LA-UR 95-1039

Conf-9505191--1

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

TITLE: **MISUSE AND INTRUSION DETECTION AT LOS ALAMOS NATIONAL LABORATORY**

AUTHOR(S): Kathleen A. Jackson, Michael C. Neuman, Dennis D. Simmonds, Cathy A. Stallings, Joseph L. Thompson, and Gary G. Christoph

SUBMITTED TO: DOE Computer Security Group Training Conference
Milwaukee, Wisconsin
May 1-4, 1995

DISCLAIMER

Los Alamos

Los Alamos National Laboratory
Los Alamos New Mexico 87545

DISCLAIMER

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

MISUSE AND INTRUSION DETECTION AT LOS ALAMOS NATIONAL LABORATORY*

Kathleen A. Jackson, Michael C. Neuman, Dennis D. Simmonds,
Cathy A. Stallings, Joseph L. Thompson, and Gary G. Christoph

Computing, Information and Communications Division
Los Alamos National Laboratory
Los Alamos, New Mexico U. S. A.

Abstract

An effective method for detecting computer misuse is the automatic auditing and analysis of on-line user activity. This activity is reflected in system audit records, in system vulnerability postures, and in other evidence found through active system testing. Since 1989 we have implemented a misuse and intrusion detection system at Los Alamos. This is the Network Anomaly Detection and Intrusion Reporter, or NADIR. NADIR currently audits a Kerberos distributed authentication system, file activity on a mass storage system, and four Cray supercomputers that run the UNICOS operating system. NADIR summarizes user activity and system configuration in statistical profiles. It compares these profiles to expert rules that define security policy and improper or suspicious behavior. It reports suspicious behavior to security auditors and provides tools to aid in follow-up investigations. As NADIR is constantly evolving, this paper reports its development to date.

1. Introduction

The Network Anomaly Detection and Intrusion Reporter (NADIR) performs misuse and intrusion detection for various systems in the Integrated Computing Network (ICN). The ICN is Los Alamos National Laboratory's main computer network. Serving over 9,000 users, it includes six Cray-class supercomputers (including a T3D), two massively parallel machines (CM200s), a cluster of sixteen IBM RS/6000s, over 10,000 smaller computers and workstations, file storage devices, network services, local and remote terminals, and data communication interfaces. If authorized to do so and using an approved access path, any user inside the Laboratory may access any host[1] computer from office workstations or terminals. The ICN consists of two completely separate, unconnected networks; a Secure (Classified) Network and an Open (Unclassified) Network. Outside users can access 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 the auspices of the United States Department of Energy.

[1]ICN host computers are computers to which normal users have controlling access.

Open Network through telephone modems, leased lines, or one of many world-wide networks.

NADIR's charter is to automatically audit, when feasible in near realtime, user activities on both the Open and Secure ICN. It uses expert system techniques to analyze data, and identifies anomalous patterns of activity, logs that activity, produces routine reports, and makes appropriate notifications. NADIR currently audits eight ICN systems, six in the Secure Network and two in the Open Network. Data from these systems is processed by five dedicated workstations. These audited systems are the "targets" of the NADIR service. For each audited target system, NADIR software resides on both the target system and the appropriate workstation. For convenience, throughout this paper the aggregate of all NADIR software and hardware is simply called "NADIR."

Because NADIR is a work in progress, development status differs among currently audited systems. In addition, because the type and functionality of the systems NADIR audits varies considerably, the reader will note implementation idiosyncrasies.

2. Framework

The goal of computer misuse and intrusion detection is to discover security violations on computer systems. Perpetrators may be either insiders (authorized users) or outsiders who clandestinely access a system. The first line of defense against all such violations is the institution of formality of operations. This is a way of doing business that emphasizes safeguards and accountability. Formality of operations includes institutional practices such as training, configuration management, and physical security measures.

However, several factors limit the efficacy of these measures. The first is human nature. Users often see security as an unwelcome diversion from the main thrust of their work. They resist learning security measures and procedures and frequently fail to apply them. Second, system managers must effect a compromise between conflicting concerns. For example, while it is more secure to compartmentalize activity, today's users require access to distributed resources. Third, systems frequently contain undetected vulnerabilities to attack and misuse. Finally, there is the threat of insiders who deliberately misuse their legitimate privileges [7].

Given these weaknesses, a second line of defense against abuse is the maintenance and analysis of system audit records. In theory, one can detect break-in attempts and other security violations by detecting abnormal or invalid user activity, changes in the system vulnerability posture, and other misuse indications. However, the traditional approach of manual analysis has generally proved unworkable. Human limitations restrict manual review to a sampling or cursory scanning of the large quantity of audit data and system status

information typically generated. This approach can target only a few obvious misuse scenarios; it may miss even these because of human error and because of the speed at which computer misuse occurs.

The limitations of manual review have long been apparent to security personnel at Los Alamos. While manual review by security auditors did reveal instances of misuse, there was no way to evaluate the general success or completeness of this effort. Large-scale manual audits of past data also proved cumbersome and time-consuming. It was obvious that automated review would be more effective. Such an analysis combines two essential components. First is the expert's knowledge of security problems. Second is the computer's ability to process and correlate, rapidly and accurately, large quantities of data. In addition, the speed of machine processing can allow an automated system to inform auditors of suspicious activity in time for them to trace and stop it. A system can even be programmed to undertake defensive measures itself, such as logging out a suspected intruder or removing a vulnerable machine from a network.

3. Overview

Los Alamos began developing NADIR in the late 1980s; it has been operational since 1990 [2, 5, 6, 11]. NADIR has evolved over time, being more-or-less constantly modified to reflect Los Alamos' changing network environment while being continuously expanded to audit more network systems. It currently audits two ICN services and four host supercomputers. The services are the Kerberos distributed authentication system and the Common File System (CFS), a centralized mass storage system. These two services are audited on both the Open and Secure Networks. The four audited supercomputers are all in the Secure Network.

NADIR analyzes the audit records kept by these systems, checking for a set of suspicious activities. It uses an expert system methodology in that its misuse scenarios were derived from interviews with security experts, and from detailed examinations of past audit record data. The expert system comprises rules that define invalid (not allowed) and anomalous (unusual) activity. Finally, the NADIR outputs reports of suspicious activity, provides the capability to perform investigations, and saves critical information.

NADIR differs significantly from the other misuse detection systems with which we are familiar, in that it combines two distinct computer security techniques. It looks both for suspicious *behaviors* and for suspicious *characteristics*. In the first category, it analyzes system audit records for evidence of suspicious behavior (as do other misuse detection systems). In the second category, NADIR analyzes the status of targeted systems for characteristics that indicate a vulnerable configuration or other evidence that misuse has taken (or is taking) place.

In continuing to expand NADIR, we maintained the design philosophy that served us well in the past [8]. NADIR is modular. It both integrates and separates information within different modules so that it is able to easily analyze data from several target systems simultaneously. This enables us to take individual target systems in or out of analysis (either deliberately or because of failures). It checks all data fed into the database for errors. It is designed to undertake data collection and analysis while avoiding any disruption of the normal conduct of business.

4. System Description

NADIR performs three basic functions. First, it provides a near realtime 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. Second, it provides the capability for ad-hoc analysis of past ICN user activity. This is useful for on-going investigations, background examinations, and audits. Third, it provides long term maintenance of a record of audit analysis, for documenting compliance with DOE security directives.

NADIR consists of three components. They are called UNICORN, KNADIR, and CNADIR. UNICORN (the UNICOS Realtime NADIR) analyzes the audit record from the Secure ICN's Cray supercomputers. KNADIR (Kerberos NADIR) analyzes the audit record from the Kerberos Network Security Controller (KNSC). CNADIR (the CFS NADIR) analyzes the audit record from the CFS. There are two KNADIRs and CNADIRs, one each for the Secure and Open Networks. At this time, there is one UNICORN, that being in the Secure Network.

Each NADIR component consists of client and server software. Client[2] software modules reside on each system that is targeted for auditing by the NADIR function. Server[3] software modules reside on dedicated workstations[4]. The client is charged only to select required data from the target system and provide it to the appropriate server in near realtime. KNADIR's client transmits data directly across the network to its server (Figure 4-1). CNADIR's client takes data from the CFS management area and moves it to a limited access file in the CFS storage area, where it is accessed by the CNADIR server on a daily basis (Figure 4.2). One UNICORN client transmits binary data directly across the network to the UNICORN server (Figure 4-1). The other three UNICORN clients have not yet been installed, and as an interim measure their data is moved weekly to the UNICORN server via the CFS.

---

[2]So called because it initiates connection over the network.

[3]So called because it waits for connection to be made.

[4]Throughout this paper we use several terms that may require clarification. "Client" and "server" refer to software modules. We use these terms when we discourse on NADIR software and the tasks it performs. A "target" system is one that is monitored by the NADIR service; i.e., is "targeted" by the NADIR service. "Target system" and "server workstation" refer to the systems upon which the client and server software respectively run.

Each server reads the appropriate audit records, summarizes this "raw" audit data into profiles, examines the profiles for signs of misuse, and reports its findings. Manual review of suspicious events takes place off-line. At regular intervals the server backs up the transmitted data, the profiled data, and all reports to long-term[5] storage on the CFS.

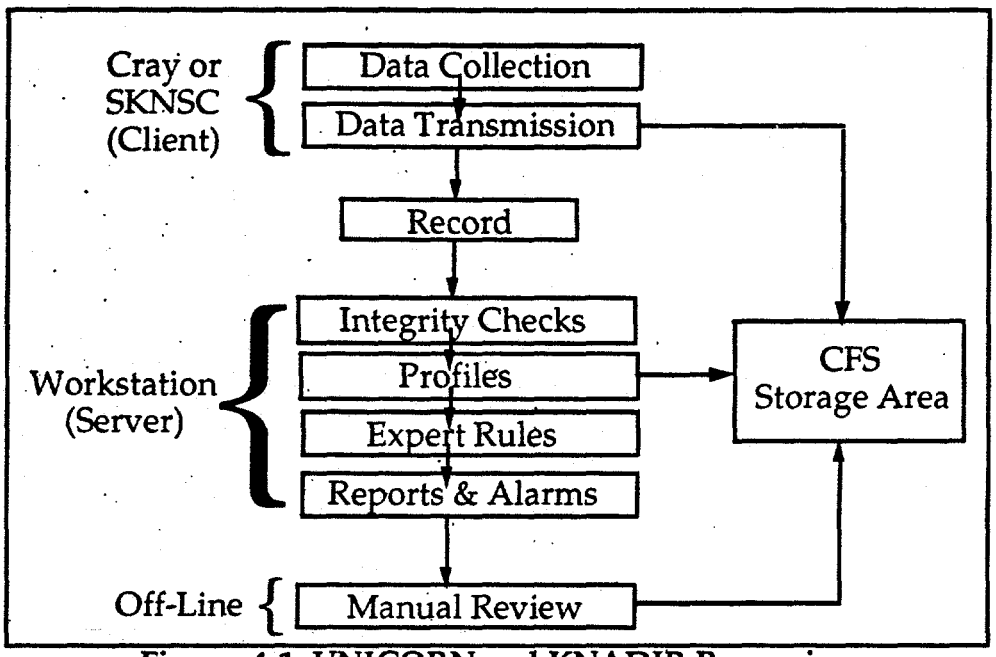

Figure 4-1: UNICORN and KNADIR Processing

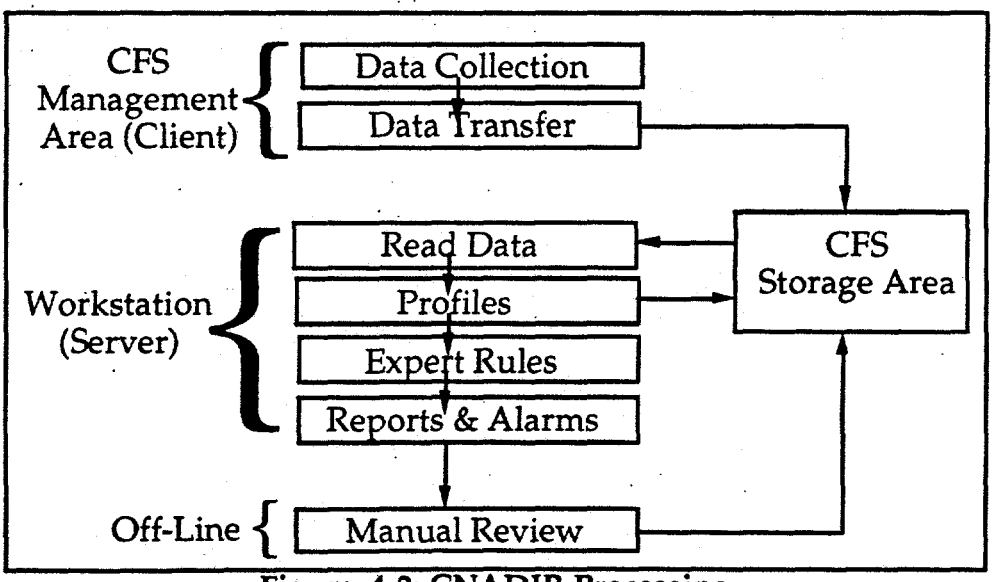

Figure 4-2: CNADIR Processing

[5]At least two years.

Isolation of the processing and alarm functions in the server has two major advantages. First, it provides a greater level of trust in the detection system. Second, it provides the capability of correlating activity from several target systems, thus increasing NADIR's sensitivity to misuse distributed over the network. NADIR could operate entirely on each target system, however the substantial increase in trust and flexibility that derives from separating the functions easily justifies the cost of a workstation and development of data transmission software.

All NADIR workstations are nodes of the two subnets of the ICN, one on the Secure Network and one on the Open Network. These subnets are called NADIR-net. NADIR-net supplies an isolated router-filtered environment for NADIR nodes in each Network. For example, Figure 4-3 illustrates the portion of the Secure ICN that is germane to the Secure NADIR-net. Data is sent from the Secure Network Cray supercomputers (Zeta, Tau, Epsilon, and Delta), the Secure KNSC, and the Secure CFS to the three NADIR nodes for processing. Notifications of critical events will soon be sent from each of the NADIR nodes to the Network Events Recording Device (NERD).

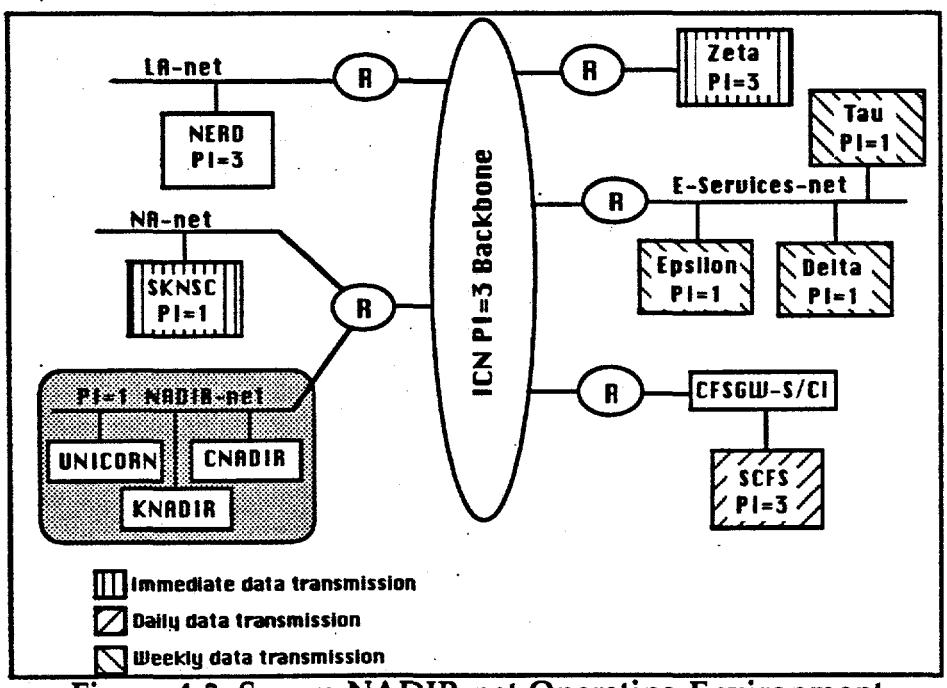

Figure 4-3: Secure NADIR-net Operating Environment

4.1 The Client

One ultimate NADIR objective is to analyze activity promptly, that is, in a near realtime[6] mode. This permits quick response to serious events. To meet

[6] At this time, we define near realtime as the detection and reporting of misuse within, at most, one hour of its occurrence.

this objective, each audited target system will eventually execute a client process that collects information and provides for its immediate transmission to a workstation-based server for database insertion and analysis. This necessary component of NADIR performs three functions:

1. On all audited systems the client collects selected audit logs. These are standard audit logs already maintained by each system. On the service nodes they are a record of the function performed by the node (i.e., ticket requests, file accesses). On the Crays they are the system logs of user activity.

2. On the Crays only, the client probes for signs of misuse and for errors in configuration (Section 4.1.1) that could potentially lead to vulnerabilities. This is done because normal users have controlling access over their space, and can knowingly or unknowingly make potentially dangerous configuration changes. The SKNSC and the management area of the SCFS are protected, limited-access systems, that allow no user codes, and may be accessed only by a small set of privileged users.

3. On all audited systems, the client transmits audit, misuse and vulnerability data (as appropriate) to the workstation-based server.

The client runs periodically[7], saving its current log position for the next collection pass. This approach permits asynchronous retrieval and transmittal of log data from the target system after system or network interruptions. With each pass, the client parses the audit logs and filters out bad data, preparing it for transmission to the server.

4.1.1 Extended Cray Auditing

Because the Crays are host computers, they allow controlling access by many normal users in a UNICOS operating system environment. Because some users may attempt to create system vulnerabilities and use them to gain privilege, this client not only audits the standard system audit logs but also runs automated security scanner. This scanner runs at regularly scheduled intervals[8]. It looks for suspicious *characteristics* on the target system (as opposed to suspicious *behaviors* noted in the system logs). In this way NADIR can identify misuse that may not show up in the standard system audit record. These include problems with system configuration, and signs of indirect user modification of the system. This component of NADIR is similar in function to Purdue University's COPS [4] and Lawrence Livermore National Laboratory's SPI [3], though it looks for a wider range of characteristics. It in-

[7]Currently, every six seconds. This collection rate is tunable.

[8]The periodicity is variable. For example, every minute the scanner sends all changed udb entries to the server, and every hour it sends the udb state (summary statistics about the udb). Then, as a sanity check, it sends the entire udb to the server once a day.

cludes enhanced Kuang (expert system) checking [1] for critical activity combinations. The security scanner checks for danger signs such as the following examples:

- Files modified by a daemon (e.g., sendmail or crontab writing to a file or changing file permissions).
- Minor changes in file permissions with indirect consequences (e.g., a user in a "system" group accidentally makes his or her home directory world writable).
- Valid file accesses with modifications that violate site policy.
- Significant changes in the /etc/udb file (e.g., security attribute additions, deletions, or modifications).
- Modification of critical system binaries.
- Flaws in critical file formatting (e.g., the /ect/group and /ect/passwd files).
- The protection status of system directories, files, and devices (e.g., world writable system directories and files, and world readable memory devices).
- Incorrect anonymous FTP configuration.
- File access permission problems (e.g., world writable files referenced by system crontab entries, the proper configuration of trust files, world writable user critical files (e.g. .rhosts, .login, .cshrc), and world readable .netrc files).
- Insecure daemons (i.e., sensitive programs such as TFTP and REXD).
- Invalid root configuration (e.g., system files and root login files owned by a user other than root, hosts.equiv and ftpusers configured incorrectly).
- World writable home directories.

4.1.2 Data Protection

We consider data protection to be an issue only on the Cray supercomputers, for the reasons covered previously in this Section. The presence of many users on each Cray poses a serious potential threat to the integrity of UNICOS audit and security scanner data, and to the client process itself. We believe this potential threat can be addressed by the security features provided by the UNICOS operating system. These features, if properly implemented, are sufficient to protect the data and process from tampering. Several protections, in particular privileged role separation, provide excellent assurance against alteration of security audit logs. Other audit logs may be protected using UNICOS MLS features, such as Mandatory Access Controls (MAC) or

Privilege Access Lists (PALs). The client software is protected by UNICOS security features such as privileged role separation. Finally, the client cooperates with the server in performing several integrity checks on all transmitted data (Section 4.1.3).

4.1.3 Data Transmission

Each client transmits all data to the appropriate server in binary format. The UNICORN client and server ensure data integrity cooperatively, using the following means:

- A *sequence number* that ensures the detection of repeated, missing, or out-of-sequence data packets. This helps detect not only data that has been deliberately tampered with, but also transmission mistakes resulting from system or client failures. The server logs all sequence failures. Such failures themselves could trigger an alarm.

- A *shared secret* that is used to verify the authenticity of each received packet. The shared secret sent by the client must match that kept by the server. The shared secret consists of a 32-bit key that can be changed as often as deemed necessary by NADIR. The server discards packets lacking the correct key, and logs all shared key failures.

- A *source identity check* that verifies the source (target system) identity of incoming data packets, and whether each packet's internal labeling matches that particular Target system machine. The server logs all invalid sources.

Transmitted data is not encrypted because we consider the network segments between target systems and server workstations to be both physically and logically secure. The only nodes (machines) allowed on these segments are special-purpose network services that are physically and internally secure. Access to them is limited to authorized network personnel. No computers accessible to normal users are allowed on these segments. Consequently, there is no way the audit data (or in the case of UNICORN, the shared secret) can be observed by unauthorized users while transiting this network segment. However, if the need arises, encryption can be implemented.

4.2 The Server

Each server resides on a SUN™ (or SUN clone) with 74-96 MBytes[9] of memory and one or two 1+ GByte disks. The Sybase™ relational database management system is used to organize the data structure and to enable easy

[9]The server currently requires only 32-48 MBytes of memory.

data access. The server software is written in the C language and Transact-SQL (Sybase's version of SQL). The server performs five functions:

1. Decodes the incoming binary data from the client and performs integrity checks on that data,

2. Formats the data for use by the server,

3. Summarizes this "raw" data into profiles of both individual user activity and composite system activity,

4. Examines the profiles for signs of misuse, and

5. Reports its findings.

4.2.1 Data Receipt

The server decodes each incoming data packet and checks its integrity as required (Section 4.1.3). It determines the type of data and activates appropriate routines for parsing and resolution.

4.2.2 Data Formatting

After the server receives an audit record, it first parses the record and places it in a canonical format. We do this to provide a standard data interface to the server. A standard interface is advantageous for two reasons. First, the server parsing function does not have to be modified to handle different data formats. All modifications are limited to this one function; the core of the server remains unchanged. Second, we wanted to implement the standard audit data interchange format currently proposed in the computer security community [10]. Widespread use of this format will allow the sharing of audit record information from different misuse detection systems. Such a common format is much desired by developers of audit record analysis tools. It includes 'wild card' fields that can be used for system-specific information, and event-specific information. Each canonical audit record describes a single event, and is formatted as is the example in Table 4-1.

4.2.3 Profiles

Each server maintains profiles for each unique user and for a composite of all users on the target system being audited. The profiles summarize the raw audit data, making it easier to store, interpret, and analyze. Profiles are saved periodically to the ICN's permanent file storage.

4.2.3.1 Profile Design

Profiles are summary statistics of activity over some defined interval. The server maintains two kinds of profiles; individual and composite. Individual

profiles summarize the activity attributed to specific users. Composite profiles summarize the activity of an entire system.

Table 4-1: UNICORN Audit Record

BASIC DATA	
Timestamp	The date and time at which the activity occurred.
Event Type	The type of event described in this audit record.
Process ID	The current process identifier.
Outcome	The event outcome. If successful, a return code indicates the type of activity. If unsuccessful, an error code indicates the type of failure.
User IDs	A full description of the subject's user identifiers.
Group IDs	A full description of the subject's group identifiers.
Session ID	The session to which the process belongs.
Security Level	The security level of the event subject, whether user or process.
Object Description	Information about the objects affected by the event, if any.
MISCELLANEOUS DATA	
Source-Specific Data:	
Host	The host on which the attempted activity occurred.
Partition	The security partition in which the attempted activity occurred (a Los Alamos specific attribute).
Event Source	The source of the activity. For example, the workstation from which a user logged on.
Compartment	The security compartment of the attempted activity.
Category	The integrity category of the attempted activity.
Event Data:	
Activity Data	The data specific to the type of activity being reported; it describes the event itself. Each Event Type has its own set of possible Activity Data values.

Each profile consists of a number of *segments*. Each segment corresponds to a certain time interval. The composite profiles are more detailed than the individual profiles: each full day's data is broken into 24 segments, one per hour. The choice of one hour for the profiles' finest granularity seemed appropriate to us, but is configurable to a shorter or longer period. Each segment has numerous *fields* that summarize some aspect of the subject of the profile (individual user or system) during that time interval. Many of these are count statistics such as the number of logon failures or ticket requests during the interval. These statistics are updated each time a relevant audit record is received.

The first two segments of both profiles describe the *current hour* and *current day* thus far. The remaining segments describe a *moving week* of data, of which the seventh day is the most recent day for which complete data are available. For example, if today is Thursday (the current day), the moving week includes data from the previous Thursday through yesterday (Wednesday). As each current hour is completed the current day segment is updated and the current hour segment is re-initialized. As each current day is completed the current moving week is updated and the current day segment is re-initialized. For example, at the end of Thursday, the moving week shifts to last Friday through Thursday.

Table 4-2: Individual Profiles

segment	interval	
1	current hour	
2	current day	
3	moving week	day 1
4		day 2
5		day 3
6		day 4
7		day 5
8		day 6
9		day 7

Table 4-3: Composite Profiles

segment	interval	
1	current hour	
2	current day	
3-26	moving week	day 1 (24 hours)
27-50		day 2 (24 hours)
51-74		day 3 (24 hours)
75-98		day 4 (24 hours)
99-122		day 5 (24 hours)
123-146		day 6 (24 hours)
147-170		day 7 (24 hours)

4.2.3.2 Individual Profiles

Individual profiles provide a summary of activity for each authorized user on the audited system. They consist of one record for each unique ICN user. Individual profile fields hold count statistics for different types of user activity. These may be derived both from the audit record, and from the active security scanner, as appropriate. The misuse recorded here, if any, is that which can be attributed to a specific user.

4.2.3.3 Composite Profiles

The composite profile provides a summary of activity and misuse indications (if any) not attributable to a single or specific user, and vulnerability posture for the whole target system. This profile consists of one record for each audited target system. Composite profile fields hold count statistics for dif-

ferent types of activity. These may be derived both from the audit record, and from the active security scanner, as appropriate.

4.2.4 Profile Analysis

The NADIR server compares the profiles to expert rules that encode our security policy and unusual or suspicious activity. One set of rules applies to individual user activity, another to composite activity.

4.2.4.1 Evaluation Schedule

Profiles are evaluated using the expert rules described in Section 4.2.4.3. The accumulated hour, day, and week profiles are evaluated separately, using different sets of rules. Evaluation is data driven; the timestamp within the incoming data is used to decide when it is time to evaluate. It is performed as follows.

At the beginning of a new hour:

1. The hour just finished is evaluated.
2. The hour's data is added to the current day.
3. The day thus far is evaluated.

At the beginning of a new day:

1. The day just finished is evaluated.
2. The oldest day is dropped from the moving week.
3. The new, just completed, day is added to the moving week.
4. The new moving week is evaluated.

This approach has a number of advantages. First, all profiles are evaluated, and thus recognizable events detected, within a maximum of one interval (currently one hour). This evaluation interval can be shortened if desired. Second, there is no discontinuity in the data being evaluated. Third, a history of past activity (at least a week) is maintained on-line. Fourth, the process lends itself well to near (within the smallest interval) realtime processing. Fifth, the data-driven approach enables NADIR to adjust easily to target system downtime, its own downtime, or missing data.

4.2.4.2 Rule Development

An important first step in developing our expert rule set was interviewing the experts -- ICN security personnel. Interviews of administrators charged

with establishing and enforcing the Laboratory's security policy were straightforward. Interviewing security auditors[10] took time but was extremely fruitful. We found that auditors rely on an undocumented combination of extensive knowledge of the ICN, experience with previous intrusions or misuses, and instinct.

Another important part of our rule development was a statistical analysis of the audit record from the target system. We spent months reviewing the raw audit data. From this review we learned enough to implement an initial set of profiles, from which we calculated the characteristics of average user and system behavior. We then studied those profiles that deviated significantly from the norm to determine which deviations comprised a suspicious event, particularly if combined with other indications.

This process of interviews and statistical analysis led to the definition of an initial rule set. This was then tested against months of audit data. This process of testing and manually revising our rule set is an ongoing one, as we continually aim to improve the accuracy of our system.

4.2.4.3 Rule Implementation

Expert rules are applied to the individual and composite profiles at the end of each interval, as described in Section 4.2.4.1. We have defined expert rules for three different intervals. *Hour rules* are applied at the end of each hour. *Day rules* are applied at the end of each hour, for the day thus far (one to twenty-four accumulated hours). *Week rules* are applied at the end of each day for the current running week (the current just-completed day plus the previous six days).

The server rule base comprises four logical rule filters; each designed to isolate certain types or levels of anomalous activities. We started by abstracting ICN security policy and well-defined invalid and suspicious behavior into rules that form the Primary Filter. Report requirements supplied rules for the Report Filter. These two filters are currently implemented; the following two are in the design stage. Analysis of Primary Filter output will result in the Event Filter. The Alarm Filter will determine the alerts resulting from each event. The server currently activates (and will activate) the rule base filters in order, as illustrated in Figure 4-4.

- The **Primary Filter** applies rules to the profiled data. These rules are straightforward descriptions of simple activities, each serving to distinguish a separate feature of anomalous behavior. The Primary Filter applies these rules individually; it does not correlate one with another. It

[10]Those who have had to manually review the audit record.

assigns a Level-of-Interest to each anomaly defined by these rules. The results of this analysis are stored in the Report Table.

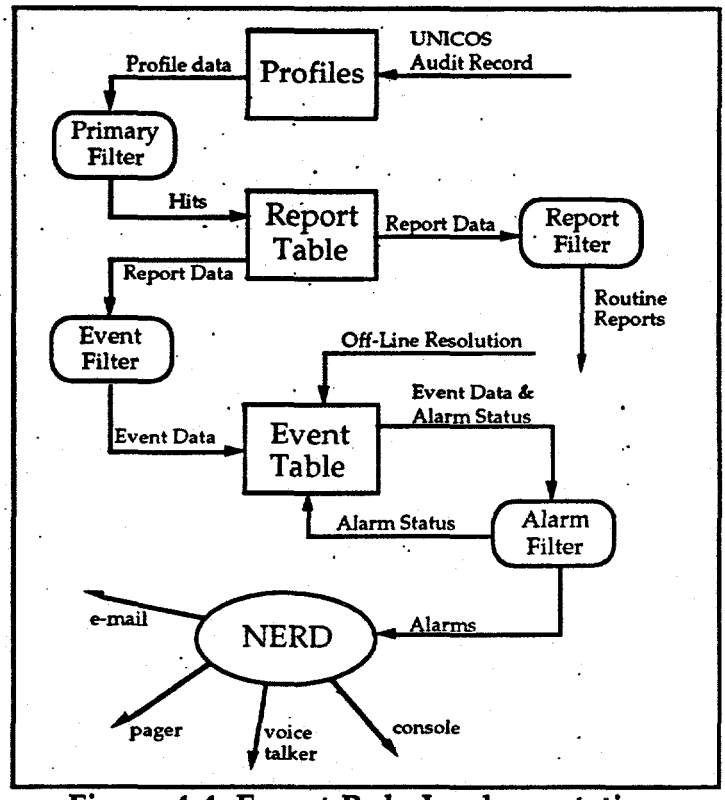

Figure 4-4: Expert Rule Implementation

- The **Report Filter** applies rules to the anomalies output by the Primary Filter, to produce routine reports of anomalous behavior.

- The **Event Filter** will apply rules to the anomalies identified by the Primary Filter. These rules will try to identify patterns of anomalous activity that have a good chance of being systematic misuse (events). They will specify what action to take when events are found, such as the scheduling and content of warning messages. The results of this analysis will be stored in the Event Table. Each event will remain 'active' in the Event Table until security auditors resolve it off-line. It will then be flagged 'inactive' by the auditors. Inactive events will be flushed from the table at regular intervals.

- The **Alarm Filter** will apply rules that manage appropriate notification of urgent or critical anomalous activity. It will determine what level of alarms should be sent, and to who, and manages their frequency.

We encode our expert rules in a condition-action (if-then) form. The condition (if) describes a suspicious profile scenario or a violation of security policy. The action (then) specifies setting a level of interest for the relevant user (or

composite user) profile. Table 4-4 gives an example of one complete rule (from UNICORN). This rule focuses on the ratio of logon failures to total logons. Variable definitions are not included because NADIR rule specifics are sensitive.

Table 4-4: RULE-CU-A-006 (Failure ratio)

IF end of an hour
ctot_hour_logons = tot_csucc_logons + tot_cfail_logons
current_ratio = tot_cfail_logons/(ctot_hour_logons)
THEN
IF (ccurrent_ratio is > $case_n$_min AND ≤ $case_n$_max)
THEN
IF ctot_hour_logons > $ctot4_n$_hour_max
THEN Set Rule CUH-006 to 4 in the Report_Table
ELSE IF ctot_hour_logons > $ctot3_n$_hour_max
THEN Set Rule CUH-006 to 3 in the Report_Table
ELSE IF ctot_hour_ogons > $ctot2_n$_hour_max
THEN Set Rule CUH-006 to 2 in the Report_Table
ELSE IF ctot_hour_logons > $ctot1_n$_hour_max
THEN Set Rule CUH-006 to 1 in the Report_Table
EXPLANATION: Greater composite failure ratio than is normal for the previous hour.

IF end of an hour
ctot_day_logons = tot_csucc_logons + tot_cfail_logons
current_ratio = tot_cfail_logons/(ctot_day_logons)
THEN
IF (current_ratio is > $case_n$_min AND ≤ $case_n$_max)
THEN
IF ctot_day_logons > $ctot4_n$_day_max
THEN Set Rule CUD-006 to 4 in the Report_Table
ELSE IF ctot_day_logons > $ctot3_n$_day_max
THEN Set Rule CUD-006 to 3 in the Report_Table
ELSE IF ctot_day_logons > $ctot2_n$_day_max
THEN Set Rule CUD-006 to 2 in the Report_Table
ELSE IF ctot_day_logons > $ctot1_n$_day_max
THEN Set Rule CUD-006 to 1 in the Report_Table
EXPLANATION: Greater composite failure ratio than is normal for the day thus far.

IF end of a day
ctot_week_logons = tot_csucc_logons + tot_cfail_logons
current_ratio = tot_cfail_logons/(ctot_week_logons)
THEN
IF (current_ratio is > $case_n$_min AND ≤ $case_n$_max)
THEN
IF ctot_week_logons > $ctot4_n$_week_max
THEN Set Rule CUW-006 to 4 in the Report_Table
ELSE IF ctot_week_logons > $ctot3_n$_week_max
THEN Set Rule CUW-006 to 3 in the Report_Table
ELSE IF ctot_week_logons > $ctot2_n$_week_max
THEN Set Rule CUW-006 to 2 in the Report_Table
ELSE IF ctot_week_logons > $ctot1_n$_week_max
THEN Set Rule CUW-006 to 1 in the Report_Table
EXPLANATION: Greater composite failure ratio than is normal for the current week (the last seven days).

4.2.5 Reports

The server can report detected activity in several ways, including scheduled routine reports and, where required, immediate alarms. It also supports ad-hoc investigations, during which it can provide detailed reports of raw or profiled data in response to auditors' specific queries.

4.2.5.1 Immediate Reports

Critical events are reported when they are detected. These are events that require prompt investigation. The server assigns a priority to each event, depending on its criticality. It then outputs an announcement to the server's console. Soon it will notify a dedicated ICN system whose function is to log and report events for the entire ICN. This system is the Network Events Recording Device, or NERD [12]. The NERD provides four levels of notification; a broadcast using synthesized speech, paging, e-mail, and console display. The NERD will undertake appropriate notification based on priority, responsible individuals, and other information supplied by the server.

4.2.5.2 Scheduled Reports

The server routinely generates reports. These reports cover the just-completed hour, the just-completed day and the just-completed running week (the just-completed day and the prior six days).

These reports consist of a one-page activity summary, e.g., the number of active users during the report interval, and the number of successful and unsuccessful user requests during that interval. There is also a set of graphs of different types of activity, plotted over time with a granularity of one hour. These are useful for representing abnormal patterns, such as an unusual spurt of off-hour usage. The rest of the report summarizes the results of the expert rule analysis. It lists suspicious users in descending priority order (from the most suspicious to the least), with a list of the rules each has triggered.

To support investigator follow-up, the server also produces a more detailed daily report that includes all raw data from the audit record. This data is the unprocessed audit record as received from the audited target system. Auditors occasionally need to review this data while attempting to ascertain what has happened during an event.

The server stores these regularly scheduled reports in a secure portion of our permanent file storage, where they can be accessed and reviewed only by authorized personnel.

4.2.5.3 Ad-Hoc Reports

The server can produce reports on demand. On-the-spot reports have proved invaluable in analyzing ongoing events. Finally, we use raw or profiled data that the server has saved to permanent file storage to perform ad-hoc background analyses of current and past activity. Authorized security personnel can examine this data using Sybase's built-in facilities, or pipe data to a statistical software package for more detailed analysis.

4.3 Off-Line Activities

At intervals throughout each day security auditors review the hour reports, the day's report, and the current running week's report. When required, they review immediate alarms. They examine each anomalous event and decide whether to investigate it further. They analyze user or system audit data and may interview indicated users. An investigation may result in a warning to a user, or the user losing, at least temporarily, their ICN privileges. More often, it results in a learning experience for the user. The auditors file a short report at the completion of each investigation, giving details of its resolution. These reports, and periodic reviews of NADIR by the security auditors, provide valuable feedback from which we continually try to improve the system.

4.4 Data Integrity

We take care to protect the integrity of the target system audit record throughout its life span. We treat it as sensitive because of its importance to security and accounting, and because its integrity is critical to ensure the validity of the intrusion and misuse detection process. Only a small set of system managers have access to the audit record on the target systems, in the file storage archive (the CFS), and throughout the process of transmitting and analyzing it. We transmit the audit records over protected lines and back them up routinely to password-protected files on the CFS. Only authorized security auditors may examine any portion of the data or the reports generated by NADIR. We treat the results of investigations as sensitive. If investigations uncover a system of network vulnerability, the report is classified. Such management activities are essential to the integrity of, and user trust in, the whole audit process [9].

5. Future Directions

We plan to continue our progress towards an optimally effective misuse detection system. During the couple of years we expect to:

- Install the UNICORN client on the three remaining Secure Network Crays, and thus move their processing to near realtime (this may require the use of more than one workstation to support four servers)
- Install the UNICORN client on the Open Network Crays and provide the required workstation servers for them
- Expand UNICORN data collection and analysis and on-line vulnerability checks (to make UNICORN more effective)
- Complete a user-friendly graphical user interface for investigative personnel

- Provide near realtime notification of critical events using the NERD
- Expand NADIR auditing to the network routers and the Laboratory's large (10,000 plus) network of workstations
- Expand NADIR auditing to MERCURY, a system that will (in the future) allow authorized users to move unclassified files from the Secure Network to the Open Network
- Add a NADIR component that will review reported anomalous activity from all other NADIR components, thus enabling the correlation and analysis of network-wide anomalous activity

Another future goal is to explore the possibility of supplementing our expert rulebase with a component that "learns" typical behavior for each user, then reports deviations from these norms.

Acknowledgments

We acknowledge with gratitude the contributions of Jimmy McClary, who introduced us to the basic concepts of misuse detection, obtained our initial funding for NADIR, and supported us throughout the various incarnations of the project. We are indebted to Sharon Wilhelmy, who has reviewed NADIR's reports for three years, and who thus has been a valuable source of feedback on NADIR's functioning. As a result of her experience with NADIR, Sharon has been instrumental in helping us design NADIR in a way that maximizes its usefulness to Los Alamos security auditors. We thank Steve Smaha (of Haystack Laboratories, Inc., Austin, TX) for his suggested standard audit record format. The format of the canonical NADIR audit record was derived directly from this standard.

References

[1] R. Baldwin. *Rule-Based Analysis of Computer Security* (Massachusetts Institute of Technology, June 1987)

[2] K. Jackson. *Development and Analysis of User Authentication Profiles for an ICN Intrusion Detection System* (Los Alamos National Laboratory, Technical Report, June 1989)

[3[*SPI - Security Profile Inspector, Installation and User's Manual* (Lawrence Livermore National Laboratory, 1989)

[4] D. Farmer, E. Spafford. *The COPS Security Checker System* (Proceedings of the Summer Usenix Conference, June 1990)

[5] K. Jackson, D. DuBois, and C. Stallings. *A Phased Approach to Network Intrusion Detection* (Proceedings of the United States Department of Energy Computer Security Group Conference, May 1991, LA-UR-91-334)

[6] K. Jackson, D. DuBois, and C. Stallings. *An Expert System Application for Network Intrusion Detection* (Proceedings of the 14th National Computer Security Conference, October 1991, LA-UR-91-558)

[7] J. Hochberg, K. Jackson, J. McClary, D. Simmonds, *Addressing the Insider Threat* (Proceedings of the United States Department of Energy Computer Security Group Conference, May 1993, LA-UR-93-1181)

[8] J. Hochberg, K. Jackson, C. Stallings, J. McClary, D. DuBois, J. Ford. *NADIR: An Automated System for Detecting Network Intrusion and Misuse* (Computers and Security, Elsevier Science Publishers Ltd., Volume 12, Number 3, May 1993, LA-UR-93-137)

[9] K. Jackson, *Management Issues in Automated Audit Analysis: A Case Study* (Proceedings of the 8th European Conference on Information Systems Security, Control, and Audit, September 1993, LA-UR-93-2520)

[10] S. Smaha. *A Common Audit Trail Interchange Format For UNIX* (Haystack Laboratories, Inc., Technical Report, May 1994)

[11] K. Jackson, M. Neuman, D. Simmonds, C. Stallings, J. Thompson, and G. Christoph. *An Automated Computer Misuse Detection System for UNICOS* (Proceedings of the Cray Users Group Conference, October 1994, LA-UR-94-3378)

[12] D. Simmons. *Network Event Recording Device: An Automated System for Network Anomaly Detection and Notification* (Proceedings of the Internet Society Symposium on Network and Distributed System Security, February 1995, LA-UR-94-2790)