Global Information Assurance Certification Paper

Information System Activity Review

Achieving Compliance with §164.308(a)(1)(ii)(D)

GIAC HIPAA Security Certificate (GHSC)
Practical Assignment
Version 1.1

Lee A. Kadel
CCA, CCSA, GSEC, MCP, MCSE, NT-CIP
GHSC Certificate Challenge
March 16, 2005

Table of Contents

Table of Figures

Abstract

This paper examines section §164.308(a)(1)(ii)(D) of the HIPAA Security Rule, and develops both policy and procedures to meet the requirements listed therein. This section of the rule requires a covered entity to examine system log files, access log files, and other system-generated logs and reports on a regular basis.

The study begins by describing the network and security environment of a fictional healthcare company, and then proceeds to examine the HIPAA requirement. Once the environment is defined, the study presents a formal policy to deal with the requirements; and having established policy, develops an implementation plan using the NetForensics Security Information Management (SIM) system. Finally, a review of the system is performed to verify that it affords compliance with the policy and the HIPAA standard.

Define the Environment

GIAC Healthcare Inc. (GHI) is a large healthcare holding company made up of three wholly owned and two jointly owned hospital systems across three states. Overall, GHI operates 17 hospitals and 63 clinics employing over 3500 doctors, both staff and contract. The system has a total staff of over 24000 clinical and non-clinical employees. Annual revenues exceed $3.8 billion; profit is over $153 million.

The server environment supports over 750 Windows® servers and over 100 Unix and Linux hosts. Microsoft© operating systems supported include Windows NT 4.0®, Windows 2000®, and Windows Server 2003®; Unix systems include HP Unix®, AIX®, and Tru64®; Linux hosts are all RedHat®, with a mixture of v.7.3, v.9.0, and AS-3. The 8800+ user desktop systems are Windows XP Professional®, with less then a dozen RedHat® Linux machines in use by individuals in the IT department, and a few Windows98® PC's required to run one legacy application that is scheduled to be retired.

The Information Services department is made up of two main groups: application analysts who manage the business aspects of all clinical and financial applications, and the network analysts who build and maintain all of the hardware and infrastructure for the enterprise. These groups are further divided into teams, each responsible for a particular application or network function. The Information Security team consists of three individuals – one lead analyst, and two security analysts; they report to the Director of Security, who in turn reports directly to the CFO (who is also the CISO).

The diagram on the next page shows the major features of the GHI network. In the interest of space, switches, hubs, internal routers, and some other features have been purposely left out. Also not shown in the diagram are two Sonet rings – the first connects the Server network to the Hospital System 1 network; the second connects the server network to the Hospital System 2 network. Another ring is planned to replace the current T1 connection to the Hospital System 3 network; the three rings will eventually be combined. In addition, each network is further divided so that each individual hospital and clinic is on its own subnet; in the larger hospitals, the subnets

are further divided so that each floor is on its own subnet.

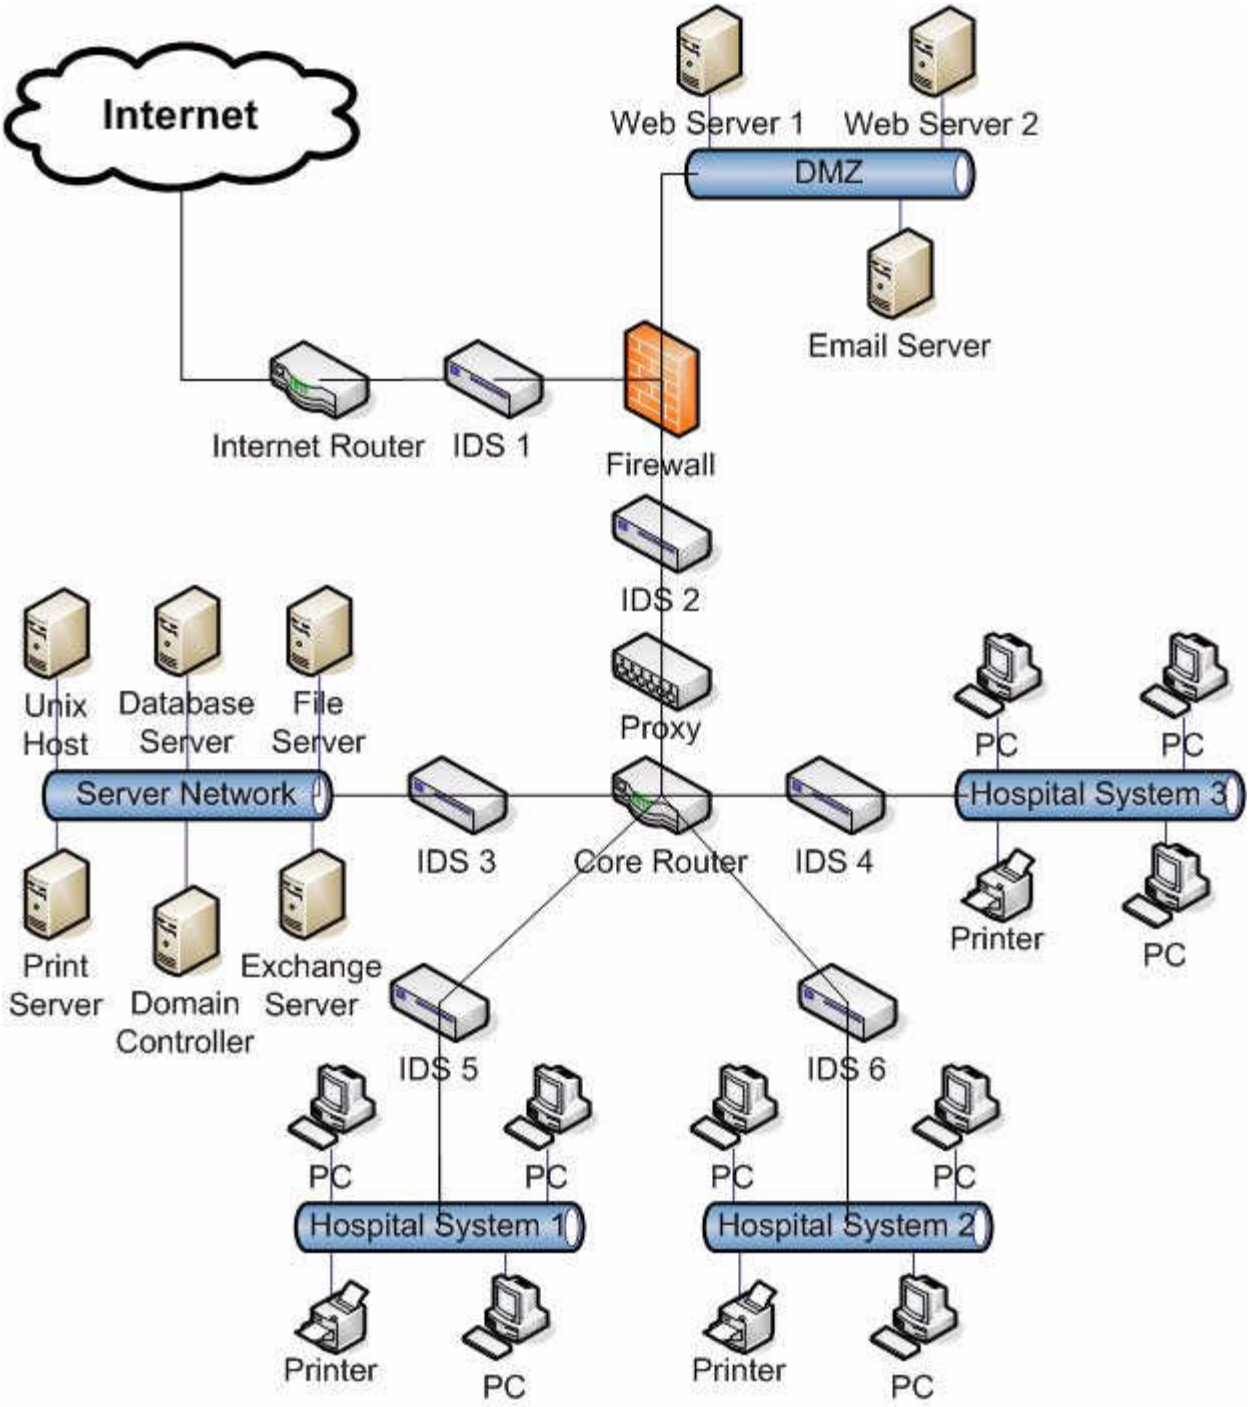

Figure 1 – GIAC Healthcare Inc. Network

Assignment One: Explanation

This specification of the HIPAA Security Rule requires covered entities (CE's) to, "Implement procedures to regularly review records of information system activity, such as audit logs, access reports, and security incident tracking reports." (45 CFR Parts 160, 162, and 164, §164.308(a)(1)(ii)(D)). The wording of this specification implies three main points: that systems can be configured to generate log files; that the log files contain sufficient relevant information; and that those log files can be collected, stored, and made available for review. The specification does not address such topics as long-term storage, data retention periods, or reporting.

The intent of this specification is to establish a mechanism to help identify policy violations, security incidents, and other suspicious behavior. It relates directly to the specification regarding incident handling, §164.308(a)(6)(ii)(A), in that identification of an issue as a result of log review will often lead to some investigative and/or remedial efforts. Indeed, the log files will most likely be one of the most important resources in any forensic examination. They are also used to identify vulnerabilities and weaknesses as part of a risk assessment and analysis (SANS, 2004).

Different device types, and even different models of devices, will generate different events and event types. The logs obtained from a firewall, for example, will be entirely different than those taken from a Windows© server – different information is being tracked, and the log files are stored in different formats. Reviewing individual log files will often result in little, if any, useful information; by examining series of events and patterns from several separate devices and/or systems, security issues may be identified, and the need for further action justified (SANS, 2004).

The volume of logs and events may often be so large that manual review becomes impossible – events may be generated faster than they can be read and reviewed, even by a dedicated team of people. In this situation, there are essentially two choices: apply restrictions and limits on the level of logging and the detail of the logs, or implement a centralized event management system to collect, aggregate, and correlate events, and report or notify when suspicious activity is found. Restricting the logging process is less desirable, because important information may be missed; however, the cost of an event management system may be prohibitive, particularly in smaller institutions. The rule does not specify any particular method for gathering and reviewing logs and reports; only that it must be done. Entities must evaluate and choose the method(s) that work best in their particular environment.

The choices made by a covered entity regarding compliance with this specification should be published in a formal policy statement, along with supporting process and procedure documents. Policies must be updated regularly, and the processes and procedures adjusted as systems and devices are added, removed, or reconfigured. In this way, ongoing compliance under the specification is assured, and the effectiveness of the security measures maintained. Failure to do so may, in itself, be considered a violation of the security rule.

Assignment Two: Policy (option B)

Log File Review Policy

1.0 Purpose
The purpose of this document is to define the GIAC Healthcare Inc. (GHI) policy regarding:

- Systematic review of system log files within the GHI network environment.
- Monitoring and auditing for unauthorized access attempts and other security incidents.

2.0 Scope
This policy covers all computer devices owned by GHI, as well as those which are under GHI's control but may not be owned by GHI. Log files include, but are not limited to, the following:

- Windows© system event logs
- Unix/Linux syslogs
- Firewall event logs
- Intrusion Detection System (IDS) event logs
- Network router and switch event logs
- VPN Concentrator access logs
- Application activity logs
- Internet proxy logs

3.0 Policy

3.1 Ownership and Responsibilities
The development and implementation of necessary systems, processes, and procedures to comply with this policy is the direct responsibility of the GHI Information Security team, who are herein granted the authority to access any/all systems and their respective log files that are deemed necessary to the implementation of this policy.

The GHI Information Security team is responsible for monitoring device and system logs, and identifying events or series of events that may constitute a security incident, as defined in the GHI Security Incident Handling Policy document. The Information Security team is also responsible for responding as necessary to investigate, control, contain, and eradicate or remediate those incidents identified, and is the assigned leader of all such efforts. Further, the Information Security team has authority to direct any/all actions, up to and including the authority to shut down all network access until such time as the incident is no longer an active threat. All responses and actions must comply with the GHI Security Incident Handling Policy and its associated processes and procedures.

3.2. Log Collection, Aggregation, and Correlation

A dedicated log collection and aggregation system will be used to collect and analyze log files in real time from any/all systems identified as High Risk based on GHI's then current risk assessment, or which contain data classified as Confidential or higher, based on GHI's standard data classification system.

- The system must be capable of collecting and retaining log files for the retention period specified below (see section 3.5).
- The system must be capable of normalizing and aggregating events collected.
- The system must be capable of creating customized events, based on the individual and aggregated events received.
- The system must be capable of alerting and notifying appropriate staff in the event of an identified security incident, as defined in the processes and procedures documents.
- The system must be capable of generating reports including, but not limited to, risk level, threat level, attacks, access attempts, and other reports as requested or required by leadership.

3.3 Service Degradation and/or Interruption

The collection and analysis of system log files must not significantly impact the speed, reliability, or accessibility of the GHI network or any devices contained therein.

3.4 System Availability

The log collection system must be available 24x7x365, excepting for normal downtimes for scheduled maintenance.

3.5 Data Retention

All data collected must be actively accessible for a period of at least 180 days for use in reporting or forensic examination. After the active period, data must be archived and stored for a period of 7 years, according to the GHI Data Archival and Storage Policy.

4.0 Enforcement

Any system found to be non-compliant with this policy, unless specifically exempted and with the proper documentation, will be shut down or removed form the GHI network until such time as it can be made verifiably compliant.

Any employee caught tampering with or otherwise changing log files will be subject to disciplinary action, up to and including immediate termination of employment.

5.0 Revision History

June 29, 2004 – Initial policy document created.
July 1, 2004 - Policy reviewed and approved.
Next scheduled review June, 2005.

Assignment Three: Procedures (option B)

After six months of research and investigation, GHI has chosen to implement the NetForensics© SIM Suite (NF) product to collect, aggregate, and correlate events, and to issue alarms and notifications when an event of sufficient severity is identified. A risk assessment and analysis was performed to determine the liability and exposure that would result from not implementing the system and it was ascertained that the cost of the system was justified.

Implementation of the NF system was in four phases. The first phase included the purchase and installation of necessary hardware and software, and configuration of those devices determined to be of greatest use in establishing a baseline. Other devices were added in the subsequent phases of the implementation.

Hardware Configuration

- (1) HP Proliant DL580, 4 x 3.2Ghz Xeon, 8Gb RAM, 2 x Gb NICs, 6 x 72Gb HDD
- (1) HP SA1000 mini-SAN, 22 x 72Gb HDD, fiber-channel connection to HP580
- (2) HP Proliant DL380, 2 x 3.2Ghz Xeon, 4Gb RAM, 2 x Gb NICs, 3 x 72Gb HDD
- (2) HP Proliant DL360, 1 x 3.2Ghz Xeon, 2Gb RAM, 2 x Gb NICs, 2 x 72Gb HDD

Phase 1 Devices

- Internet Router
- Firewalls
- IDS Devices
- Domain Controllers
- Email Servers

Phase 2 Devices

- Core Router
- VPN Concentrator
- Unix Hosts containing ePHI
- File Servers containing ePHI
- Database Servers containing ePHI

Phase 3 Devices

- Web Servers
- Print Servers
- Unix Hosts containing Financial or HR Data
- File Servers containing Financial or HR Data
- Database Servers containing Financial or HR Data

Phase 4 devices

- All other devices are evaluated on a case-by-case basis to determine exposure and liability. If sufficient cause is found, the system or device will be added.

Between each phase, an analysis was performed to verify and validate that all parts of the system were functioning as required. After phase 1, a baseline was created to establish the parameters of normal operation; information retrieved from subsequent phases was compared to this baseline to further refine it and arrive at the definition of a normal state within the environment. As devices and systems are added, removed, or changed within the environment, the baseline and resultant normal state are updated.

The NetForensics application and its associated Oracle database application are installed on the DL580; the database itself is on the SA1000 mini-SAN. The remaining servers are used to host the various agents needed to collect log files and normalize events before sending them to the database. The systems are configured as follows:

NFSYSTEM1
- Proliant DL580 w/mini-SAN
- Windows 2000 Advanced Server®
- NetForensics application
- Oracle database application

NFSYSTEM2
- Proliant DL380
- Windows 2000 Advanced Server®
- NF Engine (aggregation and event forwarding)
- NF Agent for HPUX (HP Unix events)
- NF Agent for Checkpoint Firewall-1 (firewall events)
- NF Agent for CIDS4 (IDS events)
- NF Agent for Syslog (router, switch, and VPN events)

NFSYSTEM3
- Proliant DL380
- Windows 2000 Advanced Server®
- NF Correlation Engine (event correlation, forwarding, and notification)
- NF Universal Agent (events that are not handled by included agents)

NFSYSTEM4
- Proliant DL360
- Windows 2000 Advanced Server®
- NF Windows Event Agent (Windows events)

NFSYSTEM5
- Proliant DL360
- RedHat Linux AS v.2.1
- NF Agent for Unix (Linux, AIX, and other non-HP Unix events)
- NF Universal Agent (events that are not handled by included agents)

All devices capable of forwarding messages in syslog format were configured to send copies of all INFO level or higher messages to the appropriate logging server according to the phased implementation schedule. The Windows© agent uses Windows Management Instrumentation (WMI) to collect events, also according to the schedule.

There are two main screens that GHI uses to monitor events and identify incidents in the NF system – the first screen contains the Device Status and the Event Console windows. The Device Status windows (right) lists all configured devices and the number of events received since the view was last cleared; the Event Console (left) has been configured to only display events that are classified as critical.

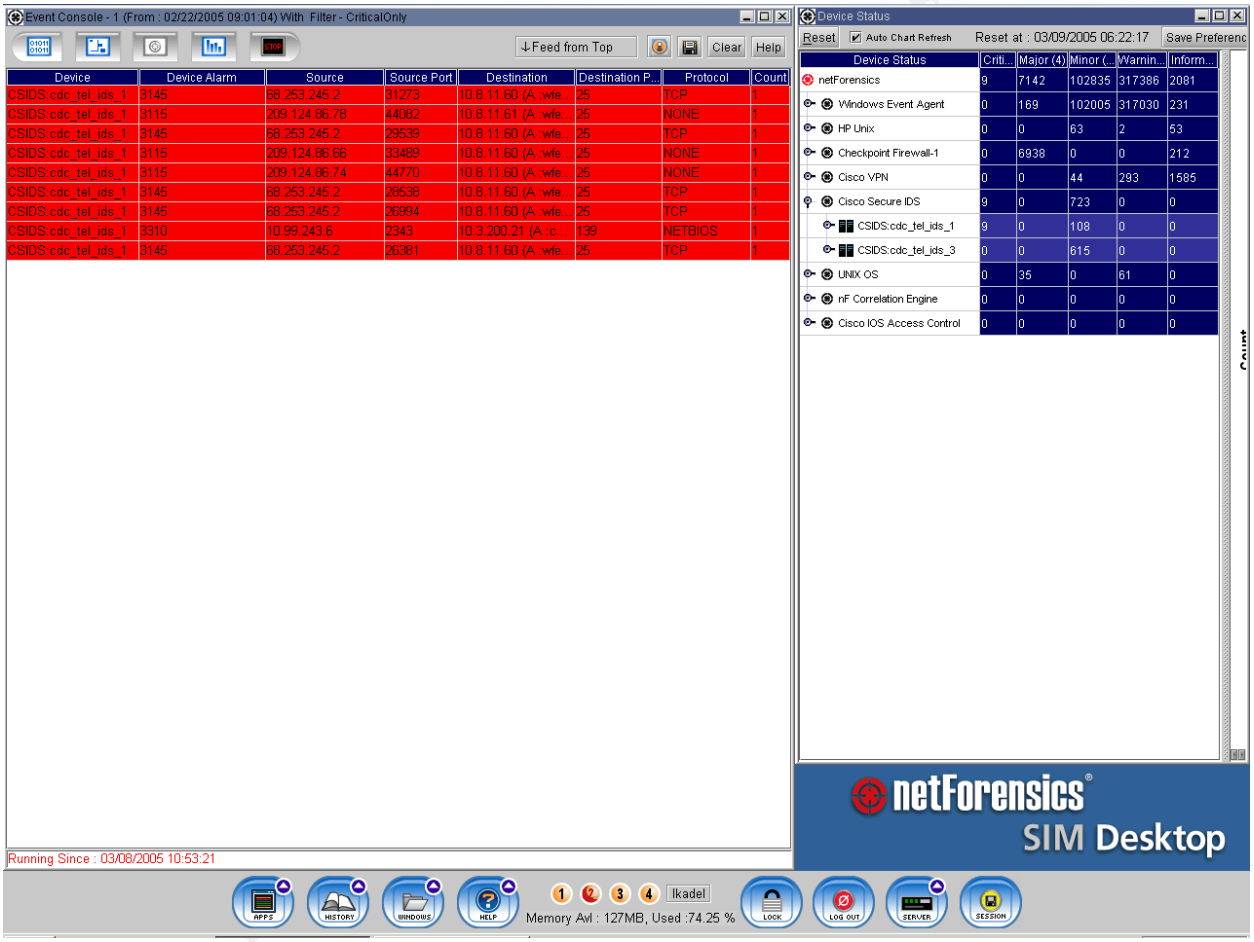

Figure 2 - NF Event Window

The Event Console can be configured to filter and display events by severity, device, device type, source and/or destination address, and more; multiple views have been created for use in different situations. From the Event Console, double-clicking on an event will display the event detail, including source and destination addresses, source and destination ports, the device alarm name, the corresponding NF alarm, the severity rating, and the entire text of the original event or message. This screen is used mainly to identify situations where there is ongoing activity – a series of vulnerability scans

from a single source against a large range of target IP addresses, for example.

The second screen, and the one used more often, is the NF Dashboard view. The Dashboard is a preconfigured display of nine windows, each displaying a different aspect of the same data. Each of these views is briefly described below:

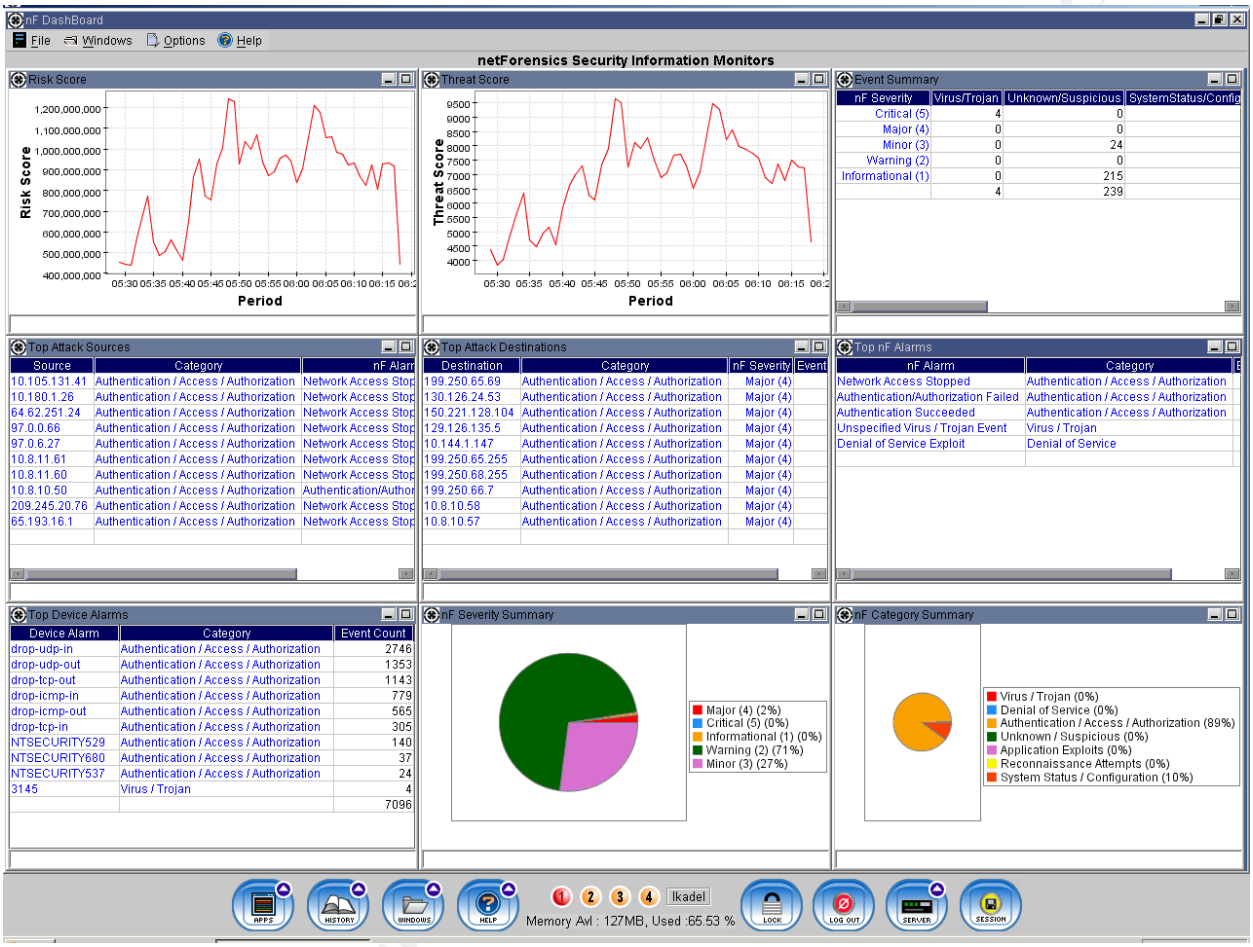

Figure 3 - NF Dashboard

- Top-Left: Risk Score, calculated from the values assigned to Assets within the system (more on Assets later)
- Top-Center: Threat Score, also calculated from Asset values
- Top-Right: Event summary by type and severity
- Left-Center: Attack sources
- Center: Attack destinations
- Right-Center: NF Alarms
- Bottom-Left: Device Alarms
- Bottom-Center: Event severity summary
- Bottom-Right: Event category summary

Because events are tracked by the system in real time, these displays can quickly

identify abnormal conditions which may require further investigation, or which may indicate an attack in progress. It has taken GHI the better part of a year to determine the best way to configure and use these two displays for maximum effectiveness, however many previously unknown vulnerabilities and threats were discovered and the issues resolved within the first 4-6 weeks of deployment.

To calculate Risk and Threat scores, the system was first populated with names and addresses of all assets active in the environment. These include all servers, routers, switches, workstations, and other IP-based devices in the GHI environment. In the NF system, a Device is anything that sends events to be aggregated and analyzed, an Asset is anything that can be either a source or a destination for an event; all Devices are also Assets, not all Assets are Devices.

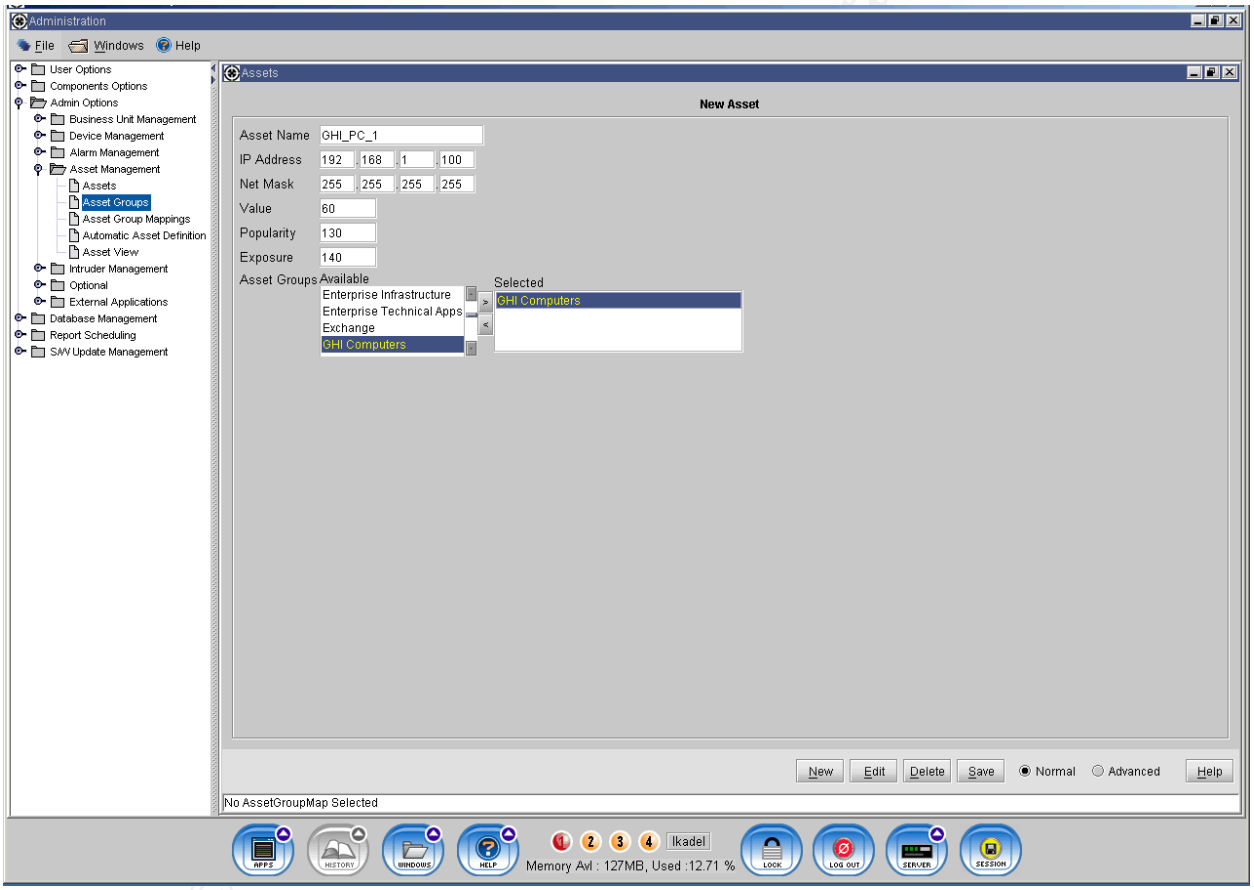

Figure 4 - Asset Configuration

Once identified in the system, each asset is assigned three values (NetForensics Admin Guide, 2003):

- **Value -** A numeric weight indicating the value of the asset, based on its cost or importance to the enterprise
- **Popularity -** A numeric weight indicating the number of events that are generated for a particular asset, similar to the number of hits on a WEB site
- **Exposure -** A numeric weight based on the functionality of the asset, assuming that a higher number of services provided by the asset equals a greater chance of corruption, loss, or unavailability

The Threat Score is calculated using the formula (NetForensics Reports Guide, 2003):

$$\text{Threat Score} = \sum_{n=1}^{5} [\text{Event Count}] * (2^{(n-1)} - 1)$$

where 'n' equals Event Severity

Having determined the Threat Score, it is then used to calculate the Risk Score in the following formula (NetForensics Reports Guide, 2003):

Risk Score = Threat Score * Popularity * System Value * Exposure

It is these two scores that are used to generate the charts in the two upper-left windows of the NF Dashboard display. As long as these graphs remain within the established baseline, no immediate action is required; it is when a spike in the value occurs that the team must investigate further, and decide what action, if any, is required. Any of the windows in the Dashboard view may be custom configured to display charts, graphs, or tables, and can have multiple filters and queries applied to create unique views into the GHI environment.

All of the above features display events in real time, but there is no ability to track related events from multiple devices. Take a simple attack scenario:

- The first step often involves password guessing; possibly a brute-force attack
- Once into the system, the attacker will most likely try for elevated privileges
- With the newly gained rights, the actual attack begins with a system exploit

Because the events generated by these actions will likely be caught by different devices at different times, it is very difficult to determine that an attack is taking place until it is too late. The rules-based correlation engine solves this by tracking patterns of events. The NF system contains twelve built-in rules; GHI has created two custom rules that better meet the unique needs of their network. The screenshots and explanations on the next pages demonstrate how this is done, and to what effect.

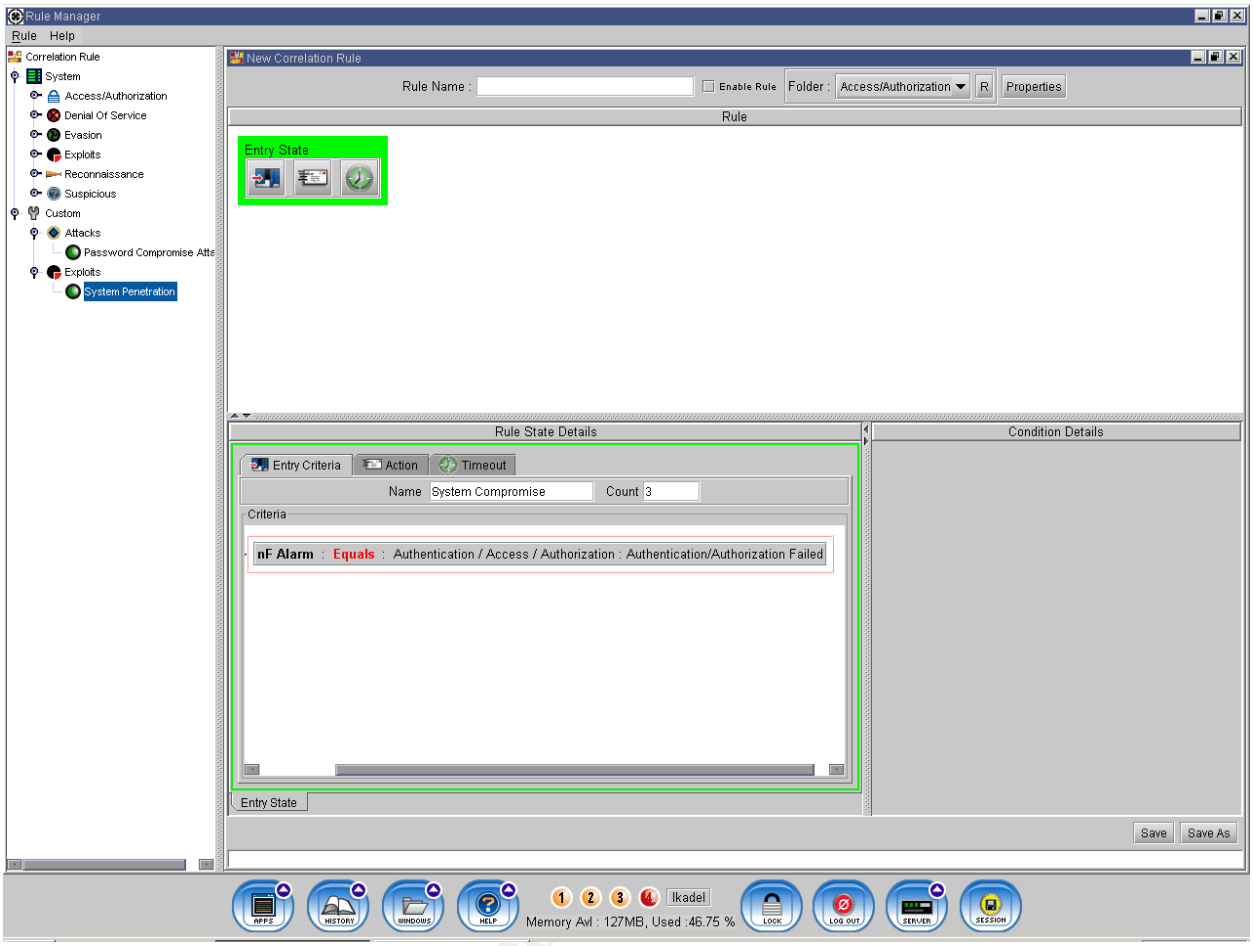

Figure 5 - RBCE Entry State

The correlation engine uses state-based rules to develop a pattern of device events upon which to generate a correlated event. In the above example, the Entry State will be triggered when the system records three authentication failures in 30 seconds. No actual action is taken at this point; nothing further is triggered until the next series of events triggers the Transient State – otherwise the rule times out and the process begins again. In the screenshots on the next page, the entry criteria tab of the Transient State portion of the rule is configured, and an associated action is taken – in this case, it triggers a correlated NF Alarm of type Planned Attack with the message, “Possible Unauthorized Access”. Note that this event will only be triggered if there are 3 or more authentication failures in 30 seconds, followed by a successful login within another 30 seconds. The correlated event is assigned a severity of 4 (Major), and it can be set to send a notification via email and/or SNMP to one or more people.

The next page shows the Entry Criteria and Action tabs for the Final State, which is triggered only after the Transient State criteria is satisfied and the first correlated event is triggered. This Final State event, based on a known Windows exploit, may trigger another notification, and possibly a custom action such as blocking the intruder, or running a script. These correlation rules are also updated as new devices are added, removed, or changed.

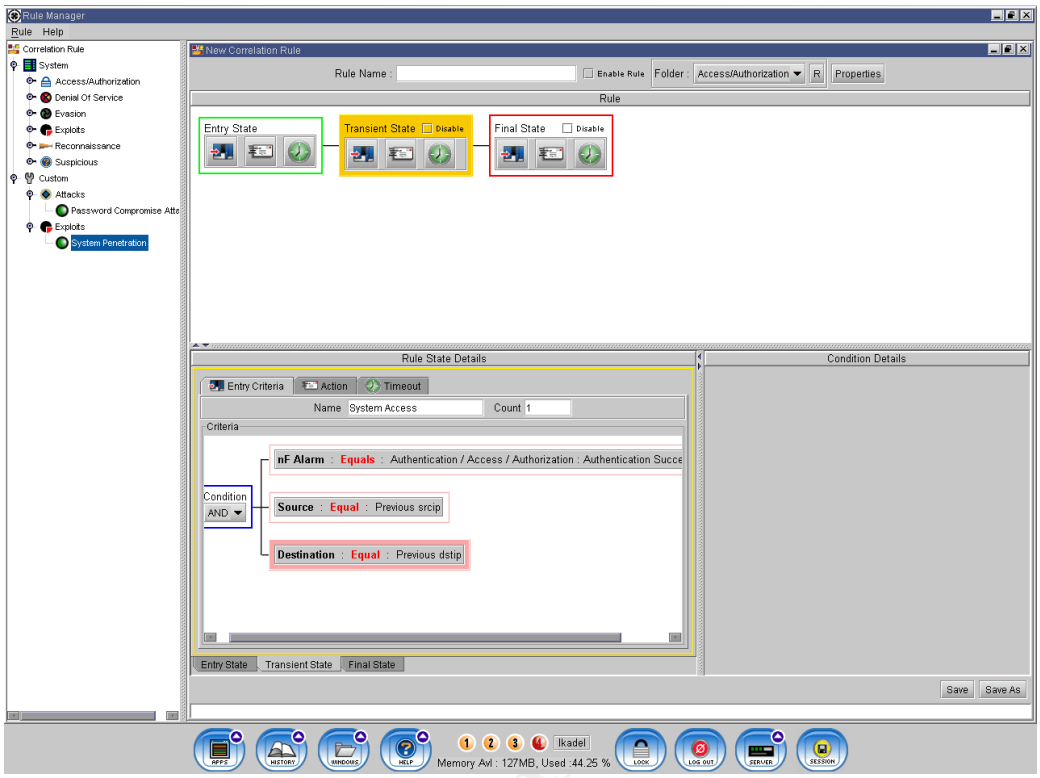

Figure 6 - RBCE Transient State Entry Criteria

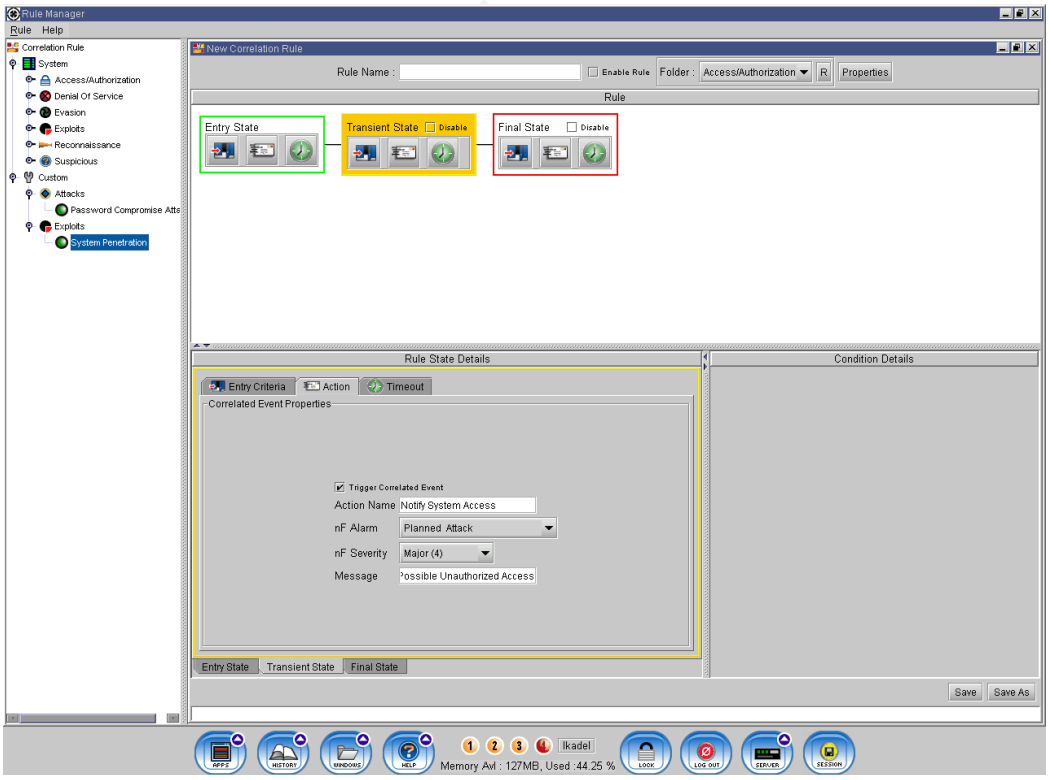

Figure 7 - RBCE Transient State Action

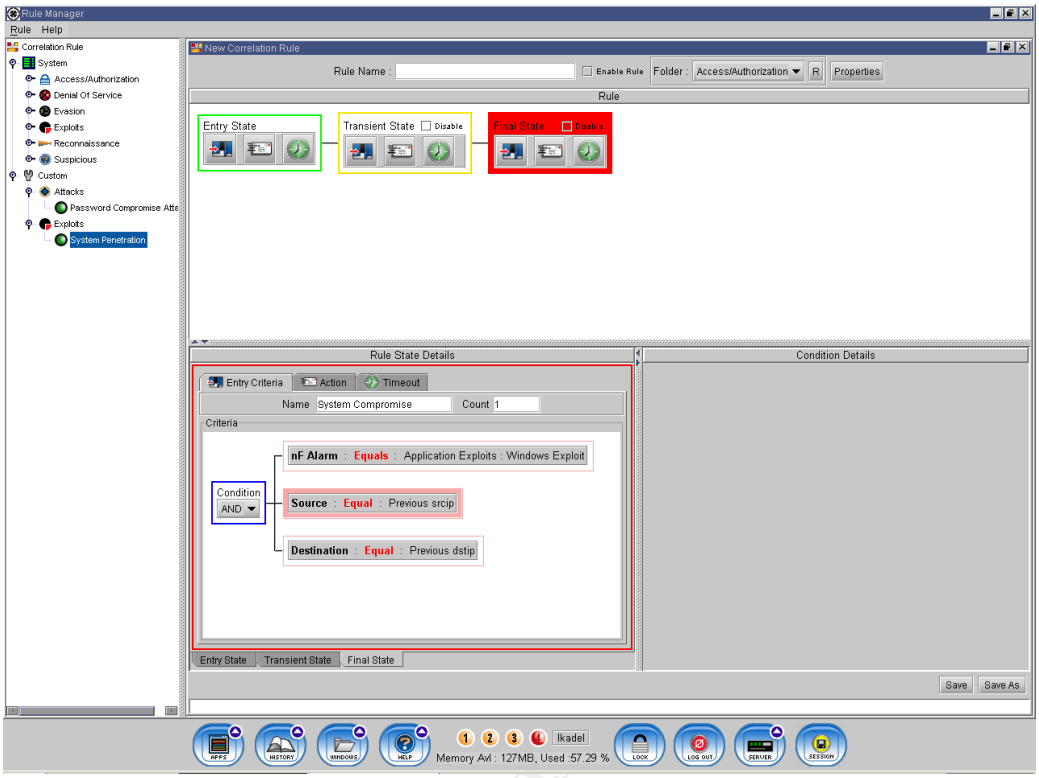

Figure 8 - RBCE Final State Entry Criteria

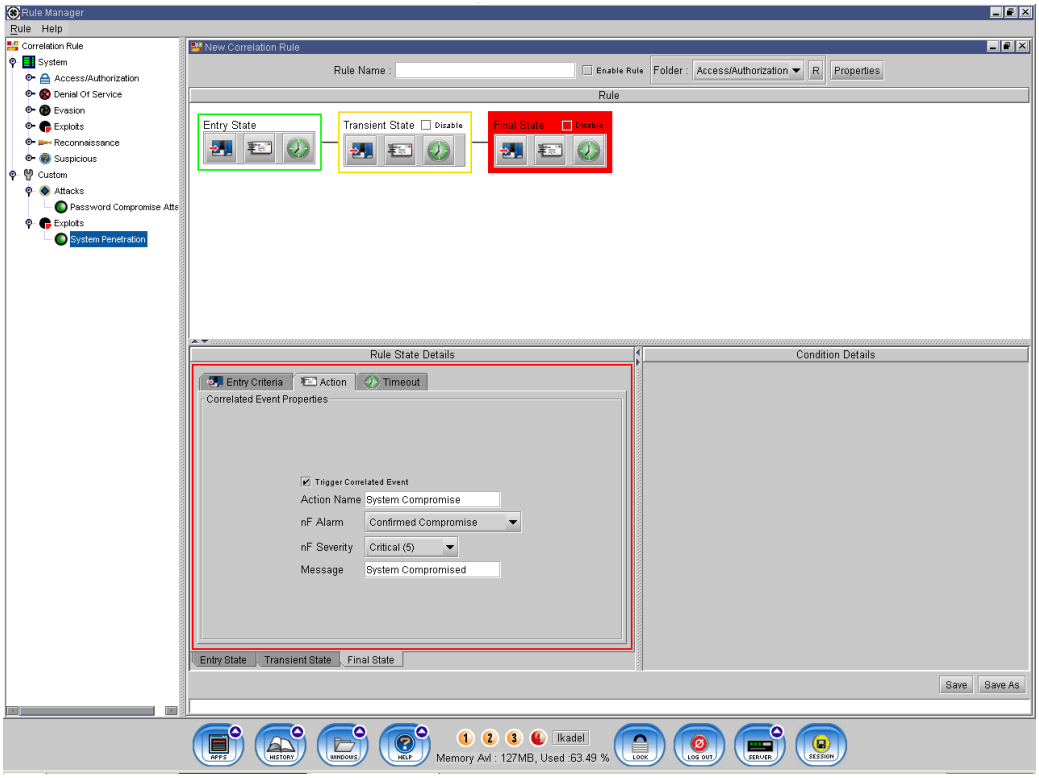

Figure 9 - RBCE Final State Action

A second Final State is then created within the same rule to catch a different avenue of attack. As shown below, this is also triggered only after the Entry and Transient states have been satisfied; in this case, if a Denial of Service is attempted.

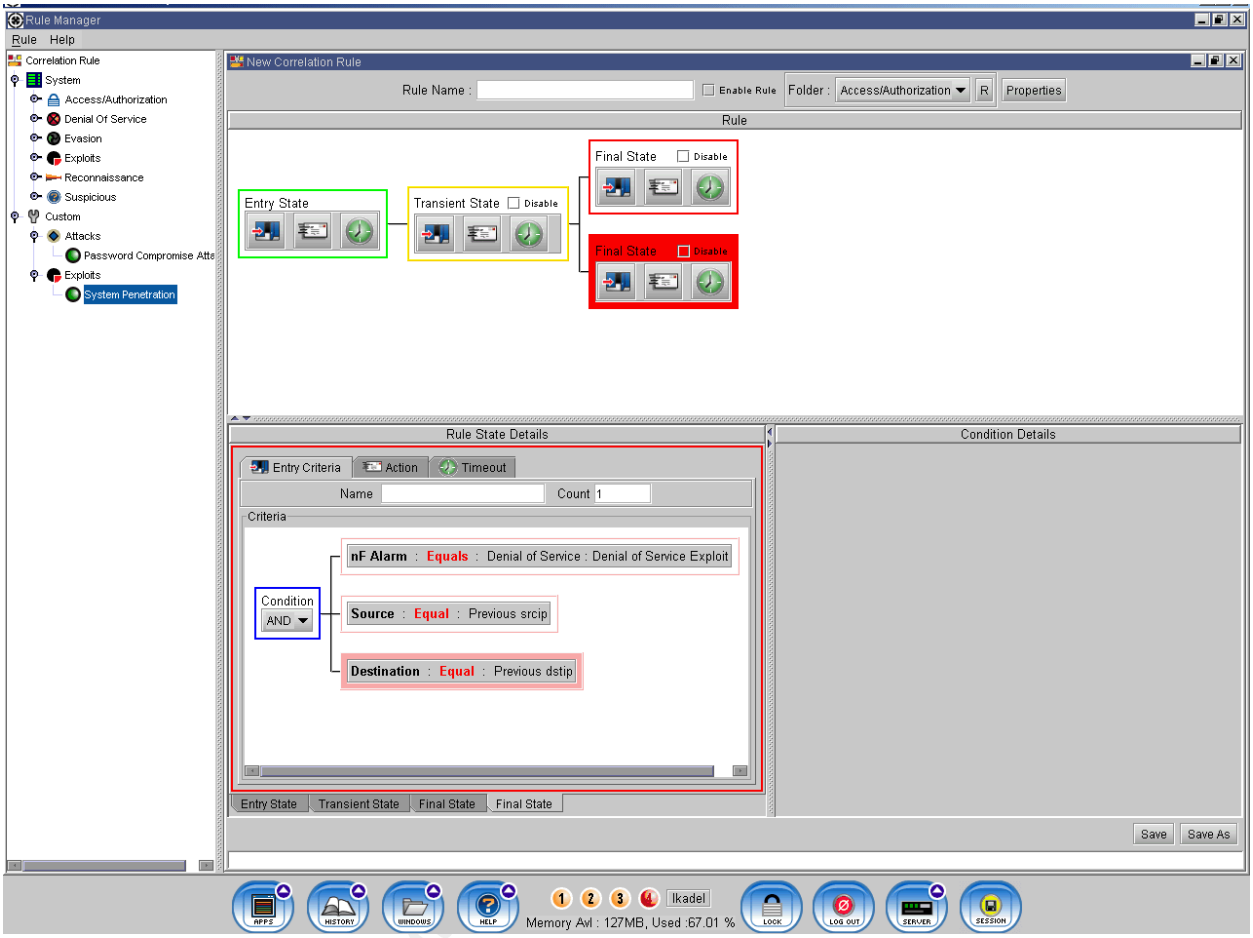

Figure 10 - RBCE Final State 2 Entry Criteria

This combination of monitored devices, assets, real-time events, and rule-based custom events, coupled with the notification system gives GHI a constantly updated view of the security state in the enterprise. When a potential incident is identified, it can be investigated and acted upon immediately, often stopping an attack before it has been fully developed and systems compromised.

All events, both device-originated and correlated, are retained in the active database for a period of 180 days; after that time, events are archived to files stored elsewhere on the system, and kept for another 180 days (one year total). These 'active archives' can be restored into the database at any time for use in forensics or reporting. After one year, the archive files are moved to long-term storage kept on AIT3 backup tapes. The system has only been active for a little over a year, so details of storage requirements have not been fully established. To date, no need has arisen to require the restoration of archived data.

There are several predefined report templates in the NetForensics system – these

are used by the Information Security team to generate reports for management and senior leadership, and for day-to-day auditing. All report generation windows are very similar, and the process to create a report is the same for all.

Similar to the NF Dashboard, the Risk and Threat Assessment reports display the risk and/or threat posture for a given timeframe. Other reports include Event Queries, IP Activity Queries, Most Frequent Events, and more; all reports can be displayed as either a table or graph, and graphs can be vertical, horizontal, or pie chart. The example report below shows a Threat Assessment based on all reconnaissance attempts from the current week.

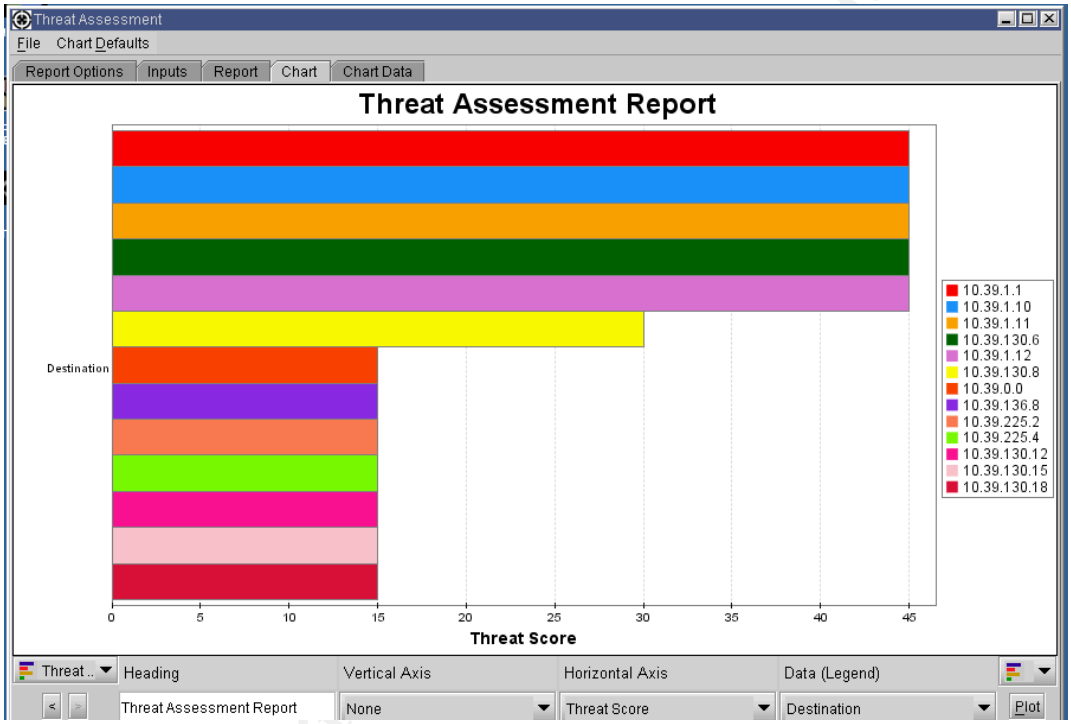

Figure 11 - Threat Assessment Report

These reports are run monthly, exported to HTML files, and stored in a secure location on the network; they are then consolidated into a standard reporting format (using Microsoft Word®), along with other reports from other systems in the enterprise, and sent to management as a quarterly status report. The reports are manually archived to ZIP files once a year; the archives are kept on AIT3 tape for a period of 7 years.

In addition to the built-in reporting capabilities of the NetForensics system, GHI is working to create custom reports using the Crystal Reports application. This process has presented some unique challenges, many of which have yet to be resolved, and the process is far from complete. It is the goal to be able to generate custom reports and graphs based on any combination of data available within the NetForensics database.

Summary

HIPAA regulations, in §164.308(a)(1)(ii)(D), require covered entities to review log files and other system activity records on a regular basis. While the implementation details of this requirement are left up to the individual company, the intent is clear: to identify risks and threats to the environment, and audit for violations of established policies and procedures.

To comply with the rule, GHI first developed and implemented a policy to address the issue of activity and event log review. Within this policy, the various types of logs are defined; responsibility for review of the logs is assigned the use of an automated system is authorized; and data retention periods are mandated. Finally, penalties for violation and/or non-compliance with the policy are declared.

In the GHI environment, it was decided to use the NetForensics SIM Suite product to collect and analyze events in real time, and notify the appropriate staff if/when an incident occurs. The system was purchased and installed, and an initial configuration deployed. By using the data garnered from the initial deployment, a baseline was created; subsequent phases of the deployment refined this baseline.

The NetForensics suite confers compliance with the HIPAA specification by:

- Collecting events from systems configured to generate logs
- Normalizing those events into a standard format
- Aggregating the events into a manageable environment
- Correlating the events to identify suspicious activity or incidents
- Notifying in real time when incidents occur
- Reporting on events, activities, trends, risks, and threats in/to the environment

While the functionality is not complete (there are some devices and applications which do not generate events or activity logs), it has been determined that these exceptions do not pose a substantial risk, and that they can be effectively monitored using other means. Therefore, GHI has concluded that this implementation satisfies the requirement under the law, and even exceeds it in many respects.

References

Health Insurance Reform: Security Standards; Final Rule, 45 C.F.R. Parts 160, 162, and 164 (2003).

SANS Institute. (2004). *HIPAA security implementation* (version 2.0). Bethesda, MD: Author

NetForensics, Inc. (2003). *NetForensics Administration Guide* (version 3.1.1) Edison, NJ: Author

NetForensics, Inc. (2003). *NetForensics Reports Guide* (version 3.1.1) Edison, NJ: Author