

Approved for public release;
distribution is unlimited.

Title: AUTOMATED INFORMATION SYSTEM (AIS) ALARM SYSTEM

Author(s): William J. Hunteman

Submitted to: 20th National Information Security Conf.
Baltimore, MD
October 6, 1997

RECEIVED

MAY 05 1997

OSTI

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED *dy*

MASTER

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Type of Submission
Title

Paper
Automated Information System (AIS) Alarm System

Abstract

The Automated Information Alarm System is a joint effort between Los Alamos National Laboratory, Lawrence Livermore National Laboratory, and Sandia National Laboratory to demonstrate and implement, on a small-to-medium sized local area network, an automated system that detects and automatically responds to attacks that use readily available tools and methodologies. The Alarm System will sense or detect, assess, and respond to suspicious activities that may be detrimental to information on the network or to continued operation of the network. The responses will allow stopping, isolating, or ejecting the suspicious activities. The number of sensors, the sensitivity of the sensors, the assessment criteria, and the desired responses may be set by the using organization to meet their local security policies.

Author(s)

William Hunteman

Organizational Affiliation

University of California, Los Alamos National Laboratory

Telephone

505-667-0096

Fax

505-665-3456

E-mail

wjh@lanl.gov

Automated Information System (AIS) Alarm System

Abstract

The Automated Information Alarm System is a joint effort between Los Alamos National Laboratory, Lawrence Livermore National Laboratory, and Sandia National Laboratory to demonstrate and implement, on a small-to-medium sized local area network, an automated system that detects and automatically responds to attacks that use readily available tools and methodologies. The Alarm System will sense or detect, assess, and respond to suspicious activities that may be detrimental to information on the network or to continued operation of the network. The responses will allow stopping, isolating, or ejecting the suspicious activities. The number of sensors, the sensitivity of the sensors, the assessment criteria, and the desired responses may be set by the using organization to meet their local security policies.

1. Introduction

The Automated Information System (AIS) Alarm System project is developing a near-real-time intrusion detection and response system that can supplement or complement an information assurance program for protecting the Department of Energy's information resources. The Alarm System will support the detection, evaluation, and response to suspicious events and possible violations of security policies in local area networks using Unix System V and Microsoft Windows NT operating systems. The Alarm System is designed for use by a network administrator or security officer who has the skills and knowledge to perform system management activities. The Alarm System will not affect LAN users, except for possibly a slight degradation in performance, when responses to suspicious activity are undertaken. The system will automatically deploy a response that will either inform the system administrator or computer security personnel that unauthorized activity has been detected or stop, isolate, or eject the unauthorized activity. The Alarm System products will be developed as a single integrated set of components that can be selected by the using organization. The user may also select the individual sensors and responses to be installed within each family of components.

1.1. Current Environment

The current environment for detecting suspicious activity is based largely on audit trail based intrusion detection. A substantial number of commercial off-the-shelf (COTS) and government off-the-shelf (GOTS) tools have been developed for analyzing audit trail and other similar information. Developments in commercial firewalls have added the ability to detect certain forms of network attack. Application-level filters in the firewall system can now inspect certain high risk protocols, searching for signatures of known intrusive behavior. These products require the collection of substantial information and recognize an attack by identifying a signature or pattern of activities. Because these products require the collection of log information or packets from a network, the recognition and reporting of an attack is often delayed until after the attack has been initiated or actually completed. Although work is being done to improve the timeliness of recognition and reporting, these products still require the collection of enough information to recognize the signature or pattern of an attack. Many of these existing tools and products are expected to complement or be incorporated into the Alarm System by functioning as additional sensors.

Other government activities somewhat related to the Alarm System project include systems to monitor the security status of individual information systems from a single control facility. These activities are based on the assumption that the hardware and software of the monitored information systems is well-known and static. The monitoring activities are designed to address needs different from the Alarm System project. These activities are based on the

assumption that once a system is secure, periodic monitoring will be sufficient to detect a change in the security posture of the monitored system. With Alarm System, such activities are complementary and remain useful as part of a defense-in-depth strategy, and such products can be incorporated as sensor or response elements.

1.2. Alarm System Goals

The Alarm System design goals include

- Perform detection of suspicious events and initiate responses based on event recognition (not audit trails) in near-real-time.
- Support the capability to “ramp-up” or change configuration of sensors based on assessment and response decisions.
- Low cost, small functional components.
- Low impact on the computing environment during normal (non-alert) conditions.
- Configuration of the Alarm System based on the site risk assessment and input by the network administrator.
- Support the use of existing intrusion detection and other analysis tools as “plug-in” sensors.
- Turn-key installation and configuration
- Simple scalable architecture.
- Self-secured
- Provide for (or not preclude) support in future versions for waste, fraud, and abuse monitoring and forensics.

2. Overview

Figure 1 illustrates the Alarm System environment and some of the important entities in that environment. The Protected Network is a local area network (LAN) protected by the Alarm System. The Alarm Administrator is a person, such as a system manager or a network administrator, responsible for the maintenance and security of the Protected Network, and configuration of the Alarm System. Barrier Defense represents other defenses, such as physical security measures or firewalls, that may be in use. The Protected Network is connected to an External Network, such as another LAN or the Internet, through the Barrier Defense. The External Network is not protected by the Alarm System, and may be an access point for Intruders. Intruders are people that attempt to illicitly access or modify systems or data in the Protected Network, including attacks on the Alarm System itself.




Figure 1. Alarm System Environment.

Figure 2 shows the logical view of the Alarm System in more detail. A sensor is a component that can recognize suspicious events and notify the central Assessment component. Assessment contains the logic to determine whether a sensor notification or a series of sensor notifications constitutes an attack, and if so, to initiate a response. Attack definitions and their associated responses are specified by the Alarm Administrator via a set of rules. The System Coordinator is a collection of software modules that manage the configuration of the system, the set of assessment rules, and the history database of sensor notifications and responses that have been initiated. The Graphical User Interface (GUI) serves four roles. First, it presents the real-time activity and status of the Alarm System. Second, it allows configuration of individual sensors. Third, it provides a means to install and securely register an Alarm component through the System Coordinator. Fourth, it provides an interface to the Assessment Rule Editor.

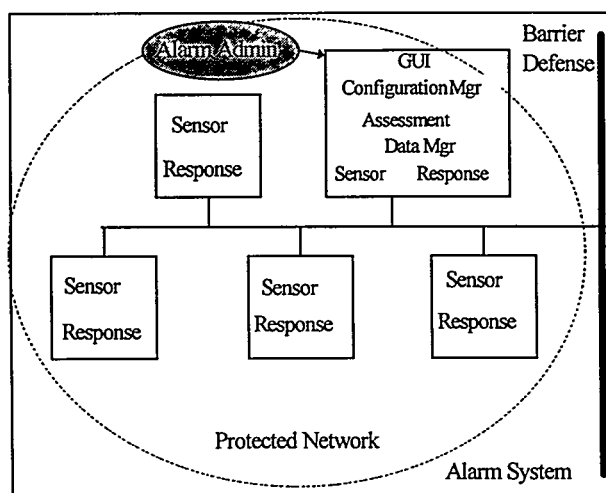


Figure 2. Logical View of Alarm System.

Figure 3 shows the general Alarm System configuration. The Command Center consists of Central Assessment, the Configuration Manager, the Data Manager and one or more local Sensors and Response components. A GUI may also reside on the Command Center machine. However, for convenience, the Alarm Administrator may specify that the GUI reside on one or more machines other than the Command Center, or that no GUI be running at all. The local Sensor(s) detect events in the Command Center. Remote Sensors detect events in other systems on the network. All Sensors report to Central Assessment. If responses are required, Central Assessment initiates responses by telling Response agents to perform some task such as informing the Alarm Administrator via pager, or isolating or ejecting the suspicious activity. The Response agents then report the success or failure of the response task. Central Assessment also provides information to the Data Manager, which manages the history database, and passes relevant information to the GUI(s). The Configuration Manager manages the Configuration Database, which keeps track of what sensors are installed on different machines on the network, the enable/disable state of those sensors, as well as other sensor parameters. The Command Center contains the most-trusted Alarm System modules, Assessment, Data Manager, and the Configuration Manager, and therefore must be well-protected.

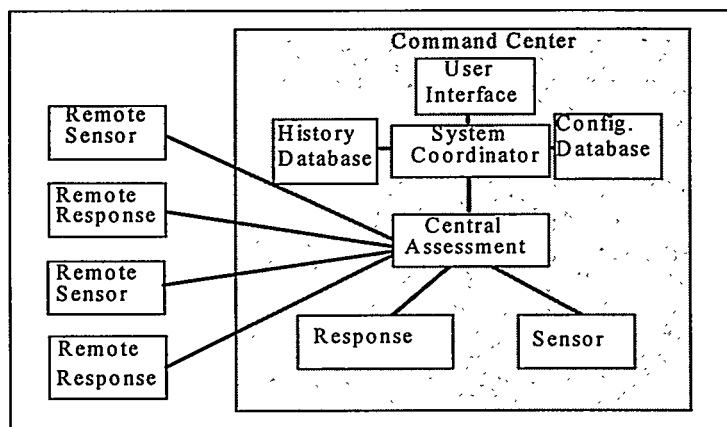


Figure 3. Alarm System Configuration.

3. Operational Concepts

3.1. Configuration

The core Alarm System consists of Assessment, the Configuration Manager, the Data Manager, the GUI, the history and configuration databases, one or more Sensors and Response agents, and the Communications Transport. The sensors are resident on the Command Center computer system and monitor host and network activity for security-related events (sense events). Assessment evaluates security events, determines the desired course of action, and initiates responses (notify others, modify security monitoring or filters, etc.)

Remote sensors are typically installed on other hosts in the network. Remote sensors connect to the core Alarm System for Command Center authentication and configuration before communicating alarms traffic. From a security policy viewpoint, the Alarm System configuration defines the desired alarm detection and response. Specifically, the strategic location of sensors on different hosts, the enable/disable state and other configuration parameters of those sensors, and the set of rules for mapping sensor events to responses, are all based on the security policy for the protected network. This information about the system configuration is stored in the configuration database, and managed through the System Coordinator. The Alarm Administrator can modify the sensors installed, their

configuration, and the assessment rules via the GUI, and the System Coordinator will update the configuration (and the configuration database) accordingly.

3.2. Hierarchical Alarm Systems

The design and implementation of the Alarm System will not preclude the use of multiple Alarm Systems in a hierarchical arrangement. A hierarchy of Alarm Systems would include a normal Alarm System operating at the lowest level in the hierarchy and the Assessment component in one level operating both as an Assessment component and as a Sensor for the next higher level Alarm System. For example, if an Alarm System was installed on a LAN the Assessment component would receive reports from Sensors and initiate responses for the LAN. The Assessment component could also act as a sensor for a higher level Alarm System by sending a sensor notification message to the higher Assessment component as part of its responses. If the Alarm System is used in a hierarchy of Alarm Systems, all conflicts in the Assessment rules must be resolved by the Alarm Administrators.

4. Alarm System Architecture

The following diagram shows the architecture for communication among different components of the Alarm System.




Figure 4. Alarm System Architecture

The diagram shows a software (SW) bus and transport, which together support secure communications among and within systems; an applications programming interface (API) to provide a common, portable interface for secure communications, and various specialized Alarm System modules. Separate API and SW bus modules are shown for each component, to emphasize that security parameters for the components are not intermixed. The lower layers (API and below) are discussed later. The left side is a client computer with a simple Alarm System sensor and a more complex third-party component integrated into the Alarm System environment by an interface translator. The right side shows a Command Center with some of the modules (Assessment, Configuration manager and Data Manager modules, and Response Components, along with the optional GUI).

Each Alarm component has a specialized role. Sensors detect and report events to Assessment, perhaps after local filtering. Assessment receives event notifications from sensors and then judges what the events mean and what

should be done. Responses carry out tasks such as notifying the Alarm Administrator or isolating or ejecting the suspicious activity. The GUI displays information about recently reported events and responses, and is used by the Alarm Administrator to query for more information, dynamically reconfigure the system, or re-load assessment rules. The Configuration Manager is responsible for configuring installed sensor agents, including both communication configuration (to whom does the agent listen and/or send information) and internal configuration (which of its capabilities does the agent perform). In addition, it updates the system configuration when directed to do so from the GUI, and maintains the system configuration and history databases.

Each sensor agent will be configured to send information to a single appropriate destination, the central Assessment agent. Hence, both information on detected events and state-of-health information will be sent to Assessment. There will be simple mechanisms to enable and disable sensor agents, as well as modify other sensor-specific configuration parameters. The Alarm Administrator will be able to request enabling and disabling sensor agents and modifying other parameters through the GUI. The Configuration Manager will send a message to Assessment which will forward it on to a sensor to modify its configuration, and then update the sensor state information in the configuration database.

When sensors report events, several kinds of feedback are possible. The response may be to (1) reconfigure the sensor, (2) enable additional sensors to look for different information, for example to "ramp-up" and collect more information if a suspicious event suggests an attack is occurring, and 3) load different assessment rules to provide more sensitive analysis of information and responses.

Trust relationships among components will probably be non-symmetrical. For example, we expect the Command Center (hence, the Configuration Manager, Data Manager, Assessment, and most Response agents) to be well-protected. We expect many sensors to reside in less-well-protected hosts, and if a host is compromised then the sensor may become untrustworthy. As a consequence, sensor components will probably be coded to trust an authenticated System Coordinator and authenticated Response Agents, but Assessment probably will not fully trust some Sensor Agents, even if authenticated.

5. Sensor Agents

Sensors are placed at critical points in the protected network to optimize detection capabilities, including in the Command Center to help protect critical Alarms components. Sensors can detect an intruder during three phases of an attack: evaluation, penetration, and information compromise. It is desirable to detect an intruder as early as possible while reducing false positives.

During evaluation, the intruder is gathering information for use in the next phase. Since information gathering is not actually an intrusion, a measured response is warranted. When a system penetration is detected, the assessment can be assured an intrusion is in progress. However, it is preferable not to allow the intrusion. The final stage for detecting an intrusion is after information or services have been compromised. Responding at such a late phase is less desirable but might enable the system to heighten protections of more sensitive machines on the network, and prevent the intruder from penetrating them.

Sensors detect activities that indicate, either alone or in combination with other sensed activities, that evaluation, intrusion, or information compromise is taking place or has occurred. In the Alarm System, many sensors are designed to be quite simple, detecting only the presence or absence of a single kind of activity. Several different sensors may have to be tripped before the pattern of behavior is considered an attack and a response is initiated.

When novel attacks occur, rapid deployment of new, unpolished sensors is likely. Consequently, other Alarm System modules will be able to handle variances in how completely a sensor implements the capabilities specified

here. In particular, a mechanism is needed to indicate which capabilities are supported and which are not. A full-featured sensor can report its configuration and status, be enabled and disabled, have its operation or configuration modified, communicate on one-way paths with specified Alarm components, have its communications paths enabled and disabled, be given updated authentication credentials, be installed, updated, and removed, and communicate with an Alarm Administrator through a Command Center.

5.1. Example Sensors

Some types of sensors that are planned for incorporation into the Alarm System are described in the following subsections.

5.1.1. Network sensors

A host attached to a LAN could contain a sensor based on a network sniffer. Such a sensor would filter and discard large volumes of uninteresting traffic. It might be configured to detect and report specific patterns or events of interest, or might be configured to report anything that did not match well-known patterns of presumably-acceptable traffic. Examples of products that could be made into Alarms sensors by adding a notification capability include:

- Gabriel, a sensor that detects port scan patterns that are indicative of a Satan attack
- Sensors that detect new hosts or hosts with incorrect addresses (ARPWATCH)

5.1.2. Network Monitoring Sensors

An SNMP-based network monitor could inquire about the status of various SNMP-aware products in a network and report if it found anything of interest.

5.1.3. Host port sensors

Various "wrappers" are available to add protection to network service ports and daemons. With the addition of Alarms communication capability, they can notify assessment of interesting events they see. Some examples include TCP_Wrappers, Overflow_wrapper, Sensors in SOCKS firewalls, SendMail wrappers

5.1.4. Host internal sensors

Specific sensors, built by adding Alarms communication to tools that look at host files will be considered for use in the Alarm System. Possibilities include integrity checkers such as Tripwire, log-file checkers that look for anomalous logged events such as bad log-ins, "zap" detector that checks if log files are being damaged, "cpm" that checks promiscuous mode behavior within a host, Disk-space watcher, and NFSwatch

5.2. Sources of sensors

5.2.1. Third-party sensors

Sensors can often be developed by adding Alarm System communications to an existing, usually security-related, product. There are at least two technical difficulties with this approach. If the product is complicated and needs proper configuration, controlling the configuration using Alarms communication may be prohibitively difficult. Especially if source code is not available, adding Alarms communication to the product may be difficult. Where a cost-benefit analysis indicates that such products are the best way to implement sensors, we will do so.

5.2.2. Custom sensors

Where an important sensor cannot be built on an existing product or the cost-benefit analysis determines that the product is not appropriate, the Alarm System project will implement the sensor either by developing new sensors or by modifying existing sensors.

6. Assessment and Response

6.1. Operational Concept for the Assessment and Response Components

The Assessment component is where policy regarding attack definitions and their associated responses are implemented. Assessment keeps track of what sensors have been triggered on what machines, and compares this list with the attack definitions in the assessment rules. If a match to a defined attack is found, Assessment automatically initiates the corresponding response(s) (specified in the rule).

The Response components are software modules that each generate one specific response action. The number and types of response actions performed will depend upon the assessment of the attack (i.e., the assessment rules).

The Alarm System design has one centralized Assessment component for the protected network. If assessment were distributed to each machine on the network, such that each assessment agent were only aware of sensor notifications on its own machine, then attacks on the network may not be recognized, because the activity on each individual machine would appear benign. In contrast, a centralized Assessment has the capability of recognizing specific attacks by recognizing specific patterns of events that are reported by multiple sensors, even on multiple machines.

One or more sensor notifications may come in from one or more systems. Assessment will, based on the assessment rules, attempt to determine what types of responses should be initiated. For example, for a certain sensor or set of sensors, response actions 1, 2, and 3 may be initiated. For other sensors, only responses 1 and 4 may be performed, and in yet another case, only response 3 will be performed.

Sensor messages are sent to Assessment from one or more Sensor components residing on various nodes. Assessment analyzes the sensor messages and sends response action messages to the necessary Response components. These Response components then perform their specified actions, such as, build and dispatch a string to an auto-dialer, build and send a command to a router, edit some particular system file, etc. Assessment also reports the response to the Data Manager, which both archives the information to the history database and reports the information to any GUIs that are currently being displayed. The Response components report back to the Assessment component on the status of their action. Assessment reports this information to the Data Manager, which can use this status to update its overall "view" of the situation, again updating the history database and GUI(s).

Assessment has the capability to be in "active" mode or in "access" mode. In "access" mode, no response actions are performed. Sensor messages are still sent to the Assessment component from various sensors, Assessment still analyzes the sensor messages, and the sensor notifications are still reported to the GUI. However, no messages are generated to the Response components. After Assessment has evaluated the current situation, it generates messages to the GUI, reporting what responses would have been invoked if Assessment had been in "active" mode. In "active" mode, response actions are initiated.

6.2. Assessment Architecture

The Assessment component is designed for easy addition of new sensors, responses, and assessment rules. Sensors can be added and/or removed from the system dynamically. As sensors are added, the Configuration Manager is

responsible for notifying Assessment of its type and location, and then Assessment is responsible for establishing communications with the sensor. When sensors are removed or disabled, Assessment is notified by the Configuration Manager.

Assessment is responsible for the assessment of single or multiple sensor notifications. Assessment is designed to receive multiple sensor messages from multiple machines. It will act as the central repository for all sensor messages that are triggered in the protected network. The evaluation of sensor notifications depends on a set of rules, either the default rules which are supplied by the Alarm System at installation, or rules reflecting local security policy which are supplied by the Alarm Administrator via the GUI. The rules conceptually are of the form "If X, then A and B" or "If X and Y and Z, then A". Here X, Y, and Z are sensor message conditions, such as the type of sensor that sent the message, number of messages received from the sensor, the source of the sensor, or the time the sensor message was received. A and B are responses to be initiated. Assessment keeps track of received sensor messages and compares the conditions to the "If" portion of the rules. If a match is found, the response or responses in the "then" portion of the rule are dispatched. The inputs to Assessment are the sensor messages (and the rules). The output of Assessment is one or more response messages. Assessment also communicates with the Data Manager to notify it of actions that have been taken. The Data Manager communicates with Assessment to notify it when sensors are being activated/deactivated and when assessment rules are altered.

6.2.1. Example Sensor/Assessment/Response Configurations

6.2.1.1. Single sensor to Multiple Response Example

It is possible that a single sensor triggered from one machine could warrant a response that should be executed on many different machines or that a single sensor could trigger multiple responses. An example of this situation would be if a user account has been compromised on one machine, an appropriate response may be to disable the compromised account on all machines that contain the account. Note that this is just a special case of the general architecture.

6.2.1.2. Multiple sensor to Single Response Example

Another situation could be that many individual sensor notifications are received from different hosts. This could imply that an attack such as Satan is being run against all or part of the protected network. This type of attack could result in a single response such as denying service at the router from the domain that is detected as the attacker's source. Note that this is just a special case of the general architecture.

The distributed nature of this architecture is the concept which makes it unique in capability. Without the centralized Assessment component, attacks against multiple systems, that individually may look benign, may not be detectable. On the other hand, individual attacks may provide sufficient information that preventative measures can be taken to protect multiple systems before the attacker can even get there. For centralized Assessment to work in these various cases, all notifications must go to the same place.

6.2.2. Response Architecture

The Response subsystem is a "plug-and-play" compatible component that can easily adapt to changes in the Sensor component, Assessment, and communication transport mechanism. A Response component is responsible for performing the desired action, or directing another process to perform the desired action, if the response needs to be run on a remote machine. An example of this might be to edit the password file on a remote machine to disable an account. The specific rule that has been satisfied will determine how the response will be structured. There is no

inherent limitation to the number of message conditions to be satisfied in a rule, nor the number of responses that can be initiated as a result of the same set of conditions.

Assessment determines which response or set of responses are to be initiated in the event of a possible attack or violation. Once the response(s) is determined, Assessment communicates with the respective Response component(s) to initiate the action. The Response component is accountable for performing the desired action or for directing another process to perform the desired action as designated by Assessment. Upon successful completion of the response action, a message is sent to Assessment indicating that the action was successful; otherwise, a different message is sent denoting failure of the action. Assessment will then forward the status of the response to the Data Manager, which archives the information to the history database and reports the message to the GUI.

The only Alarm System component that a Response component communicates with outside of the host is Assessment. A Response component will always receive commands from Assessment, and upon completion of the response action, the Response component will return a message to Assessment; the only exceptions to this being when Assessment is in "access" mode in which case the communication between Assessment and the Response component will be disabled. The purpose of the "access" mode is for verification that the Alarm System is operating properly while preserving the state of the network and alarm system.

The Response components may require some "low-level" configuration specific to each host, e.g. identifying sensitive files that require signatures. The Response components will be designed so that they are easy to install by the user and will be functionally the same across UNIX System V and Windows NT systems. The capability to incorporate a third-party Response component or a component created by the user may be considered for future versions of the Alarm System.

6.3. Example Responses

The following is a list of response alternatives that could be taken based on sensor notifications that have been received by Assessment. This list is intended only as a possible set of response options. The Response components for the Alarm System will not be limited to the functions listed, nor should there be any expectation that all functions listed will be implemented in the current development cycle. The set of responses implemented will be selected based on an evaluation of the completeness of functionality, robustness, and flexibility of the set.

6.3.1. Notification Responses

The minimal (and least aggressive) response for any intrusion detected is to notify a person or persons via one or more of the following *Console message, Pager message, E-mail message to a designated address, Fax message, Start-up the GUI, Generate an SNMP alarm, Notify other Alarm Systems, and Send a message to the GUI.*

6.3.2. Active Responses

The more aggressive responses that stop, isolate, or eject the intrusive activity are labeled as Active Responses. Possible active responses are *Close connection, Disable user account, Close application, Deny a Service, Honey-Pot Machine* (dynamically switch intruder to another system or a system attractive to intruders).

6.3.3. Passive Responses

Passive Responses are responses that do not stop, isolate, or eject the detected intrusion. The responses might collect information about the attack, turn on additional protection mechanisms if needed, or allow human intervention before another response is initiated. Possible passive responses are *Traceroute, Turn on auditing or*

other sensors, Escalate Monitoring, Check signatures on Sensitive Files, Change file permissions, filename, or file location, Time-based response (Allow specified amount of time to pass to allow human intervention before initiating another response), Email message to the system administrator of attacker's network, Request approval (via pop-up dialog box) before initiating a response, Load new assessment rules to change sensitivity of evaluation.

7. Transport

Basic transport for the Alarms system includes communications and security. The transport will be simple to implement, understand, use, and install; require little user maintenance; support for "reliable" and "unreliable" transport; prevent easy eavesdropping, and decoding of the communications; prevent attack of the Alarm System by hijacking or playback, and provide for asynchronous communication via callback mechanisms.

In the Internet, "unreliable" services tend to be more usable for some purposes when a network is behaving poorly (such as when most data is lost or corrupted due to errors, intermittent failures, or attacks). "Reliable" services tend to use complicated optimization and recovery methods, and introduce their own failure modes and points of attack.

The ability to use "unreliable" communications may be crucial for the Alarms system to function in a degraded mode under denial of service attacks, such as, broadcast storms, port and processor saturation. "Reliable" mechanisms like TCP may fail when a network is under attack. Initial development of the Alarm System will allow for homogeneous deployment of reliable (TCP), or unreliable services (UDP) between network nodes. The use of TCP or UDP will be set when the Alarm System is installed on a network.

The initial implementation of the Alarm System will use a simple, readily available secure communications mechanism. Overall system architecture and design "principles" shield the upper levels of the software from needing to know any of the secure transport details. In particular, the Software Bus API provides a standard interface to the authors of upper level components (Sensor, Assessment, System Coordinator, and Response agents). Should the secure transport mechanism change, then only the Software BUS API, and lower level transport mechanism need be modified. Interprocess communications used in Alarms will use shared memory for local host communications and an IPC-like mechanism for communication between hosts. The Alarm System will use the Secure Socket Layer protocol (SSL) for secure transport (encryption) and authentication. The message transport will also include sequence numbers and time stamps in each message.

8. Graphical User Interface (GUI)

The Graphical User Interface (GUI) enables an Alarm Administrator to communicate with an Alarm System. Using the GUI, the Alarm Administrator will be able to install, update, and remove Alarm Components, enable and disable Alarm Components, modify the operation of Alarm Components (e.g. with local configuration files), specify which Alarm Components communicate with one another, modify and dynamically reload assessment rules, look at certain logs and auditing information, ascertain the configuration of Alarm Components in the network, ascertain the status of Alarm Components in the network, ascertain correct operation of Alarm Components in the network, ascertain whether interesting events occur in the network, get notifications of specified events in near-real-time, relate the above information to the network configuration easily, and communicate sufficiently securely with the System Coordinator while doing any of the above actions.

The GUI interacts directly with the Data Manager, who can perform many of these tasks (ascertain configuration; enable and disable; modify operation; specify communication; modify rules; and install, update, and remove) because it communicates with the Configuration Manager. Other tasks (ascertain events; get notifications; look at logs) are supported because the Data Manager maintains and accesses information in the history database.

The Alarm System GUI will be developed to be independent of any specific operating system or computer system. The GUI is implemented in Java using a viewer. Alarm System information will be collected and formatted, by the Data Manager. The GUI may run on any node accessible from the protected network.

9. Configuration Manager

The Configuration Manager is the collection of software modules that are aware of all the components of the Alarm System, coordinate supervisory issues such as startup and shutdown, , manage the Alarm System configuration, and is responsible for maintaining knowledge of the Alarm System's configuration in the configuration database. Specifically, the Configuration Manager interacts with the Alarm administrator via the GUI to provide configuration information and to receive administrator requests; interacts with Assessment to provide required rule changes and required changes in sensors, be informed when Assessment causes or learns of changes in sensor behavior, and change parameters and behavior of response components; interacts with remote computer systems to control sensor behavior; installs and removes Sensor and Response components; and maintains a configuration database that contains the desired configuration information and, the actual configuration.

10. State-of-Health Monitoring

The Alarm System will insure itself against the compromise of integrity and availability of the Alarm System components themselves, such as replacing a component that does not provide its Alarm System functionality or an intruder may also attempt to intercept messages and/or modify messages from one Alarm component to another.

The Alarm System will have several complementary protections designed in to guard against such threats of compromise of the System itself. First, sensors will be installed to monitor for changes in the binary codes of the different Alarm components. If a component is modified or replaced, the sensor will detect it, and an appropriate response can be taken, such as re-installing the correct component.

Another mechanism to protect against compromise of the Alarm System will be state-of-health monitoring ("heartbeat") of the Alarm components. State-of-health monitoring is a means of checking that the components are "alive and well". State-of-health monitoring of sensors and response agents will be accomplished from Assessment. Assessment will send a message to each Sensor and Response component. If the component is "alive", it is expected to reply to the query within a designated time interval. If no reply is received, a message is sent to the GUI, or other responses could be dispatched, as defined in the Assessment rules.