A LAYERED FRAMEWORK FOR PLACEMENT OF DISTRIBUTED INTRUSION DETECTION DEVICES

Marcelo Medina
George Washington University
Washington, D.C.
mmedina@seas.gwu.edu

ABSTRACT - Network based distributed intrusion detection is a common trend in several commercial intrusion detection systems. However, network based intrusion detection requires that a security officer comprehends the dynamic and non-deterministic nature of data traffic across the network. This paper provides security officers with a brief introduction to intrusion detection techniques and classifications. The paper, then, proposes a framework for placement of distributed intrusion detection devices along its four layers: the network perimeter, the high sensitivity network components, the location of data and applications, and traffic analysis. The following sections will discuss intrusion detection and the proposed layered approach in greater detail.

INTRODUCTION

The growth of the Internet and corporate intranets has led to an era of connectivity. Most companies are connecting their computer networks to their partners' networks. The expansion of the Internet in the commercial market has forced many companies to offer Internet access to its employees. Along with this connectivity comes the concern of unauthorized use of these computer networks. This paper calls unauthorized users intruders. These intruders are classified as external, internal and misfeasors [5]. Intrusion detection systems have been developed to catch these intruders.

Intrusion detection can be characterized as host based, multi-host based and network based. The host and multi-host based implementations assume that a security officer has control of the computer systems themselves, which is not always the case. For example, the same group may not manage the business hosts and the telecommunications infrastructure of a company, therefore the security officer for the telecommunications group may not require modifications on the computer systems. Network based solutions are now commercially available and attracting the attention of corporate security officers. This paper focuses on network based security.

Recent research efforts such as EMERALD [3] and GrIDS [4] have focused on intrusion detection across a computer network. The efforts include the research of techniques on how to coordinate data generated by intrusion detection devices (IDDs) placed across a network. This research overlaps with the research done in network management, which also needs to handle multiple devices across a network, gathering and interpreting vasts amount of data flowing through a network.

This paper covers intrusion detection and several issues of interest to coordinate the cooperation of the IDDs across a network. The paper focuses on the placement of the IDDs, and presents a framework for placement of these devices.

Classification of Intruders

The media has generated much hype on network intruders. This paper does not comment on the tools or techniques such intruders use. The focus is on the coordination of activities to catch these intruders. This paper follows the intruder classification developed by Jim Anderson [5]. Anderson classified intruders as:

- External – users not authorized to use the systems
- Internal – users not authorized to use some resources
    a) Masquerades – impersonate other user
    b) Clandestine – evade auditing
- Misfeasors – who misuse their privileges

Teresa Lunt described extensively techniques against masquerades [2]. Clandestine intruders are a potential threat for weak security operating systems (OSes) and badly managed systems. The threat of clandestine intruders should be treated by the OS with better control of audit processes and more secure OS implementations. External intruders are the focus for physical security, firewalls and other techniques.

This paper treats all users as possible threats, regardless of where they are originated or how they are authenticated. This maximizes the coverage of system.

Types of Intrusion Detection

Intrusion detection consists of several techniques to trace unauthorized use of resources. These techniques are based on the study of audit trails and network traffic; Teresa Lunt [2] characterized the study of intrusion detection in three types:

- Real-time testing of audit data
- In depth off-line (after-the-fact) analysis of audit data
- Subsequent analysis of the audit data for damage assessment

This paper focuses on the real-time testing of audit data. The intent is to catch an intruder in the act, so that an immediate response can be taken. The in-depth and subsequent analysis should still be performed since they provide insight to new methods and techniques of real-time analysis.

The body of this paper consists of four sections. The first section provides an introduction to intrusion detection and its techniques. The second describes the benefits and difficulties of distributed intrusion detection systems. The third presents the common locations for intrusion detection devices. The fourth describes the framework for placement of intrusion detection devices.

INTRUSION DETECTION SYSTEMS

Intrusion detection systems monitor computer and network traffic for signs of unauthorized use. The system generates alarms (e.g., console messages, e-mail messages and pages) when it detects possible unauthorized activities. The techniques used on this detection are signature analysis and statistical profiling. Intrusion detection systems can be host based, multi-host based and network based. This section describes this classification and the techniques.

COAST's Classification of Intrusion Detection Systems

COAST [8] maintains a listing of developed intrusion detection systems. It classifies intrusion detection systems based on the origin of the data:

- Host based – audit data from a single host
- Multi-host based – audit data from multiple hosts
- Network based – network traffic data along with audit data from multiple hosts

Historically, the systems evolved from host based, through multi-host based, into network based. A few commercially available systems, such as NetRanger[1] and RealSecure[2], are strictly network based without reviewing audit logs from hosts.

Limiting Intrusion Detection Devices to the Network Level

Most network security officers have little or no control over the applications running on their network. The functional division of institutions precludes the network security group from tracking every application developed. Also, controlled environments (e.g., military facilities) might choose to have some applications classified above the network security officer. *Therefore, the network security officer can monitor only the network traffic for information on intruders.*

Intrusion Detection Techniques

Intrusion detection systems offer a variety of functions and different implementations. This paper introduces two common techniques for intrusion detection. The paper, however, does not dwell on how they are implemented. The implementation of such systems is a complex computer science problem [1, 2, 3, and 4] and is not covered on this paper.

Intrusion detection can be defined as activities using two techniques:

1. Signature Analysis
2. Statistical Profiling

Intrusion detection systems may use one or both of the techniques above. The following sections describe these in more detail and explore the benefits of distributed intrusion detection.

Signature Analysis

Signature analysis matches network traffic against known rules containing known attack traces and protocol uses. Network traffic is matched against these rules. Traffic that matches known attacks are flagged, for example, the Mitnick/Shimomura attack. Protocol traffic that does not match the protocol rules is flagged as a potential attack. For example, a normal ftp session performs one command at a time, so if a put is received while a get command is being executed it should be flagged. The algorithms have to be extremely efficient to track the sessions and activity for all network traffic. The most serious implementation issues involve maintaining a database of common attacks and protocol

[1] NetRanger is a trademark of WheelGroup Co.
[2] RealSecure is a trademark of ISS, Inc.

usage, and efficient algorithms to match traffic against the rules.

Statistical Profiling

Statistical profiling is common for host based systems, but can also be performed by network based systems. It consists of monitoring certain characteristics of user usage. For example: application, amount of data, time of usage, protocols used, source and destination address, etc. Profiles are generated for each user and subsequent uses are checked against the profile. For example, a user whose profile indicates the use of MS Word[3] only should be flagged if he edits a remote host password file using *vi*. This can also be performed at the network. For example, web browsers pool web servers for data, the server receives requests in specific http protocol format to download pages to the client browser and specific format form data so that scripts can generate a page. If the web server starts receiving remote commands and file uploads, those are not normal processes and should be flagged.

DISTRIBUTED INTRUSION DETECTION SYSTEMS

Benefits of Distributed Intrusion Detection

Originally real-time network level intrusion detection devices were stand-alone devices. They gathered, processed and took actions alone. However, the deployment of these devices on large networks has led to new developments on these products. They now are expected to cooperate, sharing information (traffic data and alarms). This has led to research and development of distributed intrusion detection systems [3, 4]. Distributing the intrusion detection devices across the network allows the security officer to have a broader view of the network. He can therefore:

- Identify intruders that scan the network (sweep) and not just a segment. For example, intruders attempting *"Doorknob rattling"* [4], a sweep that checks for vulnerable hosts.
- Correlate attack signatures among different segments on the network (e.g., sweep attacks from multiple sources).
- Coordinate counter actions by following the physical route taken by the packets to track the attack sessions (i.e., trace the user even if he provides a false source IP address based on the links he traversed). Therefore he can support counter action through the collection and correlation of event data.
- Reduce cost by sharing IDD resources.

Difficulties on Distributed Intrusion Detection Systems

Coordinating the interaction between these devices has led to a number of difficulties. Several of these difficulties were already known by work developed on network management [6]. Network management, however, tends to be centralized and to have easily defined states (i.e., link is up or down, devices answers pool or not, device configuration matches defined state or not). Intrusion detection takes this to a dynamic level in which instead of comparing snapshots the system looks at windows of time and sequences of activity.

The SRI research on EMERALD[4] [3] shows some of the issues that need to be considered. Those are: event generation and storage; state-space management and rule complexity; knowledge repositories; and inference activities.

Event generation and storage consider where are the events generated and stored, distributed or centralized. Previous intrusion detection systems have concentrated on centralized generation, storage and processing of events. This model does not scale well for large networks. The large number of events and devices distributed across the network can generate too much network traffic and too much data to be stored in one location efficiently. It also does not cover distributed services (e.g., DNS, firewalls).

State-space management and rule complexity deal with a tradeoff on how complex rules should be and the demands they impose on processing the traffic through the network. Complex rule models are more effective[5] but require complex state management and analysis algorithms. The CPU requirements to process large audit logs with complex rule sets will likely limit the real-time application of complex rules. Less complex rules process faster, which facilitates the processing on high-speed links.

Knowledge repositories are locations containing the rules. The issue is that rules need to be updated and distributed to the IDDs. Having one central point of

[3] MS Word is a trademark of Microsoft Co.

[4] EMERALD (Event Monitoring Enabling Responses to Anomalous Live Disturbances) is a distributed scalable tool suite for tracking malicious activity through and across large networks.

[5] Effectiveness is based on the accuracy of the decision model.

distribution minimizes management concerns. This model does not scale well, since rule updates need to be distributed across the network in an orderly and timely fashion.

Inference activities are the processing of event logs [3]. Where should the inference activities be performed? The centralized location model does not scale when there is a large number of IDDs. As discussed on event generation and storage, the traffic generated by logs crossing the network is too much overhead for this system. Also, having one central processing unit will require a lot of CPU power. A distributed environment will require algorithms that control the distributed processing. A centralized environment generates traffic. There is a tradeoff here.

> "Centralized analysis severely limits the scalability of the detection algorithms. In internetworks of multiple administrative domains, different domains may be unwilling to share all activity information with others. Also sufficient processing and communications resources to analyze activity in very large internetworks is unlikely to be available." [4]

Other Difficulties on Distributed Intrusion Detection

The SRI research has taken a computer science view on how to implement intrusion detection devices. Many of its considerations are well-studied computer science problems. There are, however, other practical issues that need to be considered by the network designer, the concern being on how this new service will impact the existing network infrastructure. Here is a list of four of those issues:

- Communication between intrusion detection devices across the network can generate traffic on the network (e.g., SNMP traffic). Meyer, et all [6] made a point that "expanding intrusion detection to a distributed system is likely to result in network congestion if all audit data must be sent to a central location."
- The sharing of logs across the network creates a security risk. The attacker may be able to compromise the flow of data across the network or gain knowledge of systems security by monitoring the log traffic.
- Intrusion detection devices need to be able to have a standard event description so that they can share information. An object model for events is necessary. Events should not be distributed as audit logs across the network. The distributed sensors should process the logs and share, in a distributed or centralized fashion, these events as objects.
- The modeling of the location of the intrusion detection devices is network dependent. What works in one company's network, may not work in another's. Network implementations vary widely from company to company. Some have tiered architectures, others are functionally divided and others, such as recent mergers, are a connection of pre-existing infrastructures. Determining traffic patterns is a non-trivial exercise. Modeling of network traffic patterns needs to be developed for optimum placement of the IDDs.

IP Networking

This section introduces security officers to IP networking and routing principles. The intention is to clarify the dynamic nature of the IP networking environment and its non-deterministic nature. These are fundamental to the placement of IDDs.

Switched environments (such as ISDN and X.25) create paths across the network and maintain those during the session. The session data therefore follows the same physical connections as long as the session is up. IP based networks, however, are a packet forwarding type of network. Every node on the network forwards the data packet (a piece of a message) to the next "best" node. The way it identifies the *best* node is by looking at its routing table. The routing table includes the following information: the destination, the next node to send the data to and a measurement of the distance to destination through that node (e.g., number of hops, cost of the link). These routing tables are generated by routing protocols.

Routing protocols define how network nodes communicate (i.e., frequency, and format of data) and how routing tables are calculated based on the received information. There are many routing protocols and each has specific properties. Some protocols measure distance on number of hops (e.g., RIP), others provide finer granularity using weights and costs (e.g., OSPF). The one common characteristic between them is that changes in the network status will trigger the recalculation of the routing tables. The routing tables are updated with the new paths of smallest costs (i.e., number of hops, accumulated weight). Therefore, the best path for one packet may not be the best path for another packet, even if they are part of the same session. Therefore, *there is not an obligatory path on an IP network.*

The basic principle of network design is that there shall be no single point of failure. From every source to destination there shall be at least two paths. There is no single link that all traffic must traverse. IP and its routing protocols guarantee this. If there were only one link, it would be a single point of failure - which is a non-acceptable network design. Even gateways to external networks are planned for redundancy. Therefore, for external and internal attacks, there is not single link that all traffic covers. To gather data on all traffic, IDDs must be distributed across the network.

COMMON IDD LOCATIONS

There are many locations to place intrusion detection devices across the network. This section considers three possible locations: the network perimeter, the server farms, and the backbone. Most implementations are a variation of placing devices in one or more of these locations.

The Network Perimeter

The definition of the security perimeter is a common assignment for security officers. A common perimeter demarcation is all that is internal to a network against all that is external. For example, the internal components are all the equipment that belongs to a company inside a building and the external components are the shared services, equipment owned by a third party and what is hosted on a somebody else's premisses.

Figure 1 presents the network perimeter. Usual perimeter equipment is:

- Firewalls – connect the internal network to an external network (i.e., the Internet or another company's network). Firewalls may also be used inside the network to define perimeters inside perimeters (e.g., classified network would be "firewalled" from unclassified network even though both belong to the Department of Defense).
- Access servers and modems – support dial-up users into the network. These mark a point of entry into the network.
- Commercial links – depending on the sensitivity of the data the demarcation between proprietary wiring into commercial services may also be a perimeter demarcation.

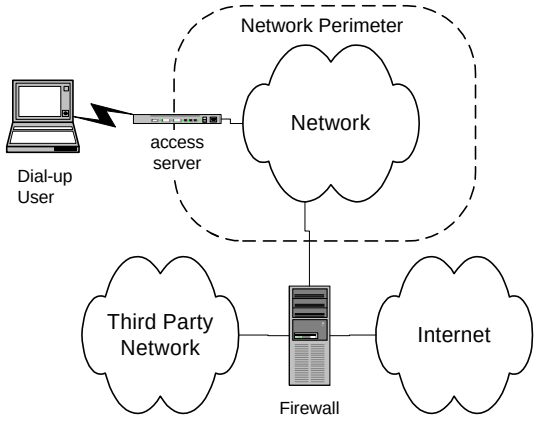

Figure 1. The Network Perimeter

The Server Farms

Server farms are network segments that host servers. Those are controlled segments where no client workstation is present, only servers. The server farms are likely targets of attacks. They tend to have valuable data and provide the highest risk for the company. They are characterized by having enormous amounts of traffic. Depending on the nature of the services this traffic may have specific characteristics, e.g., web servers, or miscellaneous traffic, e.g., file servers.

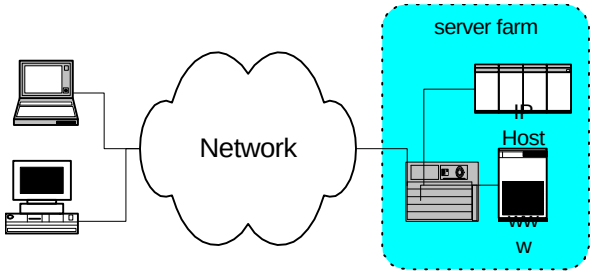

Figure 2. The Server Farms

Server farms are commonly implemented on large networks. However, one should not consider that they are the sole storage location. A lot of data tends to be maintained at local workstations and servers. Many of the new workstations and laptops come with large hard drives, allowing the user to maintain his data locally. Many users will only place data on server if this data needs to be shared.

Peer-to-peer networking makes all stations possible servers. Users that would place data on servers for sharing will now allow other users to attach to their own workstations. The local segments should be monitored in an environment where most data is kept on the local

workstation. Further analysis should be performed for intrusion detection in peer-to-peer environments.

The Backbone

The backbone is an infrastructure that provides access in between different network areas. Those areas can be geographical and/or functional. Backbones vary from low bandwidth to high bandwidth links depending on how systems are implemented across the network. Business applications may avoid backbone links because of possible delays. The traffic may be a great variety of protocols but of little variance on the processes since they tend to be written into the applications. Intruders will likely probe the network for important systems. Anomalous network traffic on the backbone may flag intruders on the network. For example, port scanning across the backbone and IP spoofing attempts are occurrences that should trigger the security officer to investigate possible intruders.

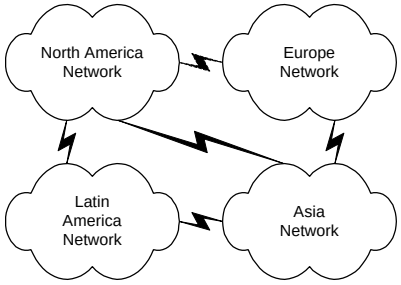

Figure 3. The Backbone

Figure 3 shows the backbone as lighting bolts connecting regional networks. All traffic leaving one region to another region will cross the backbone.

A LAYERED FRAMEWORK FOR PLACEMENT OF DISTRIBUTED INTRUSION DETECTION DEVICES

The placement of intrusion detection devices on a distributed IP network should be the result of a combination of several factors:

- the effectiveness of the intrusion detection system,
- the amount of network overhead generated, and
- the cost.

The effectiveness of the system is the only item covered in this section. The amount of network overhead generated will depend on the implementation. As for the cost, no specific budget constraints are considered. This section, however, assumes limited resources in order for us to minimize the number of IDDs. The contradiction of this assumption would be to have unlimited resources, which is rarely the case. Also, assuming unlimited resources would allow the placement of IDDs on every network segment. The management burden of such a solution would surpass its effectiveness. Therefore the goal is to limit the number of IDDs across the network while maximizes its effectiveness.

The proposed framework considers four grouping of segments and network equipment when placing intrusion detection devices. The framework considers each group as an increment to the previous group. Therefore, the framework is presented as a nesting of security implementation layers. Figure 4 presents the framework as a nesting diagram with the most interior layer being the least secure and security growing as the other layers are also covered. Those layers are:

- **The network perimeter.** The network perimeter is the demarcation point between internal and external. The demarcation point can be firewalls to the Internet, firewalls to third party networks, access servers for dial-up users, modems, the physical location of the network devices, ownership of the network devices, etc.
- **The high sensitivity network components**. These are the network components that are critical to the functioning of the business. For example, application servers for banks, network infrastructure for telecommunications companies, the data servers for an accounting firm.
- **The locations of the data and applications**. The data and applications may be stored on the desktop, on local servers, on server farms or on a combination of these.
- **Traffic analysis.** It should consider most utilized links, how much inter-region traffic exists against intra-region traffic, and type of data. Traffic analysis is where modeling needs to be run continuously so that the existing deployment does not get dated.

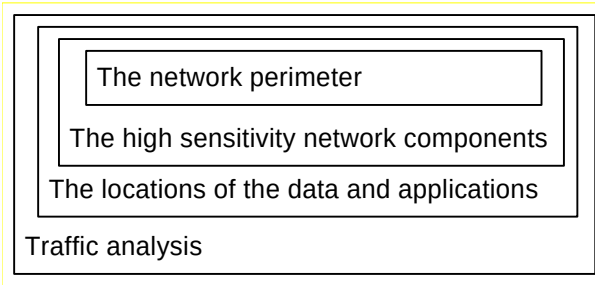

Figure 4. The Nested Framework

The Network Perimeter

If nowhere else, the network perimeter should be guarded against attacks. Regardless of the security statistics presenting the risk of internal users as higher than external attacks, on the growing connectivity environment of today one needs to protect itself against the external threat. The proliferation of test tools such as Satan [7] made breaking into systems a teenager pastime. Not tracking misuse by external users may place the IT infrastructure in the hands of potential intruders.

Facilities with more specific security concerns, such as banks and government organizations need also concern themselves with professional intruders such as industrial spies, professional hackers, etc. Those will likely be able to get around commons security implementations such as proxies and packet filters. They may also illegally acquire entry codes into the dial-in pool (e.g., stealing laptops, dynamic password generator cards, and personal papers or digital organizers).

Generic usage services such as access servers should have signature analysis monitors on their segment. This would be able to flag initial probes and attacks. Internet specific traffic that tends to be limited in nature should run signature analysis and statistical profiling. Statistical profiles will enable the security manager to track any new activity on web and mail servers, which the security manager can clear with the applications support group. As for modems, those should be prohibited by the security policy unless the connected equipment performs strong authentication. There are a variety of network based modem services that can provide dial-out capabilities and dial-in should only be performed at the access servers.

The High Sensitivity Network Components

Once the perimeter is protected, it is time to consider the high sensitivity network components. These will vary from network infrastructure for network service providers to applications and data servers for most companies. This layer assumes that network perimeter is protected. Therefore, the intruder is on the network and the security officer intends to protect the assets that are most sensitive. The assumption is that once the intruder is on the network he can use any break-in tool, such as Satan and ISS' System Security Scanner. Signature analysis should be used and the rules used should be the most complex and broad ones.

Depending on the sensitivity of these network components they may have their own perimeter which is different than the network perimeter. If they do not have a defined perimeter, one should be generated and IDDs should be placed on it. For these systems both signature analysis and statistical profiling should be performed.

It is on this layer that real-time intrusion detection intends to recover its costs. It is also on this layer that the necessity for support personnel for intrusion detection becomes noticeable. There must be an Emergency Response Team for these systems. Contingency and countermeasure plans must be in place. Intrusion detection systems may catch the attack, may even shun certain attacks. But a determined attacker will require manual intervention to protect the livelihood of the service and the company.

The Locations of the Data and Applications

The third layer of this model is the locations of data and applications. This layer assumes that the network perimeter and the high sensitivity network components are protected. These may not be classified as critical to company. They are, however, where strategical data tends to be maintained and where important infrastructure systems, such as HR databases, are commonly maintained.

Intruders will likely probe the local servers before expanding over the network. They tend to contain information about the location of important data and files like *password.xls* containing passwords for the other many systems the account holder commonly logs in. Misuse also tends to begin on the local server, where the disgruntled employee starts testing his access rights. Signature analysis should be run because intruders will likely try to exploit known security flaws on the local server. Host and multi-host based intrusion detection should also be considered for the local hosts.

Peer-to-peer environments are difficult to control. Security may benefit most by educating the users as to the security risks of maintaining important data on the local workstations. For example, Windows NT[6] has been deployed in some corporations as the panacea to workstation security, later it was learned that NT has security flaws that allow remote users to gain remote access to the workstation. Deployment of intrusion detection in a peer-to-peer environment is difficult to implement, unless the scope is small.

Server farms, however, tend to concentrate large amounts of data. The network segments where server farms are located should run signature analysis. Statistical profiling of the large number of users and the variety of services would make the CPU burden of statistical profiling unfeasible.

---

[6] Windows NT is a trademark of Microsoft Co.

Traffic analysis

Traffic analysis is the last, but most important layer for distributed intrusion detection. Traffic analysis studies the amount and type of traffic that rides on a network. It logs time, source and destination addresses, protocol used (well known and proprietary) and the amount of data. GrIDS [4] provides a possible implementation based on graphs that monitor network activity. Traffic analysis can define the areas of most traffic on the network. It can also be used to generate profiles of the network traffic for statistical analysis intrusion detection. Traffic analysis will validate the placement of the IDDs. If the devices are not capturing most of the traffic, it is likely that intruders will be missed. Signature analysis should also be performed on the areas of most traffic. Intruders are likely to try to disguise their activities by performing them in times of heavy traffic. Given enough processing power, detailed statistical analysis should be performed to help build profiles. Traffic analysis provides the insight to the network which is necessary to place the IDDs for successful counteraction. Distributed devices will need to correlate data; they will also need to coordinate the monitoring of flagged sessions (since the network traffic may be routed through an alternative path).

CONCLUSION

This paper presented a layered framework for placement of distributed intrusion detection devices on IP based networks. The model has four layers, each layer being a superset of the previous layer. The layers are: the network perimeter, the high sensitivity network components, the locations of data and applications, and traffic analysis. Network based distributed intrusion detection is the direction of several commercial intrusion detection systems.

The effectiveness of distributed intrusion detection systems depends on how much of the data traffic is captured. The amount of data captured depends on the network traffic on a distributed environment. Therefore, the relationship between distributed intrusion detection systems and traffic analysis needs to be explored. This should improve the effectiveness of distributed intrusion detection systems.

The peer-to-peer environment needs to be controlled, either through security tools, intrusion detection on the desktop or through policy. Not doing so may hinder the effectiveness of intrusion detection systems.

REFERENCES

[1] T.F. Lunt, Automated Audit Trail Analysis and Intrusion detection: A Survey. *In Proceedings of the 11th National Computer Security Conference, October 1988.*

[2] T.F. Lunt, Detecting Intruders in Computer Systems. *In Proceedings of the 19th National Information Systems Security Conference, October 1988.*

[3] P. A. Porras and P. G. Neuman, EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances. *In Proceedings of the 20th National Information Systems Security Conference, 1997.*

[4] S. Cheung, et all, GrIDS - a graph based intrusion detection system for large networks. In *Proceedings of the 19th National Information Systems Security Conference, 1996.*

[5] J.P. Anderson. Computer Security Threat Monitoring and Surveillance. *Technical Report, James P. Anderson Company, Fort Washington, Pennsylvania, April 1980..*

[6] K. Meyer, M. Erlinger, J. Betser, C. Sunshine, G. Goldszmidt, and Y. Yemini. Decentralized control and intelligence in network management. In *Proceedings of the Fourth International Symposium of Integrated Network Management (IFIP/IEEE), Santa Barbara, CA, May 1995, pages 4-16. Chapman & Hall, London, England, 1995*

[7] SATAN - Security Administrator's Tool for Analyzing Networks [On-line].
Available HTTP http://www.fish.com/satan/

[8] Purdue's COAST Project. "Intrusion Detection Systems." Computer Operations, Audit, and Security Technology intrusion detection page [On-line].
Available HTTP http://www.cs.purdue.edu
File: /coast/intrusion-detection/ids.html