

TitleCOMPUTER SECURITY EVENT MANAGEMENT SYSTEMField of the Invention

The present invention relates to an event monitoring system for monitoring computer security events.

Background of the Invention

Computer related security breaches are becoming more common. Just a few examples are denial of service attacks, viruses, worms and hacking. In addition the consequences and costs of security breaches and attacks is ever increasing. However, simply providing a firewall machine alone is insufficient. It is just as important to protect the computer system from backdoor entry and misuse by authorised personal. Furthermore, even with security products in place such as firewalls and intrusion detection systems, antivirus software and the like, these systems generate vast logs of data in a variety of formats that is not useful for macro analysis and often results in after the fact detection. This makes real time intervention difficult or impossible.

Summary of the Present Invention

In accordance with the first aspect of the present invention there is provided a computer security event monitoring comprising: receiving means for receiving externally
The incident records generating means may comprise a plurality of first modules, each first
module designed to execute at least one function for processing an event alert of one type.
The receiving means may comprise a plurality of second modules, each second
module
designed to execute at least one function for receiving an event alert of one type.
The receiving means may be operable to receive the security event alerts from an
external intrusion detection system.
The receiving means may be operable to receive the security event alerts from an
external firewall.
The system may further comprise the event reaction means, and the reaction means may
be responsive to the incident records and remotely handles a security event based on the

incident record.

The incident records generating means may perform an assessment of the severity of the event alert and prioritises incident records accordingly.

The assessment may be made by checking the alert against one or more rules for determining severity.

The assessment may be made by checking the alert against one or more rules for determining a

reaction to the security event that triggered the alert.

Any one or more of the receiving means, the incident records generating means and transmitting means may comprise a software application running on a dedicated computer

connected by a network to a computer system being monitored.

The computer system may be connected to the dedicated computer by a local area network.

The computer system may include a plurality of computers connected by a wide area network.

The computer system may be a plurality of computers connected by a global computer network.

In accordance with a second aspect of the present invention, there is provided a method of computer security event monitoring comprising the steps of:

receiving externally generated security event alerts of different types;

generating incident records in a unified format based on the receipt security event alerts;

transmitting the incident records to an event reaction unit remotely located from the computer security event monitoring system.

The security alert may be received from an external intrusion detection system.

The security alert may be received from an external firewall.

The method may further comprise assessing the severity of the event alerts for prioritising incident records accordingly.

The method may further comprise checking the event alerts against one or more rules for determining the reaction to the security events associated with the event alert.

Brief Description of the Drawings

In order to provide a better understanding, preferred embodiments of the present invention will now be described in detail, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a schematic representation of a computer security event monitoring system in accordance with the present invention;

Figure 2 is a schematic representation of an event manager of the system shown in Figure 1;

Figure 3 is a schematic representation of a first level of the event manager of Figure 2;

Figure 4 is a schematic representation of a second level of the event manager of Figure 2;

Figure 5 is a schematic representation of a third level of the event manager of Figure 2;

Figure 6 is a schematic representation of a fourth level of the event manager of

- 4 -

Figure 2;

Figure 7 is a schematic representation of a master slave relationship between a plurality of event managers; and

Figure 8 is a schematic representation of a hand-off process between slave event managers.

Referring to Figure 1 there is shown a computer security event monitoring system 10 that includes a plurality of triggers 12 that provide security event alerts to an event manager 14.

The event manager 14 converts the event alerts to incident records and stores these in a database 16 as well as forwarding them to an event reaction means in the form of a global command centre console 18.

The triggers are in the form of an intrusion detection system, a firewall program, antivirus software, an application software and/or operating systems logs. Each of the triggers 12 are configured to send an alert to the event manager 14 if a security event occurs. A security event is an event that is a real or potential threat to the proper working and use of a computer system. For example, a firewall protecting a corporate network that suffers an Internet Control Message Protocol (ICMP) flood and registers a list of violations will trigger an alarm. An intrusion detection system that detects a string of commands targeted at a corporate mail server for the purpose of exploiting administrator access will trigger an alarm. Antivirus software residing on a file server that detects an unknown variant of a virus and is unable to cleanse the infected file will trigger an alarm. Operating systems, such as Windows NT or Unix, that have a login violation in the middle of the night will trigger an alarm. Security flaws can be caused by software defects (bugs), security holes, incorrect system setup or deliberate backdoor openings. A trigger is aimed at detecting poor security or security breaches by detecting misuse or patterns that indicate real or potential security problems. Sometimes, only a macro analysis will reveal a security problem.

The event manager receives event alerts from a distributed security monitoring environment and is able to analyse various forms of communication and in particular

- 5 -

TCP/IP traffic such as raw TCP and UDP connections as well as other protocols such as SNMP, SMTP, HTTP, FTP and other application non-streaming protocols and non-application streaming protocols. These are converted into a uniform format referred to as transportable incident format (TIF) records that are stored in the database 16. Specific
5 modules of the event manager may be installed depending on the computer systems requirements.

The TIF records may also be assessed to determine an appropriate reaction to the event that triggered the alert. A response may be based on standing orders provided by a set of
10 rules. In more serious cases human intervention can be instigated to counteract a security breach. Thus real time continuous analysis of the security of a computer system can be monitored, thereby allowing the ability to provide a real time assessment defence/reaction to a security incident. Furthermore, storing alerts in a uniform format enables reporting capabilities, trend analysis, and vulnerability analysis of the recorded incidents.

15

Monitoring devices include firewalls that report rejected requests or an intrusion detection system (IDS) that monitors the activity of a computer network. Generally the firewalls and IDSs are configured to send alerts to the event manager by providing a directly resolvable IP address rather than relying on DNS resolution. This is done to prevent reliance on the
20 DNS service and to reduce vulnerability in the system.

Referring to Figure 2, the event manager 14 receives alerts 20 from monitoring devices 12 (triggers). These alerts are streamed into input/output (I/O) plug-ins 22 that communicate with the monitoring devices. An alert from a technology specific monitoring device is
25 received by each I/O plug-in 22, which forward the data to technology specific plug-ins 26. A technology specific plug-in 24 appropriate to the type of monitoring device interprets the data and converts it into TIF records. The plug-ins 26 then broadcast the TIF record.

30 The I/O and technology specific modules broadcast raw data and TIFs respectively, such that other modules can receive them, process them or ignore them as per their

- 6 -

configuration.

There is a push/pull mechanism in this communication. The various modules push the data with only the module that is configured to process the specific data pulling the data, while other modules ignore the data. For example, the TIF data broadcast by the TS modules may be pulled by storage modules, or an inference engine module. When the inference engine is used then the storage module ignores the TS module broadcast and only pulls inference engine broadcasts.

Preferably, but optionally, an inference engine 30 is provided. If the inference engine 30 is provided, it receives the TIF record broadcast and forwards it to a secure transmission protocol (SXP) plug-in 28. The inference engine 30 also filters TIF records according to the degree of severity of the security event that triggered the alert, with only serious incidents being forwarded to a storage plug-in 32 for storage. The SXP plug-in 28 ignores TIF records broadcast by the plug-in 26 and only accepts TIF records broadcast by the inference engine 30. The storage plug-in 32 acts similarly.

If the inference engine 30 is not used, the SXP plug-in 28 receives the TIF records from the plug-ins 26. The storage plug-in 32 also receives the TIF records directly from the plug-ins 26.

The SXP plug-in 28 forwards the TIF record to a remote SPX plug-in 34 of the reaction means 18 by secure TCP/IP communication. Again it is preferred that an inference engine 34 be used at the reaction means 18. Again the inference engine 34 is able to act according to a set of rules to categorise alerts with false alarms or low priority/severity alerts given less importance. Filtered TIF records are then provided to a storage plug-in 38. If the inference engine 36 is not used then unfiltered TIF records are provided to the storage plug-in 38. The storage plug-in 38 converts the TIF records into database storable format for storage in the database 16.

30

The reaction means 18 is configured to detect false alarms according to a computer system

- 7 -

specific checklist. It is able to reject false alarms as requiring no action. If a real security event is detected a number of counter measures can be initiated. These include shutting of an IP port or blocking an intruder by adding the intruders IP address to the computer system's routers access control list. This will immediately block traffic from the IP address specified. In addition, collation and analysis of incidents can be conducted. Some security problems may only be detected by macro analysis.

Incident types can be defined according to the nature of the alert. Such alerts may includes probes, FTP application attacks or remote procedure call attack attempts.

10 Multiple or repeated threats from the same IP source can be monitored and acted upon.

Incidents of over reaction to a false alarm are reduced with some low level security problems being placed under observation to see if further recurrences occur. An automated inference engine can deal with some security threats, but if the inference engine cannot handle a particular security situations then it may be escalated in priority to dealt with a consultant. The consultant may, for example, contact with a particular owner of a domain or ISP to deal with the incident.

Records of incidents stored in the database 16 can be analysed to provide a correlation of security incidents rather than multiple discrete alerts. A log of security incidents may be useful for forensic analysis in case law enforcement is required. In addition, the database 16 may be useful in profiling of security incidents.

The event manager 14 is provided in a modular design with four levels, each level having one or more plug-ins/modules. The event manager can be configured to operate in a distributed environment with a plurality of nodes, each node dedicated to perform a different function. One event manager node is set up as a master to distribute the load over various slave nodes.

30 Referring to Figure 3, the first layer of modules comprises input/output plug-ins 22 that receive data in various kinds of format including , but not limited to, SMTP, TCP, UDP,

- 8 -

HTTP and FTP. Each module is dedicated to a particular input stream protocol. The plug-in removes the raw data and broadcasts it to technology specific modules 26 in the second layer. In addition there is an application layer proxy (ALP) module and a content integrity agent (CIA) module. The ALP module intercedes between host server and connecting clients and is capable of filtering out traffic (IP packets) based on a user defined wordlist (dictionary) or complex regular expression list. The ALP module operates at the highest network layer, namely the application layer. The monitoring and packet rejection hinges on a list of forbidden words which can be specified. The CIA protects publicly available data, such as a website, by exposing a monitored folder to public access and if tampered with, synchronises it with an insulated master folder.

Referring to Figure 4, a second technology specific (TS) module layer is shown including the TS plug-ins 26. Each TS modules 24 converts raw input (alerts) into transportable incident format and broadcasts it to the inference engine or SPX and storage modules.

15

Raw alerts mainly come from two types of device, these being firewalls and intrusion detection systems. When an I/O module receives alerts it broadcasts the alerts to all TS modules, which will listen at the same time. Each module takes the relevant alerts and ignores the rest of them. Thus all the modules collectively process the entire data.

20

Intrusion detection system (IDS) modules deal with alerts from various intrusion detection systems. Each TS module decodes and processes an alert from an IDS product or a firewall product of a particular software vendor. ISS decoder module processes alerts received from "ISS Real Secure". NFR decoder processes alerts from "Network Flight Recorder". SW3 decoder processes alerts from "CA Session Wall-3".

25

FW1 decoder process alerts from "Checkpoint Firewall 1". The net screen SMTP decoder processes alerts from "Net Screen Firewall". Raptor flight decoder processes alerts from "Raptor Firewall" for Windows NT or Sun Solaris. Sonic wall decoder processes alerts from "3 COM Sonic Wall". Watch guard SMTP decoder processes alerts from "Watch Guard Firebox". WebTrends SysLog decoder processes alerts from "Web Trends

30

- 9 -

SysLog".

Referring to Figure 5, a third level 70 of the event manager is shown. This includes the inference engine (IE). The IE performs a preliminary check on the data to determine whether the event should be logged or dropped based on a set of user defined rules. A basic inference engine outputs TIF data without altering it only removing unwanted alerts. Where an event manager is configured with an inference engine module, the TIF output from TS modules is ignored by storage and SXP modules. Only post inference engine TIF data (those only containing serious alerts) is accepted for processing by the storage or SXP modules. The SPX module transmits TIF records to the global command centre 18 database in an encrypted format, such as secure TCP.

Referring to Figure 6, a fourth layer 80 of the event manager is shown which includes a storage (ST) module that supports use of common databases for storing incidents. It converts processed TIF events into data specific format and is easily retrievable table in database, such as Microsoft SQL server.

General modules may be provided including a log master that logs all incoming traffic, if configured to do so, and a device health check module that automatically checks the status of other modules, and a device health check module that automatically checks the status of other modules.

Referring to Figure 7, when a distributed event manager environment is used, as indicated by 90, one of the event managers 92 is configured as a master event manager. Other event managers 94 and 96 are configured as slave managers. The master event manager distributes raw data equally among slaves and the local TS module within itself. The TS modules only accept redirected raw data being broadcasted by I/O modules. The TS modules of the master and each of the slaves then completes processing of the raw data by converting all the data into TIF records which then may be handed off.

30

A handoff occurs by providing an IP addresses specifying machines to which incidents are

- 10 -

to be forwarded. A handoff is useful where a client side event manager has limited processing such as of I/O and TS modules only. If SXP is configured to hand-off the processed incidents (TIF) to a designated global command centre event manager further processing and storage can be performed.

5

A re-direct can be provided that passes raw data without performing any processing so that raw data can be shared amongst several event managers. A re-direct is suitable for environments with a large amount of raw incidents, whereby a single event manager is insufficient. This is typically used in an event manager farm with a master and multiple slave event managers sharing the processing load.

10

The master event manager receives raw data incidents from the security devices through the I/O modules and transmits the incident to the machines listed under a re-direct list via the SXP module. All the event managers to which the data is re-directed are configured with exactly the same modules, typically, several technology specific modules. The decision of which event manager to transmit to is done fairly and equally. The master event manager keeps a journal of what has been transmitted to each event manager in the group and transmits the raw incidents to one which receive the least number of raw incidents. Upon receiving the data the event manager that data is redirected to processes the raw data, converting it into TIF records. TIF records may be further processed by an inference engine configured on the same event manager or may be stored in a repository by an ST module. They can also be handed off the another event manager module to, for example, be categorised by an inference engine.

15

20

Referring to Figure 8, TIF data is handed off to a master event manager's SXP module. The TS module within the master broadcasts TIFs and the master passes all data to the storage module and inference engine. The hand-off event manager can be another slave attached to the master or the master itself. Alternatively there may be two hand-off IP address which could be the slave or the master. If the first hand-off fails then the second hand-off is attempted.

25
30

- 11 -

The skilled addressee will appreciate that the present invention has advantages over the prior art. The present invention allows security events to be dealt with in real time, as well as providing a means of analysis of security events for trend analysis, vulnerability analysis or forensic analysis in case law enforcement is required.

5

The skilled addressee will appreciate that modifications and variations can be made to the present invention without departing from the basic inventive concept. Such modifications may include various forms of security event triggering devices, different methodologies and rule sets for dealing with and prioritising incidents. The system may provide a local
10 security system or the security system may be distributed or may be entirely remotely based. Such modifications are intended to fall within the scope of the present invention, the nature of which is to be determined from the foregoing description and appended claims.

-12-

Claims:

1. A computer security event monitoring comprising:
receiving means for receiving externally generated security event alerts of different types;
incident records generating means for generating incident records in an unified format based on the received security event alerts; and
transmitting means for transmitting the incident records to an event reaction unit remotely located from the computer security event monitoring system.
2. A computer security event monitoring system according to claim 1, wherein the incident records generating means comprises a plurality of first modules, each first module designed to execute at least one function for processing an event alert of one type.
3. A computer security event monitoring system according to claim 1 or 2, wherein the receiving means comprises a plurality of second modules, each second module designed to execute at least one function for receiving an event alert of one type.
4. A computer security event monitoring system according to any one of claims 1 to 3, wherein the receiving means is operable to receive the security event alerts from an external intrusion detection system.
5. A computer security event monitoring system according to any one of claims 1 to 4, wherein the receiving means is operable to receive the security event alerts from an external firewall.
6. A computer security event monitoring system according to any one of claims 1 to 5, further comprising the event reaction means, and the reaction means is responsive to the incident records and remotely handles a security event based on the incident record.

-13-

7. A computer security event monitoring system according to any one of claims 1 to 6, wherein the incident records generating means performs an assessment of the severity of the event alert and prioritises incident records accordingly.

8. A computer security event monitoring system according to claim 7, wherein the assessment is made by checking the alert against one or more rules for determining severity.

9. A computer security event monitoring system according to claims 7 or 8, wherein the assessment is made by checking the alert against one or more rules for determining a reaction to the security event that triggered the alert.

10. A computer security event monitoring system according to any one of claims 1 to 9, wherein any one or more of the receiving means, the incident records generating means and transmitting means comprise a software application running on a dedicated computer connected by a network to a computer system being monitored.

11. A computer security event monitoring system according to claim 10, wherein the computer system is connected to the dedicated computer by a local area network.

12. A computer security event monitoring system according to claims 10 or 11, wherein the computer system includes a plurality of computers connected by a wide area network.

13. A computer security event monitoring system according to claims 10 or 11, wherein the computer system is a plurality of computers connected by a global computer network.

14. A method of computer security event monitoring comprising the steps of:
receiving externally generated security event alerts of different types;

-14-

generating incident records in an unified format based on the received security event alerts; and

transmitting the incident records to an event reaction unit remotely located from the computer security event monitoring system.

15. A method of computer security event monitoring according to claim 14, wherein the security alert is received from an external intrusion detection system.

16. A method of computer security event monitoring according to claims 14 or 15, wherein the security alert is received from an external firewall.

17. A method of computer security event monitoring according to any one of claims 14 to 16, the method further comprising assessing the severity of the event alerts for prioritising incident records accordingly.

18. A method of computer security event monitoring according to claim any one of claims 14 to 17, the method further comprising checking the event alerts against one or more rules for determining the reaction to the security events associated with the event alert.