

LA-UR 90-3726

CONF-910596--1
MAY 1991

DEC 13 1990

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

LA-UR--90-3726

DE91 004823

TITLE: NADIR-A PROTOTYPE NETWORK INTRUSION DETECTION SYSTEM

AUTHOR(S): KATHLEEN A. JACKSON
DAVID H. DuBOIS
CATHY A. STALLINGS

SUBMITTED TO: 1991 IEEE SYMPOSIUM ON RESEARCH IN SECURITY AND PRIVACY

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

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

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

FORM NO 836 R4
ST NO 2629 5/81

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

pe
MASTER

NADIR - A Prototype Network Intrusion Detection System

Kathleen A. Jackson, David H. DuBois, Cathy A. Stallings

Computer Network Engineering Group
Computing and Communications Division
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Abstract - The Network Anomaly Detection and Intrusion Reporter (NADIR) is an expert system which is intended to provide real-time security auditing for intrusion and misuse detection at Los Alamos National Laboratory's Integrated Computing Network (ICN). It is based on three basic assumptions: 1) that statistical analysis of computer system and user activities may be used to characterize normal system and user behavior, and that given the resulting statistical profiles, behavior which deviates beyond certain bounds can be detected, 2) that expert system techniques can be applied to security auditing and intrusion detection, and 3) that successful intrusion detection may take place while monitoring a limited set of network activities such as user authentication and access control, file movement and storage, and job scheduling. NADIR has been developed to employ these basic concepts while monitoring the audited activities of more than 8,000 ICN users.

The Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36. This work was performed under auspices of the United States Department of Energy.

1 Introduction

In a large, complex, and rapidly changing computer network it is not realistic to expect that all security loopholes will be identified, and if identified, can be closed. In the case of some known vulnerabilities, it is impractical to do so. A primary reason for this is that a balance must be struck between security and the requirement that we provide reasonably convenient services to network users, who are in fact our reason for being.

The authentication and access control system in any network is the initial defense against intruders from outside. Authentication is the identification of a user with reasonable assurance that the user is who he or she claims to be. This is usually accomplished by means of a user ID and password combination. Access control is a mechanism of restricting access by authenticated users to those portions of the network consistent with their clearance and need-to-know. Given the industry-wide frequency of break-ins by outsiders, it is unfortunately obvious that authentication and access control mechanisms can be compromised or bypassed and that they alone cannot be completely relied upon to ensure that no penetration by outsiders occurs. In addition, even the most secure systems are vulnerable to abuse by insiders who misuse or attempt to misuse their privileges. This is obvious from the number of well publicized reports in the last few years of incidences of unauthorized access and removal of classified information by insiders from otherwise secure computer systems.

An auxiliary line of defense against both intrusions by outsiders and insider misuse is the maintenance and review of an audit record of significant network activity. In the absence of an automated system, security personnel must attempt to review huge quantities of printed output in an often futile attempt to spot invalid activity. The sheer volume of data makes it nearly impossible to detect suspicious activity that does not conform to a few obvious intrusion or misuse scenarios, and even these may be missed. What is needed is the capability for automated security analysis of the audit record; a capability which combines the knowledge of security experts with a computer's capability to process and correlate large quantities of data. When this analysis is done in near real-time, security personnel may be notified of suspicious activity in a timely manner, and direct action taken to trace and stop an identified penetration attempt or other misuse.

2 Target System

The Integrated Computing Network (ICN) is Los Alamos National Laboratory's main computer network. It consists of host computers (which execute user programs), file storage devices, network services, local and remote terminals, data communication interfaces, and distributed processors (DPs). DPs are remote processors which range from workstations (personal and mini computers) to full-scale computers. The "core" of the ICN is considered to be the main host computers and their support devices, while the DPs are considered to be the "extended" network. Through the ICN, any user inside the Laboratory may access any host computer (if the user has authorization to do so and the access path is approved) from office workstations or terminals. Outside users typically access the ICN through telephone modems, leased lines, or one of multiple world-wide networks. The core ICN has more than 8,000 validated users.

The ICN consists of four "partitions"; the Open, Administrative, National Security, and Secure partitions. ICN Partitions are dedicated to specific levels of processing, and are limited to users cleared for the most sensitive information processed in a given partition. A host computer may be in only one ICN partition. The Open partition is available to anyone who has a legitimate need to compute at Los Alamos and is limited to unclassified, non-sensitive computing. The Administrative partition is primarily dedicated to processing sensitive unclassified data or data subject to a privacy act. The National Security partition is dedicated to processing of DOD classified and unclassified data. The Secure partition is dedicated to processing DOE classified and unclassified data. Partitioning is enforced throughout the network by a system of dedicated ICN nodes. These nodes are physically protected, have tightly restricted access, are limited to only that software needed to perform a specific function, and do not execute user programs. Only these dedicated nodes are allowed to service multiple ICN partitions. Each of these nodes is required to maintain a complete audit record of its activity.

Use of the ICN is controlled by a security system which authenticates users and grants access only to those authenticated users who possess a validated password, who are located in an area under security control consistent with the ICN partition to which access is sought, and who also have a clearance level consistent with that partition. This implementation of user authentication is brought about through the use of a Network Security Controller (NSC), which is a dedicated, single-function com-

puter through which all ICN user authentications must pass. In compliance with DOE orders, an audit record is maintained of all attempts to gain access, or "logon" to the ICN. Each day of the weekday NSC audit record normally contains twelve to fourteen thousand entries, while each day of the weekend audit record contains six to seven thousand entries, each representing an access attempt.

In order to keep project parameters at a manageable level, and because of its critical network user authentication function, it was decided that the NSC was a logical choice as the initial target system for the automated analysis of the audit record for anomaly and intrusion detection. The system described henceforth in this paper is the first phase of the ICN intrusion detection system, as it has been applied to NSC audit data.

3 History

Until recently, most security auditing of NSC activity was performed by manually scrutinizing system logs and thus identifying potential security violations. Given the magnitude of the NSC audit records, manual review of these records was limited to a small sampling or a very cursory scanning. It was for this reason that development of a computer program to analyze these records was initially undertaken at Los Alamos in 1983-84 [1]. On a non-real-time basis, this program sought to detect attempts to penetrate the NSC and to gain unauthorized access to the computers and files of the ICN. It checked the audit record for a small set of entries, or combination of entries, which were anomalous enough to raise concern. At the time of this development, no body of knowledge on the nature of attempts to penetrate a computer security system similar to that at Los Alamos existed, so the rules defining significant events in the NSC were few, and were conjectural in origin, rather than empirical. Even so, the results of this program development were encouraging.

Further development at Los Alamos of an automatic audit record analysis program was influenced by the initial research of Dorothy Denning and her colleagues [2, 4, 5], and the current IDES research and development being carried out by Teresa Lunt and her colleagues at SRI International [6, 7, 10, 11, 14]. They have demonstrated that 1) the statistical analysis of computer system activities may be used to characterize "normal" system and user behavior and, given such statistical profiles, that user and system activity that deviates beyond certain bounds is detectable, and 2) known intrusion scenarios, exploitation of known system vulnerabilities, and violations of a system's security policy are detectable

through use of an expert system rule base. Their approach puts a primary emphasis on the detection of deviations from normal user and system behavior by statistical means, combined with an expert system which encodes intrusion scenarios which are intended to catch invalid activity which may be missed by the first means. Another approach has been demonstrated by the development of the Multics Intrusion Detection and Alerting System (MIDAS), which was implemented on the National Security Center's Dockmaster system [7, 8]. Although heavily influenced by the work at SRI, the major emphasis on MIDAS was to encode a set of *a priori* rules that define invalid activity and intrusion scenarios. This approach has also been successfully applied to security audit log analysis in an expert system (AudES) developed at IBM [13].

In late 1988, an intrusion detection feasibility study was undertaken at Los Alamos. Its purpose was to look into the possibility of developing an intrusion detection system for the Los Alamos network. To restrict the scope of the problem to a manageable level, this study was limited at this time to an analysis of NSC audit data. Automated tools which provided for both on-line and off-line analysis of the NSC audit record were developed. These tools were not only used for the feasibility study; they were also put into production for use by ICN security personnel. It was determined that an expert system approach, using a set of pre-determined rules, to the problem of ICN audit record analysis would work, that invalid user activity could be detected, and in fact such a system would be relatively easy to implement [8]. In the spring of 1989, with the receipt of funding from the Operational Security Division at Los Alamos, the Network Anomaly Detection and Intrusion Reporter (NADIR) project was initiated.

The major goals for the development of the NADIR system were to:

- Develop a better understanding of the patterns and range of user activity on the ICN, for future planning and development.
- Develop a means by which to detect attempted or successful security violations and unanticipated security vulnerabilities.
- Provide a more efficient method of user authentication audit record review, which was required by ICN security personnel.
- Develop a near real-time method by which to detect a range of security relevant events, including attempted break-ins to the ICN by outsiders and invalid activity or abuses by insiders.
- Expand the means to provide near real-time detection of attacks, automated or otherwise, for the ICN.

In addition to the stated project goals, providing useful tools to network and security personnel during each phase of development was given a high priority.

4 Working Prototype

NADIR is implemented on a dedicated SUN SPARCstation¹ with two 327 MByte disks. It uses the Sybase² relational database management system and a Los Alamos designed expert system, which is programmed almost entirely in Transact-SQL³, an enhanced version of the SQL database language provided by Sybase. Transact-SQL provides such capabilities as stored procedures, triggers, system administrator tools, and control flow language features which were used extensively in NADIR. The only other programming language used is C, which was needed for a portion of the user interface. Sybase provides tools which are used to structure, maintain, and display all data on the system. NADIR communicates with the target system (the NSC) over a dedicated secure ethernet link, and receives the NSC audit record in the order, and at the time, that each user authentication occurs. From this record it constructs and maintains statistical individual user and composite (of all user activity on the system) profiles. After each profile update, NADIR performs an immediate expert system intrusion detection analysis on the updated individual user profile, and on the composite system profile. Immediate reports of suspicious activity are output to the system terminal, and the information saved. A user interface allows a choice of built-in queries or allows ad-hoc queries against both the raw audit data and the individual user and composite profiles. It allows the review of all the audit data associated with a particular user, a particular machine, or any other parameter over any selected period of time. When requested, background analysis may be performed and various reports generated. A complete audit record, starting in October 1989, and continuing to date, is formatted for NADIR and is readily available. Older audit data is archived and available, but would require processing to be usable by NADIR.

NADIR was designed to be easily integrated with different types of target systems. In general, all

¹ SUN SPARCstation and SUN workstation are trademarks of SUN Microsystems, Inc.

² Sybase is a trademark of Sybase Corporation.

³ Transact-SQL is a trademark of Sybase Corporation.

NADIR target systems are, and will be, required to install the capability to collect the appropriate audit record of user activity, put the data in a specified fixed format, and transmit it to NADIR. It is expected that audit records from different target systems will vary in format and contain, in most respects, unique data. This is because each of the ICN target systems performs functionally different tasks; i.e. user authentication and access control, file movement and partition changes, file storage, and batch job scheduling, etc. To support the addition of new target systems to NADIR, the software was designed in a modular fashion, so that upgrades can be handled with a minimum of effort. Communication with NADIR requires the installation of Sybase provided interface software (DB-Library⁴, which is available for many languages), and the use of a standard DECnet⁵ or TCP/IP protocol.

NADIR currently monitors NSC user authentication and access control activity on the ICN. The NSC is a DEC-8250⁶ machine, which runs the VMS⁷ operating system. The changes required to the NSC system were minimal. An implementation of TCP/IP under VMS was provided by the Multinet⁸ software package. Interfaces to Sybase were provided by DB-Library packages for Fortran and C. The system code was changed only to format the audit record for NADIR, and to provide for the transmission of a record of each user authentication immediately after its occurrence. The NADIR required data processing on the NSC has not resulted in any measurable degradation in system performance.

NADIR uses 100 MBytes of disk space for the database generated from the NSC audit record. From the time it is received, it is able to process an NSC audit record and report any suspicious behavior found within .25 seconds. Since the NSC performs a user authentication every 3.6 seconds during peak times, and NADIR can handle 4 every second, NADIR has demonstrated that it can easily handle the load from this one target system. (As additional target systems are added to NADIR, we envision a network of SUN workstations, each processing the audit record from one or more target

⁴ DB-Library is a trademark of Sybase Corporation.

⁵ DECnet is a trademark of Digital Equipment Corporation.

⁶ DEC-8250 is a trademark of Digital Equipment Corporation.

⁷ VMS is a trademark of Digital Equipment Corporation.

⁸ Multinet is a trademark of TGV, Inc.

systems and each contributing to a distributed database.)

5 System Design

The design approach taken with NADIR was to duplicate on the system those audit record review activities which had previously been undertaken by security personnel. We wished to replace the manual review of audit logs with an automated system. As a result of this, the system was designed to duplicate, and hopefully improve on, the auditor's activities by the application of expert rules.

As a first step, all user and system activities are summarized into profiles. These profiles are maintained for each individual ICN user, and for a composite of all ICN users. The second step in the process is the application of expert rules to the profiles. Next, anomaly reports are generated of all invalid or suspicious behavior. These reports are then sent to security personnel for review. The Los Alamos security group investigates each anomaly and provides feedback to NADIR designers as to the results of their investigations. Finally, where indicated, the expert rules on NADIR are modified to improve the discrimination and judgement of the system.

5.1 Target System Audit Record

The NADIR monitors target system activity as it occurs and is recorded in audit records generated by the target system. Each Audit Record describes a single event. The data in an Audit Record reflects the type of activity on the target system, and thus may vary in format and content. Whatever the system, the Audit Record will contain a unique ID for the ICN user, the date and time of the user's activity, fields which describe the activity, and any errors which may have occurred.

Each NSC audit record describes one attempted ICN authentication, both successes and failures. Each record, as sent to NADIR, consists of the following data:

- **Event Timestamp** - Date and time at which the authentication attempt occurred.
- **User Number** - The unique ID of the user requesting authentication.
- **Logon Level** - The user requested computing level (classification) for the authentication attempt. The level at which the user wishes to compute.
- **Logon Partition** - The user requested partition for the authentication attempt. The partition in which the user wishes to compute.

- **Source Partition** - The partition from which the authentication attempt originated.
- **Source Machine** - The network address of the machine from which the authentication attempt originated.
- **Destination Machine** - The destination of the authentication attempt, when known. In some cases, an authentication is to a particular partition, not a specific machine.
- **Authentication Status** - The status of the authentication. This may be successful, or one of 29 different errors.
- **ACP Address** - The ICN address of the access control point from which the attempted authentication originated.
- **ACP Port** - The port on the access control point from which the attempted authentication originated.
- **Charge Code** - An accounting parameter.

5.2 User Profiles

NADIR maintains profiles for both individual ICN users and for a composite of all ICN users. A profile is a description of current user authentication behavior, with respect to a set of defined parameters. The profiles are updated as each record of activity occurring on the target system is received by NADIR. As users alter their behavior, their profiles will change. Rather than the detail contained in each audit record resulting from a user's activity, the profiles contain count statistics which summarize the activity. This was done to pre-process the data and reduce data storage to a manageable level. When a new audit record is received, the data is parsed and the appropriate counts in the profiles are incremented. At this point in development, new profiles are generated for each week. Past weekly profiles are maintained for comparison purposes and as a permanent record.

5.2.1 Individual User Profiles

Each Individual User Profile is initialized when the user becomes valid on the ICN. From that point on, the profile provides a complete record of the user's authentication activity. The profile contains a the User Definition, a User History, and a record of User Activity.

User Definition - This is the basic definition for each ICN user. It is initialized when a user is first entered as a valid ICN user. After that the user number is never modified, and the other information only as circumstances require.

- **User Number** - The unique ID of the valid ICN user.
- **User Name** - The user's name.

- **User Type** - One of various types of ICN users, some of which have special privileges.
- **User Group** - The group or organization for which the user works.
- **Mail Stop** - The user's Los Alamos mailing address.
- **Citizenship** - A parameter indicating the user's citizenship.

User History - This provides a history of individual user authentication activity. From this may be determined what activities are normal for the user. This data is updated with each attempted user authentication.

- **Sources** - The different sources from which the user has logged on to the ICN.
- **Destinations** - The different destinations requested by the user when logging on to the ICN.
- **Charge Codes** - The different Charge Codes used while logging on to the ICN.
- **Blacklist History** - The number of times and dates upon which a user has been blacklisted⁹ by the NSC.

User Activity - This provides a description of individual user authentication activity. These are count statistics which are updated with each attempted user authentication.

- **Successful Source Open** - Successful logons for sources in the Open partition.
- **Failed Source Open** - Unsuccessful logons for sources in the Open partition.
- **Successful Source Administrative** - Successful logons for sources in the Administrative partition.
- **Failed Source Administrative** - Unsuccessful logons for sources in the Administrative partition.
- **Successful Source NS** - Successful logons for sources in the National Security partition.
- **Failed Source NS** - Unsuccessful logons for sources in the National Security partition.
- **Successful Source Secure** - Successful logons for sources in the Secure partition.
- **Failed Source Secure** - Unsuccessful logons for sources in the Secure partition.
- **Successful Destination Open** - Successful logons for destinations in the Open partition.
- **Failed Destination Open** - Unsuccessful logons for destinations in the Open partition.

- **Successful Destination Administrative** - Successful logons for destinations in the Administrative partition.
- **Failed Destination Administration** - Unsuccessful logons for destinations in the Administrative partition.
- **Successful Destination NS** - Successful logons for destinations in the National Security partition.
- **Failed Destination NS** - Unsuccessful logons for destinations in the National Security partition.
- **Successful Destination Secure** - Successful logons for destinations in the Secure partition.
- **Failed Destination Secure** - Unsuccessful logons for destinations in the Secure partition.
- **Successful Unclassified** - Successful logons at the Unclassified computing level.
- **Failed Unclassified** - Unsuccessful logons at the Unclassified computing level.
- **Successful Confidential** - Successful logons at the Confidential computing level.
- **Failed Confidential** - Unsuccessful logons at the Confidential computing level.
- **Successful PARD** - Successful logons at the PARD (Protect As Restricted Data) computing level.
- **Failed PARD** - Unsuccessful logons at the PARD computing level.
- **Successful Secret** - Successful logons at the Secret computing level.
- **Failed Secret** - Unsuccessful logons at the Secret computing level.
- **Successful Day** - Successful logons during the day shift.
- **Failed Day** - Unsuccessful logons during the day shift.
- **Successful Swing** - Successful logons during the swing shift.
- **Failed Swing** - Unsuccessful logons during the swing shift.
- **Successful Night** - Successful logons during the night shift.
- **Failed Night** - Unsuccessful logons during the night shift.
- **Successful Weekday** - Successful logons on a weekday.
- **Failed Weekday** - Unsuccessful logons on a weekday.
- **Successful Weekend** - Successful logons on a weekend.
- **Failed Weekend** - Unsuccessful logons on a weekend.

⁹ Blacklisting is applied to an individual user with the occurrence of five sequential authentication failures. A blacklisted person is denied access to the ICN by the NSC. Removal of the blacklist must be approved by security personnel.

5.2.2 Composite User Profiles

The Composite User Profile maintains user authentication data both for the ICN as a whole, and for each separate Access Control Point¹⁰ (ACP). It maintains the following data for each hour of the day, each day of the week, for the whole ICN and for each ACP:

- **Valid Logon** - The total number of successful logon attempts to the ICN.
- **Invalid Logon** - The total number of unsuccessful logon attempts to the ICN.
- **Invalid Header** - The number of logons rejected because the network header was invalid.
- **Improper Format** - The number of logons rejected because the logon message format was improper.
- **Invalid Field** - The number of logons rejected because a logon message field content was invalid.
- **Unknown User** - The number of logons rejected because the user number was not that of a valid ICN user.
- **Blacklisted Users** - The number of users blacklisted by the NSC.
- **Previous Blacklisted Users** - The number of logon attempts by users who have already been blacklisted, either by the NSC or by security personnel.
- **Invalid Operator** - The number of logons rejected because of ICN operator logon requests by users who were not operators (an operator is a user with special privileges).
- **Invalid Password** - The number of logons rejected because an invalid password was entered by a known ICN user.
- **Invalid Password Level** - The number of logons rejected because an invalid computing level was entered for known password.
- **Invalid Level** - The number of logons rejected because a requested computing level was invalid for either the source terminal or the requested host.
- **Improper Location** - The number of logons rejected because a requested computing level was not allowed from the source of the logon. For example, using a Secret password at an Open terminal.
- **Invalid Access** - The number of logons rejected because the requested computing level, source partition, or destination partition was not allowed by the user's assigned access (individual users have ICN access limitations). For example, a user

¹⁰ An Access Control Point is a dedicated ICN node which authenticates users through the NSC and provides access to the ICN for authenticated users.

with Open access attempting to use a Secure terminal.

- **Invalid Partition** - The number of logons rejected because the logon partition is not consistent with the source partition, or the host partition is not valid for the partition of the terminal. For example, to logon to the Secure partition from an Open terminal.

5.3 Anomaly Data

A record of each anomaly is generated when a user's behavior causes the activation of one or more expert rules. It consists of two tables, as follows:

- **Anomaly Record** - A table which contains the user number of each valid ICN user, and an indicator for each simple expert rule. The indicator is turned on as each rule is fired by the user's behavior, thus enabling the quick identification of patterns of rule violations. This table provides a history of each user's anomalous behavior.
- **Anomaly Trail** - For each anomalous user, this table contains the time at which the user performed the activity which was deemed anomalous, and the anomalous activity. It provides a record of every new rule violation by each anomalous user.

5.4 NADIR Processes

NADIR consists of six top-level processes which interact with each other through the Sybase relational database management system.

- **Access Control** - A feature of the Sybase system is the ability to administer and control databases independent of any specific database application. The system administrator maintains database integrity through use of username, password, and access privileges. For any user (and this includes target systems) to select tables, insert data, modify data, delete data, execute procedures, or modify privileges, that user must have been given the privilege to do so. The Access Control module enforces all these functions on the NADIR system.

- **Define User** - This process provides for the initial definition of each valid ICN user. When a new user is added to the NSC database, a message is transmitted to NADIR which contains the initial User Definition. Occasionally, parts of the User Definition are modified on the NSC (such as the user's type, group, or mail stop). When this happens, this process provides an update to the NADIR User Definition.

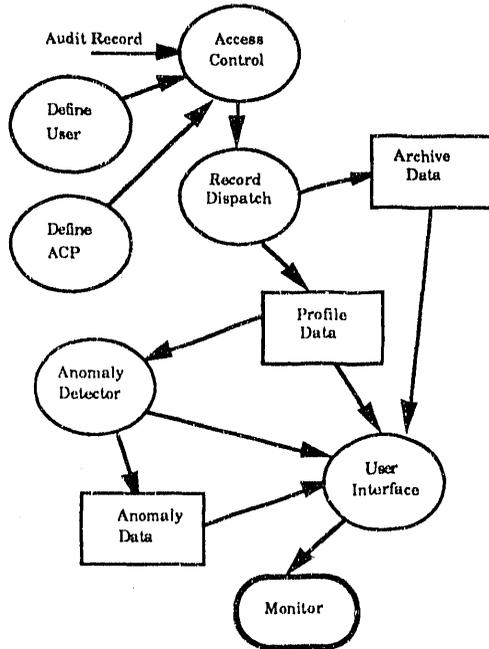


Figure 1: NADIR Functional Flow

- **Define Access Control Point** - This process provides for the definition of each valid ACP. As the ICN changes, these nodes are added and removed from the network, and changes are made to definition tables on the NSC. When this happens, this process keeps NADIR network definitions current with those on the NSC.

- **Audit Record Dispatch** - This process takes the input Audit Record from the target system, parses it, checks for a valid ICN user number, and calls procedures to update the appropriate Individual User Profile and the Composite User Profile. The activities of unknown users (not valid ICN users) are maintained in a separate table (Unknown User), and are also used to update the Composite User Profile. This process also places the input Audit Record into a table where it is maintained locally for two weeks.

- **Anomaly Detector** - This process checks the appropriate Individual User Profile (the one which has just been updated) and the Composite User Profile for any anomalous or suspicious activity by applying the expert rule base to them. If anomalous activity is found, the Anomaly Record and Anomaly Trail are updated, and if required, a message is sent to the system terminal.

- **User Interface** - This process uses Sybase front end tools, graphics packages, and Los Alamos designed routines (using C and DB-Li-

brary) to provide a preliminary interface for the knowledgeable user. For users who have been provided the appropriate access and privilege (only system developers and selected security personnel), the user interface allows a choice of built-in queries or allows ad-hoc queries against both the raw audit data and the individual user and composite profiles. It allows the review all the audit data associated with a particular user, a particular machine, or any other parameter over any selected period of time. Data may be displayed in a variety of ways, including graphically, and reports generated.

5.5 Archiving

NADIR generated user profiles and anomaly reports are regularly backed up from NADIR to the Secure partition of the ICN Common File System (CFS). A Sybase utility dumps the entire NADIR database into operating system files. The files are then archived and placed on the CFS, where they are maintained indefinitely. When required, up-to-date profiles and reports may be moved any time. This stored data may be accessed whenever needed by security personnel and the ICN CSSO for review on NADIR. The profiles are time tagged for easy access of the desired period of time to be reviewed. In addition, the raw NSC audit record is also backed up to the CFS, with updates made every few minutes.

6 Expert System Development

Probably the most difficult part of starting a knowledge engineering task, especially for those who have never done it before, is gathering and organizing the pertinent expert data. There is also the difficulty of the "AI Mystique", where terms like "artificial intelligence" give one the impression there is something magical involved in building an expert system. We found that an organized, step-by-step, well documented methodology [3] was essential to the development of a successful system. Also, rather than try to define a complete rule base on the first pass, we employed extensive testing of our developing rule base against real audit data.

6.1 Knowledge Acquisition

The first step in knowledge acquisition for NADIR was to determine which user authentication activities could be used to characterize normal user behavior, and what activities indicated possible or certain invalid activities. In other words, we needed to know just what an intrusion or invalid activity might look like in terms of NSC audit data, and how to differentiate between valid and invalid use. The

second step of the process was to define a set of expert rules for application against the user authentication audit record. Finally, these expert rules had to be tested against the audit record, and a determination made of their effectiveness.

6.1.1 Audit Record Analysis

We first needed to expand our knowledge of current ICN user authentication behavior from NSC audit records. Over a period of time, many evaluations of user activity were reviewed by the NADIR designers and ICN security personnel. In some cases, evaluations were a result of suspicious activity on the ICN which had been detected by other means, and the evaluations were used as part of the analysis of the event. In addition, the audit record was analyzed statistically, to determine normal ICN activity.

6.1.2 Interviews

In addition to audit record analysis, a great deal of knowledge was obtained by the standard method of interviewing ICN security personnel. These included those whose responsibility is to establish and enforce the Laboratory's security policy, as well as those whose responsibility was the manual review of the NSC audit record. Because the Laboratory's security policy is well defined, the first portion of knowledge acquisition turned out to be quite straightforward, and a set of expert rules was quickly defined. The auditors, however, mostly relied on extensive knowledge of the ICN, knowledge of past invalid activity, and gut instinct to identify any but the most obvious scenarios of suspicious activity. Only a portion of this expertise was documented in a way that was useful in the definition of expert rules. It was with these individuals that the interview process was of most value, and the development of an initial set of rules was completed within several weeks.

6.1.3 User Profile Analysis

A statistical profile of an average ICN user was developed to enable NADIR to detect when an individual user deviates from the overall average behavior of the group of all users on the ICN. Groups of different types of users were profiled to enable us to detect when an individual user in a particular group differs from normal group behavior. A composite of all user activity on the ICN was developed to enable NADIR to detect network-wide deviations in behavior.

Statistical analysis of both individual user and composite profiles was undertaken in an attempt to obtain a more complete and empirical definition of

what is normal overall ICN user authentication behavior. To accomplish this goal, standard summary statistics were applied to the profiles. Summary statistics provide concise descriptions of variables, and permit simple comparisons of one variable to other variables or to external standards. The results were combined to provide a statistical profile of a normal ICN user and of normal composite user behavior. Then a search was made of the user profiles for users who deviated significantly from the normal profile. Review of these user events enabled us to determine that this kind of deviation, particularly if combined with other indications, frequently comprised a significant event.

Dynamic graphical data analysis, which permits active analyst intervention, was also applied to the user profiles. It was ideal for spotting extraordinary usage profiles within a large data structure such as the ICN user profiles. Such statistical exploratory notions as rotating, brushed, and sliced plotting capabilities were used. As a development tool, it was used to help locate and identify anomalous ICN user activity, and with it we hoped discover unanticipated patterns and relationships. While this approach did not identify any new user profile relationships, it both confirmed what was discovered in the earlier user profile analysis, and provided an alternate visual method for identifying anomalous user activity.

6.1.4 Rule Definition

The process of knowledge acquisition led to the definition of about 100 rules which applied ICN security policy, defined normal individual and composite user behavior, and identified invalid and suspicious events. These rules were used repeatedly against user profiles generated from the NSC audit record, with an intriguing result. With analysis of the user profiles, the audit record, and the anomalies detected by our expert rules, we recognized violation conditions which had been missed, and vulnerabilities which had not previously been identified. A number of new event types, either suspicious or clearly invalid, were identified. In addition manual auditing had limited security personnel to very simple rules. Since a computer has no such limitations, much more elaborate rules than those which had been previously performed by security auditors were successfully implemented.

Throughout this process, our knowledge of both the ICN and of user authentication activity on it grew enormously. Not only was an expert rule base developed for the NADIR system, but cases of insider misuse were identified and corrected and modifica-

tions were made to the way some activities were performed on the network.

6.2 Expert Rules

NADIR rules try to detect attempted break-ins by outsiders, masqueraders, and misuse by insiders. To detect attempted break-ins by outsiders, NADIR uses rules involving such things as logon failures by known users (especially password failures), black-listing of known users, logon failures by unknown users, large numbers of failures from a single source, failures from dial-up lines, and a very high rate and/or precise timing of attempted logons (automation). To detect masqueraders, NADIR uses rules involving unusual or abnormal user logon parameters (time, location, partition, computing level, etc.), especially when logon failures are combined with these parameters, and such things as simultaneous (or nearly) logons from physically separate locations. To detect misuse by insiders, NADIR uses rules involving attempted access to classified or sensitive partitions, suspicious movement of files between partitions, automated logons, large rates of logons, and misuse of restricted (special usage) user numbers. At the time of writing, 211 expert rules are defined, of which over 100 are implemented on the working prototype. NADIR rules fall into four basic categories:

Immediate – These rules are generally the implementation of ICN security policy, and were obtained by interviewing security personnel and reviewing documentation. They are intended to detect individual events which are potential or certain security violations, or which because of the activity type, are inherently interesting and must be included in periodic reports. A simple example of a security violation would be:

IF an "Improper Location" error has occurred,
AND the terminal used is in the Open Partition,
AND the password used is classified,
THEN Report a security violation.
EXPLANATION: Use of a classified password from an unprotected terminal is considered reason enough to consider the password compromised. The password will be immediately invalidated.

Individual Anomaly – These rules are applied to individual user profiles, to detect when a user's behavior departs from that which has been determined to be normal and valid ICN user behavior. These rules were obtained by means of a statistical analysis of the past behavior of all individual ICN users, and by interviewing security personnel. An example of an individual anomaly would be:

IF the "Failure Ratio" of a user is $> .5$,
AND the user has logged on >50 and ≤ 100 times,
THEN Log the event, and assign it a weight.
EXPLANATION: If a user has logged on to the ICN enough not to be a new user, and the average ICN user has an average Failure Ratio¹¹ of .04, then a Failure Ratio of .5 is considered significant. A sliding scale of urgency, balanced between the total number of logons and the Failure Ratio, is applied to this rule, with the numbers above as one example.

Composite Anomaly – These rules are applied to composite user profiles, to detect when the composite of all user activity departs from the pattern which has been determined to be normal and valid for the system. These rules were obtained by means of a statistical analysis of the past behavior of the composite of ICN users. An example of a composite anomaly would be:

IF the number of "Unknown User" errors is > 40 /hour, OR > 120 /day, OR > 480 /week,
THEN Log the event, and output an urgent message to the system terminal.
EXPLANATION: The normal number of attempted authentications which contain a user number which is not valid for the ICN is statistically very consistent (1-5 an hour during peak activity, less at night and on the weekends, is normal). Extreme variations from this expected activity could be an indication of a break-in attempt. A sliding scale of urgency is applied to this rule, with the numbers above as one example.

Attack Scenario – These rules define one event, or a sequence of events, which have a low probability of occurring, and which indicate a known or postulated attack. Attack scenarios were obtained from security personnel and other experts in system penetration. Individual and composite user authentication profiles are tested for evidence of attacks (automated or otherwise) on the NSC which could result the compromise of passwords, denial of service, or "swamping" of the system. The rules which make up an attack scenario are individual rules which have been already defined in the previous three rule types. It is the sequence and combination of these rules that make for an increasing certainty that an attack may be under way. Attack scenarios are in the definition stage for NADIR.

¹¹ Ratio = $\frac{\text{Invalid_Logons}}{\text{Successful_Logons} + \text{Invalid_Logons}}$

6.3 Expert Application

The expert system portion of NADIR acts on the user profiles, whenever there is a modification to a profile. That is, it is evoked after each audit record is received from the target system. NADIR examines the user profiles and determines whether any profile is anomalous with respect to the expert rules. NADIR flags as anomalous behavior which has triggered one or more of its rules. The anomalies are weighted, based on the number, type, and combination of the triggered rules. The greater the weighting, the more suspicious the event. The triggering of one rule might not be enough to raise an alarm, but if combined with other rules it may indicate anything from an interesting event to a critical intrusion attempt. When an anomaly is detected, an Anomaly Record and Anomaly Trail are recorded in the data base, and are used in the generation of a regular report.

7 Report Generation and Follow-Up

7.1 Routine Reports

Currently, NADIR generates routine reports on a weekly basis. Summary hardcopies are routed to project developers and security personnel. A complete report, which includes data from the audit record to support the findings of the summary, is stored in the Secure partition the ICN's Common File System, where it may be accessed and reviewed electronically by authorized personnel. The weekly hardcopy report is 18-20 (two sided) pages in length, and contains:

- Summary statistics of all activity on the ICN for the week (one page).
- Graphical representation of all user activity for the week, including anomalous activity, plotted over time with a granularity of one hour (eighteen plots).
- A list of all anomalous users for the week (usually 65-85 users), listed in order of suspicion level. Of the total, 7-10 will be very suspicious, 20 or so moderately suspicious, and the rest various levels of interesting. For each user, the list contains a weight (level of suspicion), a user number, and the user's name, group, and type.
- A detailed description of each user's anomalous activity, including which rule(s) were triggered.
- A list of all users who moved files from a higher to a lower level ICN partition (a high-risk activity which is closely monitored), sorted on the basis of the classification of their computing activity.
- Two pages of descriptive boilerplate.

7.2 Critical Reports

If a critical event is detected, security personnel are contacted as quickly as possible. An appropriate short report is generated, the contents of which depend on the nature of the event. Detailed follow-up reports may be requested as part of an investigation.

7.3 Follow-Up

Upon receipt of a NADIR report, whether it be critical or routine, security personnel perform a review of all anomalous activity, even the relatively uninteresting. In order to process the weekly reports in a timely manner, specific security personnel are assigned responsibility for various categories or types of ICN users. Each anomalous user's activity is reviewed in detail, and a decision made whether further investigation is required. This may include interviewing the user. If the user's activity warrants it, the user is blacklisted during the investigation. A short report is filed at the completion of each investigation, giving details of its resolution. This information is provided to the NADIR developers, so they may have immediate feedback on system performance. Periodic reviews are held with security personnel to evaluate the system's effectiveness and to make recommendations for improvements.

8 Background Analysis

Security personnel at Los Alamos frequently have the need to perform background reviews of user activity on the ICN, based on information received from a variety of sources, and for many different reasons. These reviews usually involve one individual ICN user, but have at times involved such things as all users from a particular source. To support this need, NADIR provides the capability for background analysis of current and past activity for a particular user or users, or any other parameter in the database, over any specified period of time. The audit data required for background analysis is maintained indefinitely at Los Alamos.

9 Results

NSC audit data has been continuously processed for invalid activity since November of 1989 (for part of this time in weekly batch mode), using a growing and improving expert system. The NADIR working prototype has been in operation since June of 1990. Reports have been generated on a weekly basis for this entire time period, and statistics of ICN activity maintained. Rather than try to validate the system by use of artificially constructed test cases and in-

trusion scenarios, we used the audit data normally generated by the target system, and a process of extensive evaluation of the results. Rather than try to guess what real event scenarios would look like, we waited to see what the system would find for us. As a result, we:

- Identified invalid activity by unknown (presumably external) users.
- Identified numerous cases of misuse or suspicious behavior by insiders, including automated logons, misuse of special use user numbers, apparent (unsuccessful) attempts to logon using another person's user number, attempted logons (unsuccessful) from terminals in partitions to which the user did not have access, and attempted use (unsuccessful) of computers in partitions to which the user did not have access.
- Uncovered unanticipated network problems which had not previously been identified, which have been remedied where possible or are being closely monitored.
- Identified misuse conditions which had not previously been identified. These resulted in the definition of new rules.
- Provided support in the background analysis that was required during investigations of a number of current and past ICN users.
- Were able to define more complex rules by using the sequence and type of simple rules which were triggered by certain kinds of events. For example, a straightforward automated logon event (where a program performs a rapid, continuous, evenly spaced series of logons) always triggered the same eight simple rules. From this we were able to define a rule which would identify this kind of event specifically, on the basis of quantitative evidence.

In addition to benefits in the area of anomaly detection, NADIR has provided unanticipated benefits. It has enabled us to:

- Detect problems with some nodes of our network as they occurred. For example, a surge of invalid network messages from an Access Control Point could be the first indication of a hardware or software failure, rather than a user induced problem. We were able to tell the difference between the two types of activity and encode it into our rule base.

- Provide detailed reports upon request of network activity that was useful to personnel such areas as accounting and networking, which included statistics of network and computer usage.

We have found it difficult to come up with a number that accurately describes our "false positive rate". It's true that most of the flagged individuals and events are not intruders, spys, or even users deliberately misusing the system. It's also true that their behavior, for one reason or another looked suspicious, and for our security personnel, that's reason enough for at least a preliminary investigation. We believe that as long as the list of flagged users and events is short enough for quick review, it is better to have "false positives" than to miss anything significant.

10 Future Directions

Anomaly notification currently consists of terminal messages and periodic reports. For serious security events, the ultimate goal is to have notification on a near real-time basis. This notification will be broadcast to the Los Alamos Network Operations Center (NOC), which is manned 12 hours a day, with personnel who are reachable 24 hour a day.

Future targets will be other network nodes which control file access, storage, and movement, and operations control such as job scheduling. We plan to develop a network of SUN workstations, each processing the audit record of one or more nodes, distributing the functional applications and database, and thus optimizing performance. Each node to be added to the system has or will undergo a development process similar to that of the NSC. Currently, two nodes are in the process of being added to NADIR. The user activity record from the ICN Security Assurance Machine (SAM), which enables (and restricts) the movement of files between ICN partitions, is currently being added to the user profiles on NADIR. The Facility for Operator Control and User Statistics (FOCUS), which provides operations control, batch job scheduling, and accounting control, is undergoing analysis in preparation for adding its data to NADIR. As new nodes are added to NADIR, their user activity record is being correlated with previously included nodes to produce more complete profiles of each ICN user. This will eventually allow us to track the activity of individual users as they enter the ICN, move from host to host, access and move files, and run jobs, until they leave the ICN. New expert rules are being, and will be, defined which take into account the expanded information available, and which describe more

elaborate scenarios of invalid or suspicious user activity.

Since some kinds of invalid user activity, if allowed to continue, could result in break-ins or denial of service to legitimate users, another goal is the notification of appropriate ICN node(s) of extremely suspicious activity, and the development of responses by the node(s) to that activity. This would consist of taking direct action to stop an identified penetration attempt. The node's actions will have to be proportional to the extent that the monitored activity has deviated from what is considered valid, what damage could result from allowing an invalid activity to continue, and denial of service considerations. The criteria for such a response have yet to be determined.

Finally, we would like to identify and use a rigorous method by which to validate and verify the performance, consistency, and completeness of the NADIR expert rule base. This has become an even greater concern as the system is expanded to additional ICN nodes, and the resulting rule base has become correspondingly more complex.

11 Summary

NADIR demonstrates the feasibility of the automation of security auditing on a distributed environment such as the ICN, and the benefits of applying an expert system to the problem. It demonstrates the benefits of a phased approach to applying intrusion detection in a distributed environment. The working prototype is a start towards a longer-range goal of expanding the system to additional ICN nodes, and cross correlating their information to produce more complete profiles of user activity on the ICN.

12 Acknowledgments

We wish to acknowledge the contributions of Jimmy McClary (the ICN CSSO) who introduced us to the basic concepts, organized our funding, contributed enormously to our expert rule base, and supported us throughout the project. Valuable contributions to our rule base were made by members of the Operational Security Division, including Doty Alexander, Charlene Douglass, Lois Sylvia, Donna Stevens, and Mona Wecksung. We are indebted to Harry Martz for his expertise in statistics, and to Dorothy Merrigan and Steve Ruud for their contributions to the implementation of the NADIR system.

13 References

- [1] C. Browne. *DOTTYE - An Interlisp Program for the Analysis of the Network Security Controller Files* (Los Alamos National Laboratory, January 1984).
- [2] D. Denning and P. Neumann. *Requirements and Model for IDES - A Real-Time Intrusion Detection Expert System, Final Report* (Computer Science Laboratory, SRI International, August 1985).
- [3] M. Freiling, J. Alexander, S. Messick, S. Rehfuss, S. Shulman. *Starting a Knowledge Engineering Project: A Step-by-Step Approach* (The AI Magazine, Fall 1985).
- [4] D. Denning. *An Intrusion Detection Model* (IEEE Proceedings, 118-131, April 1986).
- [5] D. Denning, D. Edwards, R. Jagannathan, T. Lunt, P. Neumann. *A Prototype IDES: A Real-Time Intrusion Detection Expert System* (Computer Science Laboratory, SRI International, August 1987).
- [6] T. Lunt and R. Jagannathan. *A Prototype Real-Time Intrusion-Detection Expert System* (Proceedings of the IEEE Symposium on Security and Privacy, April 1988).
- [7] T. Lunt, R. Jagannathan, R. Lee, S. Listgarten, D. Edwards, P. Neumann, H. Javitz, A. Valdes. *IDES: The Enhanced Prototype A Real-Time Intrusion Detection Expert System* (SRI International, October 1988).
- [8] T. Lunt. *Automated Audit Trail Analysis and Intrusion Detection: A Survey* (Proceedings of the 11th National Computer Security Conference, October 1988).
- [9] M. Sebring, E. Shellhouse, and R. Whitehurst. *Expert Systems in Intrusion Detection: A Case Study, Draft* (National Computer Security Center).
- [10] T. Lunt. *Real-Time Intrusion Detection* (Proceedings of COMPCON, Spring 1989).
- [11] T. Lunt, R. Jagannathan, R. Lee, A. Whitehurst. *Knowledge-Based Intrusion Detection* (Proceedings of the 1989 AI Systems in Government Conference, March 1989).
- [12] K. Jackson. *Development and Analysis of User Authentication Profiles for an ICN Intrusion Detec-*

tion System (Los Alamos National Laboratory, June 1989).

[13] G. Tsudik and R. Summers. *AudES - An Expert System for Security Auditing* (Proceedings of AAAI Conference on Innovative Applications in AI, May 1990).

[14] T. Lunt, A. Tamaru, F. Gilham, R. Jagannathan, P. Neuman, C. Jalali. *IDES: A Progress Report* (Proceedings of the 6th Annual Computer Security Applications Conference, December 1990).

- END -

DATE FILMED

04 / 03 / 91

