprise networks are key systems in corporations and they carry the vast majority of mission-critical information. As a result of their importance, these networks are often the targets of attack. To guarantee the information security in the network of computers, an intrusion detection system is needed to keep track of the running status of the entire network and identify scenarios that are associated with potential attacks or malicious behaviors.

[0004] At the host level, detection systems collect rich information about process/program events (e.g., when a program opens a file) on a particular host or machine. While this information enables intrusion detection systems to monitor intrusive behavior

accurately, signature-based detection techniques fail to detect new threats, while anomaly-based detection techniques either focus on detecting a single abnormal process or need an offline model to be built from training data with purely normal events.

[0005] More importantly, intrusion detection systems often rely on a coordinated or sequential, not independent, action of several system events to determine what state a given system is in. The system monitoring data is typically made up of low-level process events or interactions between various system entities, such as processes, files and sockets (e.g., when a program opens a file or connects to a server) with exact time stamps, while attempted intrusions are higher-level activities which usually involve multiple different process events. For example, a network attack called Advanced Persistent Threat (APT) is composed of a set of stealthy and continuous computer hacking processes. APT first attempts to gain a foothold in the environment. Then, using the compromised systems as an access into the target network, APT deploys additional tools that help fulfill the attack objective. The gap existing between the levels of process events and intrusion activities makes it hard to infer which process events are related to real malicious activities, especially considering that there are massive, “noisy” process events happening in between. Hence, conventional attack detection techniques that identify individual suspicious process events are inadequate to address this scenario.

SUMMARY

[0006] A method for detecting malicious processes includes modeling system data as a graph comprising vertices that represent system entities and edges that represent events between respective system entities. Each edge has one or more timestamps corresponding

respective events between two system entities. A set of valid path patterns that relate to potential attacks is generated. One or more event sequences in the system are determined to be suspicious based on the graph and the valid path patterns using a random walk on the graph.

[0007] A system for detecting malicious processes includes a modeling module configured to model system data as a graph that has vertices that represent system entities and edges that represent events between respective system entities. Each edge includes one or more timestamps corresponding respective events between two system entities. A malicious process path discovery module includes a processor configured to generate a set of valid path patterns that relate to potential attacks and to determine one or more event sequences in the system to be suspicious based on the graph and the valid path patterns using a random walk on the graph.

[0008] These and other features and advantages will become apparent from the following detailed description of illustrative embodiments thereof, which is to be read in connection with the accompanying drawings.

BRIEF DESCRIPTION OF DRAWINGS

[0009] The disclosure will provide details in the following description of preferred embodiments with reference to the following figures wherein:

[0010] FIG. 1 is directed to a network graph representing communities and roles of nodes in accordance with the present principles.

[0011] FIG. 2 is a block/flow diagram of a method of discovering community and role memberships and detecting anomalies in accordance with the present principles.

[0012] FIG. 3 is a block diagram of a host-level analysis module in accordance with the present principles.

[0013] FIG. 4 is a block/flow diagram of a method for detecting suspicious host-level event sequences in accordance with the present principles.

[0014] FIG. 5 is a segment of pseudo-code for detecting suspicious host-level event sequences in accordance with the present principles.

[0015] FIG. 6 is a block diagram of a processing system in accordance with the present principles.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0016] In accordance with the present principles, the present embodiments provide malicious process path discovery, detecting abnormal process paths that are related to intrusion activities. This is accomplished using process traces. A set of valid sequence patterns is generated and a random walk based process is used to learn system functions and discover suspicious process sequences. To eliminate score bias from the path length, the Box-Cox power transformation is applied to normalize the anomaly scores of process sequences.

[0017] The present embodiments thereby provide complete evidence of an attacker's activity trace (i.e., the process path) after an attack has occurred. In addition, the present embodiments more accurately detect malicious process paths, reducing the number of false positives and false negatives in less time and with less computational complexity. A compact graph structure may be used to reduce memory load, a set of valid sequence patterns may be generated and used to reduce the search space, and a random walk

approach may be used to reduce the computational cost. Furthermore, the present embodiments are able to detect new attacks because no training data is needed.

[0018] Referring now in detail to the figures in which like numerals represent the same or similar elements and initially to FIG. 1, an automatic security intelligence system (ASI) architecture is shown. The ASI system includes three major components: an agent 10 is installed in each machine of an enterprise network to collect operational data; backend servers 200 receive data from the agents 10, pre-process the data, and sends the pre-processed data to an analysis server 30; and an analysis server 30 that runs the security application program to analyze the data.

[0019] Each agent 10 includes an agent manager 11, an agent updater 12, and agent data 13, which in turn may include information regarding active processes, file access, net sockets, number of instructions per cycle, and host information. The backend server 20 includes an agent updater server 21 and surveillance data storage. Analysis server 30 includes intrusion detection 31, security policy compliance assessment 32, incident backtrack and system recovery 33, and centralized threat search and query 34.

[0020] Referring now to FIG. 2, additional detail on intrusion detection 31 is shown. There are five modules in an intrusion detection engine: a data distributor 41 that receives the data from backend server 20 and distributes the corresponding to network level module 42 and host level module 43; network analysis module 42 that processes the network communications (including TCP and UDP) and detects abnormal communication events; host level analysis module 43 that processes host level events, including user-to-process events, process-to-file events, and user-to-registry events; anomaly fusion module 44 that integrates network level anomalies and host level anomalies and refines the results for

trustworthy intrusion events; and visualization module 45 that outputs the detection results to end users.

[0021] Referring now to FIG. 3, additional detail on the host level analysis module 43 is shown. The host level analysis module 43 includes a hardware processor 312 and a memory 314. In addition, the host level analysis module 43 includes one or more functional modules that may, in one embodiment, be stored in memory 314 and executed by hardware processor 312. In alternative embodiment, the functional modules may be implemented as one or more discrete hardware components in the form of, e.g., application specific integrated chips or field programmable gate arrays.

[0022] The host level analysis module 43 includes a number of different analysis and detection functions. A process-to-file anomaly detection module 302 takes host level process-to-file events from data distributor 41 as input and discovers abnormal process-to-file events. These events may include, e.g., reading from or writing to a file. A user-to-process anomaly detection module 304 takes all streaming process events as input from the data distributor 41 and models each user's behavior at the process level, identifying suspicious processes run by each user. A USB event anomaly detection module 306 also considers the streaming process events and identifies all USB device-related events to detect anomalous device activity. A process signature anomaly detection module 308 takes process names and signatures as input from data distributor 41 and detects processes with suspicious signatures. Finally, a malicious process path discovery module 310 takes current active processes from the data distributor 41 as starting points and tracks all of the possible process paths by combing the incoming and previous events in a time window. The

malicious process path discovery module 310 detects anomalous process sequences/paths as described in greater detail below.

[0023] Referring now to FIG. 4, a method for malicious process path detection is shown. A blueprint graph is used as input. The blueprint graph is a heterogeneous graph constructed from a historical dataset of communications in a network, with nodes of the blueprint graph representing physical devices on an enterprise network and edges reflecting the normal communication patterns among the nodes. Block 402 performs graph modeling, using a compact graph structure to capture the complex interactions between system entities as an acyclic, multipartite graph. Block 404 then generates a set of valid sequence patterns based on a maximum sequence length. The maximum sequence length may be set by a user or an optimal value may be determined automatically. Forming the valid sequence patterns dramatically reduces the search space size of the graph.

[0024] Block 406 then scans the graph to determine candidate event sequences that is consistent with the patterns. A “pattern” refers to an ordered set of system entity types, while a “sequence” refers to an ordered set of specific system entities. Thus, a sequence is consistent with a pattern if its series of system entities’ respective types match those of the pattern. Sequences are alternatively referred to herein as “paths.”

[0025] Block 408 applies a random walk to extract the characteristics of every entity. Based on a discovered entity’s nature, block 408 calculates an anomaly score of each candidate process to evaluate how abnormal the process is. As there may be multiple different sequence patterns of different lengths, and as the scores of two paths with different lengths are not directly comparable, the anomaly score distribution for each sequence pattern is transformed into a single distribution by block 410 using, e.g., a Box-Cox power

transformation. Block 410 measures the deviation between the suspicious sequences and the normal sequences, reporting those sequences that have a higher-than-threshold deviation.

[0026] Attacking behavior often involves multiple system entities (e.g., processes, files, sockets, etc.). Therefore the present embodiments incorporate the interactions between multiple different system entities.

[0027] The amount of information provided by system monitoring can be extremely large, making direct storage and access of that information in memory impractical. However, the information has a large degree of redundancy. Redundancy can be found first in attributes, as each event record includes not only the involved entities but also the attributes of these entities. Storing the attributes repeatedly for each event is redundant. Second, saving events that involve the same entities repeatedly with only time stamps changing between them is redundant. Third, storage of records that are irrelevant to the detection of intrusion attacks is unneeded.

[0028] The graph model is therefore generated by block 402 to incorporate the meaningful information from the monitoring data in a compressed way. The graph model is represented as a directed graph $G = (V, E, T)$, where T is a set of timestamps, $E \subset V \times V \times T$ is the set of edges, and $V = F \cup P \cup U \cup S$ is the set of vertices, where F is the set of files residing in the computer system, P is the set of processes, U is the set of UNIX® sockets, and S is the set of Internet sockets. For a specific edge (v_i, v_j) in E , $T(v_i, v_j)$ denotes the set of timestamps on the edge. For each event e , if the corresponding edge already exists in G , a timestamp t is added to the edge timestamps. If not, block 402 builds such an edge to G with the timestamp set $T(v_i, v_j) = \{t\}$. In this structure, the attribute

values for a unique entity are stored only once. For each event sequence of a length l , there is a corresponding path through G of l edges.

[0029] To extract the most suspicious paths from the graph G , a naïve approach would be to explore all of the existing paths. However, it is impractical to enumerate all the possible paths from a densely connected graph. To provide guidance to candidate searching in block 406, block 404 generates a set of valid path patterns B . Only those paths that conform to the valid path patterns are related to potential attacks—others may be discarded.

[0030] Each path pattern B of length l includes l entities and/or entity types. Thus the path pattern B may include both specific entities (e.g., particular files) as well as more general designations of entity type in the same path. B is determined to be a valid path pattern only if there exists at least a path p in G that is consistent with B . Considering that l may be a small number, all of the possible paths can be enumerated. Searching the graph G permits all of the valid patterns to be extracted.

[0031] The valid path patterns B may be, for example, generated by experts using their experiences with previous intrusion detection attacks. However, it may be difficult to obtain an accurate and complete set of valid path patterns from such experts. As such, path patterns may be automatically generated. Each entity is set as a specific system entity type. For all the paths that correspond to information leakage, they must begin with a File entity (F) and end with an Internet socket entity (I). Given a path $p \in G$ and a path pattern B of G , $p[i]$ and $B[i]$ represent the i^{th} node in p and B respectively. The path p is therefore consistent with B , denoted as $p \prec B$, if p and B have the same length and if, for each i , $p[i] \in B[i]$ (i.e., the specific entity $p[i]$ belongs to the entity type $B[i]$). Then B is a valid path pattern if there exists at least one path p in G such that $p \prec B$. In one example

following the above constraints, there are four potential path patterns of length three using the four entity types described above: $\{F, F, I\}$, $\{F, P, I\}$, $\{F, U, I\}$, and $\{F, I, I\}$. Because only a Process node can connect a File node to an Internet socket node. In this manner, all of the valid patterns in G can be discovered.

[0032] Based on the generated valid path patterns B , block 406 searches for paths in the multipartite graph satisfying the patterns. Given an event sequence $seq = \{e_1, e_2, \dots, e_r\}$, there must be an equivalent path $p = \{v_1, v_2, \dots, v_{r+1}\}$ in the graph G . Because the events follow time order, a time order constraint is applied to candidate path searching. By applying the path patterns and the time order constraint to a breadth first search, the candidate paths can be discovered with a one-time scan of G . The candidate paths C are defined as:

$$C = \{p | p \in G, \exists b \in B \text{ s.t. } p < b\}$$

[0033] Even with a filtering policy based on path patterns and the time order constraint, there can still be a large number of candidate paths remaining in the graph G , most of which will be related to normal behavior. As such, the present embodiments extract suspicious paths from the larger set of candidate paths.

[0034] A candidate path is determined to be suspicious by block 408 if, in the path, the involved entities behave differently from their normal roles. In a computer system, the information senders and receivers are identified as entity roles. The sender and receiver score should be accurately learned from the computer system, as they are used to set the profile of normal behavior. To achieve this, a random walk is applied to the graph G . From G , an $N \times N$ square transition matrix A is calculated as:

$$A[i][j] = \text{prob}(v_i \rightarrow v_j) = \frac{|T(v_i, v_j)|}{|\sum_{k=1}^N |T(v_i, v_k)||}$$

where N is the total number of entities and $T(v_i, v_j)$ is the set of timestamps on which the event between v_i and v_j has ever happened. Thus $A[i][j]$ denotes the probability that the information flows from v_i to v_j in G .

[0035] As A is the matrix representation of the multipartite graph G , A can also be denoted as:

$$A = \begin{array}{c|ccccc} & P & F & S & U \\ \hline P & 0 & A^{P \rightarrow F} & A^{P \rightarrow S} & A^{P \rightarrow U} \\ F & A^{F \rightarrow P} & 0 & 0 & 0 \\ S & A^{S \rightarrow P} & 0 & 0 & 0 \\ U & A^{U \rightarrow P} & 0 & 0 & A^{U \rightarrow U} \end{array}$$

where zero represents a zero sub-matrix, and where the arrow operator indicates a direction of information flow. For example, $P \rightarrow F$ indicates a flow of information from process to file. It should be noted that the non-zero sub-matrices of A only appear between processes and files and between processes and sockets, but not between respective processes, because process-process interactions do not come with interaction flow. These are constraints set by UNIX® systems.

[0036] Letting X be a sender score vector, with $X[i]$ denoting v_i 's sender score, and letting Y be the receiver score vector, then each entity's sender and receiver scores can be iteratively generated as:

$$\begin{cases} X_{m+1}^T = (A \times A^T) \times X_{m-1}^T \\ Y_{m+1}^T = (A^T \times A) \times Y_{m-1}^T \end{cases}$$

with initial vectors X_0 and Y_0 being randomly generated and with m referring to the number of the current iteration. Stated generally, an entity that sends information to a large number of entities having high receiver scores is itself an important information sender, and similarly an entity that receives information from a large number of entities having high sender scores is an important information receiver. The sender and receiver scores for an

entity are therefore iteratively calculated by accumulating the receiver and sender scores related to the entity. For example, the file `/etc/passwd` on a UNIX® system will have a high sender score and a low receiver score, because it is sent to many processes for checking access permissions, but it is rarely modified.

[0037] A result of this iterative refinement is that the learned score values will depend on the initial score values. However, the effect of the initial score values can be eliminated using the steady state property of the matrix. Given a general square matrix M and a general vector π , the general vector π can be repeatedly updated as:

$$\pi_{m+1}^T = M \times \pi_m^T$$

A convergence state is possible such that $\pi_{m+1} \approx \pi_m$ for sufficiently large values of m . In this case, there is only one unique value which can reach the convergence state:

$$\pi_n^T = M \times \pi_n^T$$

The convergence state has the property that the converged vector is only dependent on the matrix M , but is independent from the initial vector value π_0 .

[0038] To reach the convergence state, the matrix M needs to satisfy two conditions: irreducibility and aperiodicity. A graph G is irreducible if and only if, for any two nodes, there exists at least one path between them. The period of a node is the minimum path length from the node back to itself, and the period of a graph is the greatest common divisor of all of the nodes' period values. A graph G is aperiodic if and only if it is irreducible and the period of G is 1.

[0039] As the system graph G is not always strongly connected, the above-described iteration will not always reach convergence. To ensure convergence, a restart matrix R is

added, which is an $N \times N$ square matrix, with each cell value being $\frac{1}{N}$. A new transition matrix $\bar{A}$ is defined as:

$$\bar{A} = (1 - c) \times A + c \times R$$

where c is a value between 0 and 1 and is referred to as the restart ratio. $\bar{A}$ is guaranteed to be irreducible and aperiodic, providing converged sender score and receiver score vectors. The convergence rate can be controlled by controlling the restart rate value. One exemplary value for a number of iterations to use to ensure convergence is about 10.

[0040] Based on the sender and receiver score, and given a path p , the anomaly score for the path is calculated as:

$$Score(p) = 1 - NS(p)$$

$$NS(p) = 1 - \prod_{i=1}^r X(v_i)A(v_i, v_{i+1})Y(v_{i+1})$$

[0041] As noted above, the anomaly scores for paths of different lengths have different distributions. Therefore, to compare the suspiciousness of paths of different lengths, a transformation is performed to put paths of different length on equal footing. The path anomaly score can have an arbitrary distribution and is generally not a normal distribution. The suspiciousness of a path of length r can be defined as:

$$susp(p|r) = Prob(T(Score(p')) \leq T(Score(p)) | |p'| = r)$$

where T is the normalization function.

[0042] The top k suspicious paths are those having the largest suspiciousness score. Mathematically, it is feasible to get the transformation to convert a normal distribution to any other distribution, but it is difficult to get the inverse function. To solve this problem, a Box-Cox power transformation is used as the normalization function. In particular, letting

$Q(r)$ denote the set of anomaly scores calculated from paths of length r , for each score $q \in Q(r)$,

$$T(q, \lambda) = \begin{cases} \frac{q^\lambda - 1}{\lambda} : \lambda \neq 0 \\ \log(q) : \lambda = 0 \end{cases}$$

where λ is a normalization parameter. Different values for λ yield different transformed distributions. The aim is to select a value of λ that yields a normalized distribution as close as possible to the normal distribution (i.e., $T(B, \lambda) \sim N(\mu, \sigma^2)$).

[0043] The top k suspicious paths are not considered related to intrusion attacks unless they are sufficiently distinctive from the normal paths. To measure the deviation from suspicious paths from normal paths, the t -value is calculated between the two groups of paths. Due to the large number of normal paths, an efficient solution based on a Monte Carlo simulation is used to calculate the expectation and variance from a sample of relatively small size without computing the summation.

[0044] Once a suspicious path has been detected, the host-level analysis module 43 provides information regarding the anomaly, generating a report that may include one or more alerts. The anomaly fusion module 44 integrates these host-level alerts with other host- and network-level anomalies and automatically filters out false alarms. The resulting list of anomalies is provided via visualization module 45 to a user. In an alternative embodiment, certain anomalies or classes of anomalies may be addressed automatically, for example by deploying security countermeasures or mitigations. In one specific example, an automatic response to a detected anomaly may be to shut down a device showing the anomalous behavior until it can be reviewed by an administrator.

[0045] Referring now to FIG. 5, pseudo-code for discovering the top k suspicious paths “SP.” Sender and receiver score vectors X and Y are created using the random walk process. A queue of files F_X is sorted according to descending X and a second queue of files F_Y is sorted descending Y . Queues of processes P_X and P_Y , queues of UNIX® sockets U_X and U_Y , and queues of internet sockets S_X and S_Y are created in the same way. The paths are then processed to find a path set that conforms to the event sequence pattern and time constraint.

[0046] Embodiments described herein may be entirely hardware, entirely software or including both hardware and software elements. In a preferred embodiment, the present invention is implemented in software, which includes but is not limited to firmware, resident software, microcode, etc.

[0047] Embodiments may include a computer program product accessible from a computer-usable or computer-readable medium providing program code for use by or in connection with a computer or any instruction execution system. A computer-usable or computer readable medium may include any apparatus that stores, communicates, propagates, or transports the program for use by or in connection with the instruction execution system, apparatus, or device. The medium can be magnetic, optical, electronic, electromagnetic, infrared, or semiconductor system (or apparatus or device) or a propagation medium. The medium may include a computer-readable storage medium such as a semiconductor or solid state memory, magnetic tape, a removable computer diskette, a random access memory (RAM), a read-only memory (ROM), a rigid magnetic disk and an optical disk, etc.

[0048] Each computer program may be tangibly stored in a machine-readable storage media or device (e.g., program memory or magnetic disk) readable by a general or special purpose programmable computer, for configuring and controlling operation of a computer when the storage media or device is read by the computer to perform the procedures described herein. The inventive system may also be considered to be embodied in a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner to perform the functions described herein.

[0049] A data processing system suitable for storing and/or executing program code may include at least one processor coupled directly or indirectly to memory elements through a system bus. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code to reduce the number of times code is retrieved from bulk storage during execution. Input/output or I/O devices (including but not limited to keyboards, displays, pointing devices, etc.) may be coupled to the system either directly or through intervening I/O controllers.

[0050] Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modem and Ethernet cards are just a few of the currently available types of network adapters.

[0051] Referring now to FIG. 6, an exemplary processing system 600 is shown which may represent the analytic server 30, the intrusion detection system 31, and/or the host level analysis module 43. The processing system 600 includes at least one processor (CPU)

604 operatively coupled to other components via a system bus 602. A cache 606, a Read Only Memory (ROM) 608, a Random Access Memory (RAM) 610, an input/output (I/O) adapter 620, a sound adapter 630, a network adapter 640, a user interface adapter 650, and a display adapter 660, are operatively coupled to the system bus 602.

[0052] A first storage device 622 and a second storage device 624 are operatively coupled to system bus 602 by the I/O adapter 620. The storage devices 622 and 624 can be any of a disk storage device (e.g., a magnetic or optical disk storage device), a solid state magnetic device, and so forth. The storage devices 622 and 624 can be the same type of storage device or different types of storage devices.

[0053] A speaker 632 is operatively coupled to system bus 602 by the sound adapter 630. A transceiver 642 is operatively coupled to system bus 602 by network adapter 640. A display device 662 is operatively coupled to system bus 602 by display adapter 660.

[0054] A first user input device 652, a second user input device 654, and a third user input device 656 are operatively coupled to system bus 602 by user interface adapter 650. The user input devices 652, 654, and 656 can be any of a keyboard, a mouse, a keypad, an image capture device, a motion sensing device, a microphone, a device incorporating the functionality of at least two of the preceding devices, and so forth. Of course, other types of input devices can also be used, while maintaining the spirit of the present principles. The user input devices 652, 654, and 656 can be the same type of user input device or different types of user input devices. The user input devices 652, 654, and 656 are used to input and output information to and from system 600.

[0055] Of course, the processing system 600 may also include other elements (not shown), as readily contemplated by one of skill in the art, as well as omit certain

elements. For example, various other input devices and/or output devices can be included in processing system 600, depending upon the particular implementation of the same, as readily understood by one of ordinary skill in the art. For example, various types of wireless and/or wired input and/or output devices can be used. Moreover, additional processors, controllers, memories, and so forth, in various configurations can also be utilized as readily appreciated by one of ordinary skill in the art. These and other variations of the processing system 600 are readily contemplated by one of ordinary skill in the art given the teachings of the present principles provided herein.

[0056] The foregoing is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that those skilled in the art may implement various modifications without departing from the scope and spirit of the invention. Those skilled in the art could implement various other feature combinations without departing from the scope and spirit of the invention. Having thus described aspects of the invention, with the details and particularity required by the patent laws, what is claimed and desired protected by Letters Patent is set forth in the appended claims.

WHAT IS CLAIMED IS:

1. A method for detecting malicious processes, comprising:

 modeling system data as a graph comprising vertices that represent system entities and edges that represent events between respective system entities, each edge comprising one or more timestamps corresponding respective events between two system entities;

 generating a set of valid path patterns that relate to potential attacks; and

 determining one or more event sequences in the system to be suspicious based on the graph and the valid path patterns using a random walk on the graph.
2. The method of claim 1, wherein the system entities comprise files in the system, processes in the system, UNIX® sockets in the system, and Internet sockets in the system.
3. The method of claim 1, wherein the valid patterns are determined based on properties of the system entity types.
4. The method of claim 3, wherein the valid path patterns are determined based on definitions provided by security experts according to their experiences based on previous intrusion attacks.
5. The method of claim 1, wherein determining the one or more event sequences to be suspicious comprises performing a breadth first search of candidate paths within the graph.

6. The method of claim 5, wherein the breadth first search comprises a time order constraint based on the edge timestamps.
7. The method of claim 1, wherein determining the one or more event sequences to be suspicious comprises determining that the entities on an edge deviate from the normal roles for those entities.
8. The method of claim 7, wherein determining that entities deviate from the normal role for those entities comprises determining a sender score for a sender entity and a receiver score for a receiver entity.
9. The method of claim 8, wherein determining one or more event sequences to be suspicious comprises calculating an anomaly score based on the sender score and receiver score for each entity in each event sequence.
10. The method of claim 9, wherein determining one or more event sequences to be suspicious comprises normalizing anomaly scores using a Box-Cox power transformation.
11. A system for detecting malicious processes, comprising:

a modeling module configured to model system data as a graph that comprises vertices that represent system entities and edges that represent events between respective

system entities, each edge comprising one or more timestamps corresponding respective events between two system entities; and

a malicious process path discovery module comprising a processor configured to generate a set of valid path patterns that relate to potential attacks and to determine one or more event sequences in the system to be suspicious based on the graph and the valid path patterns using a random walk on the graph.

12. The system of claim 11, wherein the system entities comprise files in the system, processes in the system, UNIX® sockets in the system, and Internet sockets in the system.

13. The system of claim 11, wherein the malicious process path discovery module is further configured to determine valid patterns based on properties of the system entity types.

14. The system of claim 13, wherein the malicious process path discovery module is further configured to determine valid path patterns based on definitions provided by security experts according to their experiences based on previous intrusion attacks.

15. The system of claim 11, wherein the malicious process path discovery module is further configured to perform a breadth first search of candidate paths within the graph.

16. The system of claim 15, wherein the breadth first search comprises a time order constraint based on the edge timestamps.
17. The system of claim 11, wherein the malicious process path discovery module is further configured to determine whether the entities on an edge deviate from the normal roles for those entities.
18. The system of claim 17, wherein the malicious process path discovery module is further configured to determine a sender score for a sender entity and a receiver score for a receiver entity.
19. The system of claim 18, wherein the malicious process path discovery module is further configured to calculate an anomaly score based on the sender score and receiver score for each entity in each event sequence.
20. The system of claim 19, wherein the malicious process path discovery module is further configured to normalize anomaly scores using a Box-Cox power transformation.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

Require: $G = (V, E, T)$, where $V = F \cup P \cup U \cup I$, k ,
 Ensure: SP that contains the top- k suspicious paths of G .
 Initialize SP as an empty priority queue;
 Apply random walk on G to calculate sender score vector x and receiver score vector y ;
 Let F_X be the files sorted in descendant order of the sender score. Let F_Y be the files sorted in descendant order of the receiver score.
 Create P_X, P_Y, U_X, U_Y, S_X and S_Y in the same way.
 Let E' be the edges sorted descendingly by $A[i][j]$
 while $P_X, P_Y, U_X, U_Y, S_X, S_Y$ and E' are not empty
 do
 $f_x = F_X.pop(), f_y = F_Y.pop()$
 $p_x = P_X.pop(), p_y = P_Y.pop()$
 $u_x = U_X.pop(), u_y = U_Y.pop()$
 $s_x = S_X.pop(), s_y = S_Y.pop(), e = E'.pop()$
 Assume p is the path involving $f_x, f_y, p_x, p_y, u_x, u_y, s_x, s_y, e'$.
 if $Score(p) \leq \min\{Score(p') | p' \in SP\}$ then
 break
 else
 Find the path set $\mathcal{P}$ such that every path p confirms to the event sequence pattern and time constraint, and p involves at least one node in $\{f_x, f_y, p_x, p_y, u_x, u_y, s_x, s_y\}$ or e' .
 for all $p \in \mathcal{P}$ do
 if $Score(p) > \min\{Score(p') | p' \in SP\}$ then
 Insert p into SP
 if SP contains $k+1$ paths then
 Remove the $(k+1)$ -th path from SP .
 end if
 end if
 end for
 end if
 end while
 return SP

FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/043040**A. CLASSIFICATION OF SUBJECT MATTER****G06F 21/55(2013.01)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 21/55; G06F 11/30; G06Q 30/00; G06F 17/30; G06F 12/14; H04L 29/06; G08B 23/00; G06F 15/18

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: graph, suspicious, valid, path, detect, malicious, edge, timestamp, box cox, potential

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2007-0209074 A1 (THAYNE RICHARD COFFMAN) 06 September 2007 See paragraphs [0013], [0021], [0040], [0051], [0119], [0127], [0131]; and claims 2, 10, 15; and figure 1.	1-20
Y	EP 1919162 A2 (JUNIPER NETWORKS, INC.) 07 May 2008 See paragraphs [0004], [0022], [0037]; and figure 3.	1-20
Y	US 2010-0312669 A1 (PETER HANUS et al.) 09 December 2010 See paragraphs [0007], [0039]; and figure 1A.	10, 20
A	US 2008-0178293 A1 (ARTHUR KEEN et al.) 24 July 2008 See paragraphs [0009], [0022]; and figure 1.	1-20
A	US 2009-0222435 A1 (REID MARLOW ANDERSEN et al.) 03 September 2009 See paragraphs [0005], [0024]; and figure 1.	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

20 October 2016 (20.10.2016)

Date of mailing of the international search report

21 October 2016 (21.10.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office


189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

CHIN, Sang Bum

Telephone No. +82-42-481-8398



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/043040

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007-0209074 A1	06/09/2007	US 7624448 B2 WO 2008-021585 A2 WO 2008-021585 A3	24/11/2009 21/02/2008 19/06/2008
EP 1919162 A2	07/05/2008	CN 101175078 A CN 101175078 B EP 1919162 A3 US 2008-0101234 A1 US 7768921 B2	07/05/2008 26/10/2011 17/06/2009 01/05/2008 03/08/2010
US 2010-0312669 A1	09/12/2010	US 2006-0241901 A1 US 7773812 B2	26/10/2006 10/08/2010
US 2008-0178293 A1	24/07/2008	US 8161550 B2	17/04/2012
US 2009-0222435 A1	03/09/2009	US 8010482 B2 WO 2009-111212 A2 WO 2009-111212 A3	30/08/2011 11/09/2009 05/11/2009