Paper 190-30

A SAS® Framework for Network Security Intelligence

Michael Protz, SAS Institute Inc., Denver, CO
Joseph G. Zilka, SAS Institute Inc., Cleveland, OH
Jeff Mudd, SAS Institute Inc., Cary, NC

ABSTRACT

An organization's computer network is constantly under attack, facing threats such as probes, worms, viruses, and host-system root access attempts. The threats originate outside the network from sources worldwide, as well as inside network boundaries. To counter threats, most organizations employ various types of point products, such as firewalls, routers, and intrusion detection sensors to alert them of hostile activity. Some organizations employ event correlation engines to manage these point products and provide a single point of control for reacting to hostile activity. The end goal is obvious—provide an impenetrable Computer Network Defense (CND) that prevents interruption of network services and protects an organization's critical data.

This paper presents a SAS-based framework that enhances the real-time data provided by point products and event correlation engines, provides a consolidated view of CND data in a historical context, and most important, improves an organization's understanding of the health of its computer networks. This paper details the key components of the framework: a centralized repository to integrate and manage all CND-related data, scalable storage architecture to efficiently handle extremely large volumes of data, and role-based interfaces to deliver critical CND information to the right consumers at the right time. This paper also explains the potential use of advanced analytics for the detection of hostile activity undetectable by the signature-based point products. The last section contains a case study of a real-world implementation of this framework at the Navy Computer Incident Response Team in Norfolk, Virginia.

INTRODUCTION

Computer systems and the networks they communicate through are an integral part of every organization. These systems are critical assets that must be vigilantly protected from the ever-growing threats designed to disrupt and corrupt an organization's cyber infrastructure. The following sections present the numerous threats to network security, the steps that most organizations take to protect themselves, and a framework based on the core competencies of the SAS Intelligence Platform to complement and enhance an organization's threat response and protection activities. This paper is targeted to those with a background in or an understanding of network security concepts.

THREATS TO NETWORK SECURITY

The number of threats to an organization's network security is at an all-time high, and new threats are being identified every day. Because the Internet consists of hosts that are distributed across the globe, threats can originate from any source with a connection to the Internet, or without a connection. If an organization's network is compromised, it can cost millions of dollars, and potentially endanger lives.

Threats can exploit a wide range of vulnerabilities. Some threats are hardware- and software-related, and other threats attack human naivety and trustfulness. Threats can be as harmless as a scan of open ports, while others are logic bombs designed to corrupt data and mission-critical applications. Some threats are actively trying to alter a system or data; others exploit improper configurations, low-quality code, or poor information security policies. Threats can be created by a novice hacker on the other side of the globe, or by the employee in the next cubicle.

THREATS OUTSIDE THE NETWORK

A typical organization's network can be exposed to a number of threats that originate outside the network. Simple threats include scans and floods. A host's ports can be scanned for network services that listen on "well-known" ports to determine which network services are running. After these ports are discovered, a potential hacker can probe these ports for suspected vulnerabilities associated with the network service. Networks can be exposed to floods such as ping floods or SYN floods. Floods send more requests to a specific host than the host is configured to handle and can cause the host to crash or deny services to legitimate requesting clients. Floods are classified as Denial of Service (DOS) attacks. Buffer overflows attempt to send more data than a temporary storage space can hold. Once the capacity of the buffer is exceeded, critical data in adjacent buffers can be overwritten, which creates unpredictable results for the host.

More complex threats include Trojan horses, viruses, and worms that execute malicious code. Trojan horses are programs that appear useful to the host or user, but they have hidden functionality known only to the creator. The Trojan horse can execute malicious code or exploit other system weaknesses, possibly circumventing existing host security. It can create a new backdoor into a compromised system. A well-known example of a Trojan horse is Back Orifice. Hackers search for hosts infected with Back Orifice and commandeer the host with an administration tool to conduct attacks such as a DOS and a Distributed Denial of Service (DDOS).

Viruses are similar to Trojan horses in that they destroy or manipulate data, but viruses also attempt to reproduce and spread themselves. Viruses rely on host programs to execute them. In 1999, the well-known Melissa virus was a macro that e-mailed itself to the first 50 entries in the victim's Microsoft Outlook address book. A boot sector virus appeared in 1991. This virus deleted files on machines that contracted the virus after booting from an infected floppy disk. It deleted files on March 6th, Michelangelo's birthday, so it was aptly named the Michelangelo virus.

Worms exploit security holes in computer networks. After a worm infects a host, it scans the network for additional vulnerable hosts and replicates itself on a new host, from which it repeats the process of scanning that host's network for vulnerable hosts. The Code Red worm garnered attention in 2001 as it replicated itself very quickly. The Code Red worm was designed to replicate itself the first 20 days of the month, deface Web pages on infected servers, and conduct a DDOS on the White House Web server the remaining days of the month.

A less technical, but damaging threat, is social engineering. Social engineering involves preying on human trustfulness. Some Trojan horses and viruses rely on social engineering for their perpetuation. The ILOVEYOU virus, which is actually a worm, entices its victims to open an attachment that, when executed, proliferates the virus. Social engineering can be used to obtain usernames and passwords, which, in turn, allows an unauthorized user access to a server or confidential files.

THREATS INSIDE THE NETWORK

Threats can originate from within an organization. Sabotage, espionage, misuse, social engineering, and software design and configuration pose significant threats. Internal threats can originate from employees who have legitimate access to a system. A disgruntled employee at a manufacturing facility left a "time bomb" that, at a specified time, deleted mission-critical programs for running a production line and caused irreparable loss of time and money. Unhappy employees can steal information or trade secrets for another organization, resulting in lost revenue or compromised national security. Although data might not be destroyed, the results can be damaging.

Misuse of company networks might not have a direct impact on security, but it can reduce performance and throughput, along with unnecessarily exposing the network to other threats. For example, employees who read personal e-mail from service providers on their PCs at work can expose the PCs and the entire network to destructive viruses and worms.

Social engineering can be used by employees to gain access to protected or confidential information such as salaries and personnel information. This can occur if one employee convinces another employee to provide passwords to systems the requesting employee is not authorized to use.

Poor software design and configuration can be a threat to security. Low-quality software with security flaws is often made public and can be exploited rapidly and remain vulnerable until a patch is applied. An organization relies on the system administrator's knowledge of the vulnerability of the software, the availability of a patch, which machines are affected, and the timely application of a patch to repair security holes. Patching might interfere with an existing mission-critical application and break it, which could prohibit application of the patch. These same applications might require specific services or ranges of ports to remain open, which contributes to the host's vulnerabilities. The demands of keeping up with security flaws and application requirements can lead to overworked system administrators who involuntarily neglect other systems' and applications' security configurations, which can lead to an even greater degradation of network security.

IMPACT OF NETWORK SECURITY THREATS

Generally, network security has four different aspects: confidentiality, integrity, availability, and accountability. The threats previously described can affect one or all of these aspects and have a negative effect on an organization, consuming dollars from budgets or bottom lines.

Confidentiality of data means that only parties authorized to access information can access it. Data that has lost its confidentiality is compromised. A list of subscribers' credit card numbers or secret military information has a level of confidentiality and is often targeted for theft or espionage.

Data integrity means that data has not been modified by anyone other than authorized parties. Trojan horses, viruses, and worms attack data integrity. Modified file systems or a changed course grade are examples of data whose integrity has been corrupted.

Availability implies that information and resources are accessible to the intended users, on request. A DOS overloads a Web server to the extent that it cannot field legitimate requests, which affects the availability of the service.

Accountability requires that users of information and resources be identified and activities be logged, making actions and access patterns traceable. Users are then accountable for their actions.

How does an organization suffer if any of these four aspects fail? Data corruption or loss can render systems unusable and have a huge impact if an effective backup strategy is not properly implemented and tested. Loss of the physical machine can be miniscule, compared to the loss of information that was stored on that machine. Corrupted data affects integrity and can affect availability. Stolen credit card information from a bank and Social Security information stolen from a human resources system are examples that affect not only the organization, but also the individuals who entrusted their privacy to the organization. Stolen data impacts an organization's confidentiality and accountability, and there can be legal liability. Individuals who have had their sensitive information compromised might pursue legal action, not against the unknown hacker, but against the organization that failed to ensure confidentiality. Loss of productivity by systems or unavailable information can cause major problems. A manufacturing company lost millions of dollars because a production line shut down after an intentional deletion of mission-critical programs. Loss of revenue and opportunities are inevitable if a DOS is successful against a popular Web site or a Web site that provides e-commerce as a large portion of its services. The damage to an organization's reputation after a widely publicized hack can be difficult to overcome. In 2001, a number of sensitive Federal Bureau of Investigation (FBI) documents were sent to outside parties as a result of malicious hacker activity, which is an embarrassing incident for any organization, but especially for the FBI.

REAL-TIME COMPUTER NETWORK DEFENSE

In response to the large number of threats facing computer networks, large-scale organizations, including Federal Government Agencies and Fortune 500 companies, have established a comprehensive cyber security defense strategy, which is anchored by the formation of a Security Operations Center (SOC) and a Computer Incident Response Team (CIRT). The SOC, which is often a subset organization within the Network Operations Center (NOC), provides front-line defense against cyber threats and is the nucleus of all information and Internet security operations. The SOC provides continuous protection, detection, and response capabilities against threats, remotely exploitable vulnerabilities, and real-time incidents on the networks.

The CIRT provides real-time network monitoring in support of the SOC, but the CIRT is primarily responsible for the incident handling and forensic activities of a potentially malicious attack. *Incident handling* refers to the response of a person or organization to an attack. A *security incident* is an adverse event in an information system or network or the threat of such an event. Incidents can include, but are not limited to, unauthorized access, malicious code, network probes, and DOS attacks.

In addition, organizations have architected a layered approach to CND. Also known as "defense in depth," a layered CND architecture segments the network into categories such as automated applications, enclaves, outsourced IT processes, and major system interconnections. Furthermore, the network gets broken down into subnets, such as Demilitarized Zones (DMZs), with multiple layers of screening routers, perimeter firewalls, internal firewalls, Virtual Private Networks (VPNs), anti-virus protection at the desktop or server-level, and Intrusion Detection System (IDS) sensors placed strategically throughout the network.

In many cases, a layered CND strategy introduces a variety of point products (Figure 1) at each layer such as routers, firewalls, IDS sensors, anti-virus protection, and vulnerability scanners. For each type of point product, organizations might deploy solutions from multiple vendors. This strategy creates a challenge when trying to manage a cohesive strategy for CND. Every point product produces data (logs/events), and most point products come with their own management console.

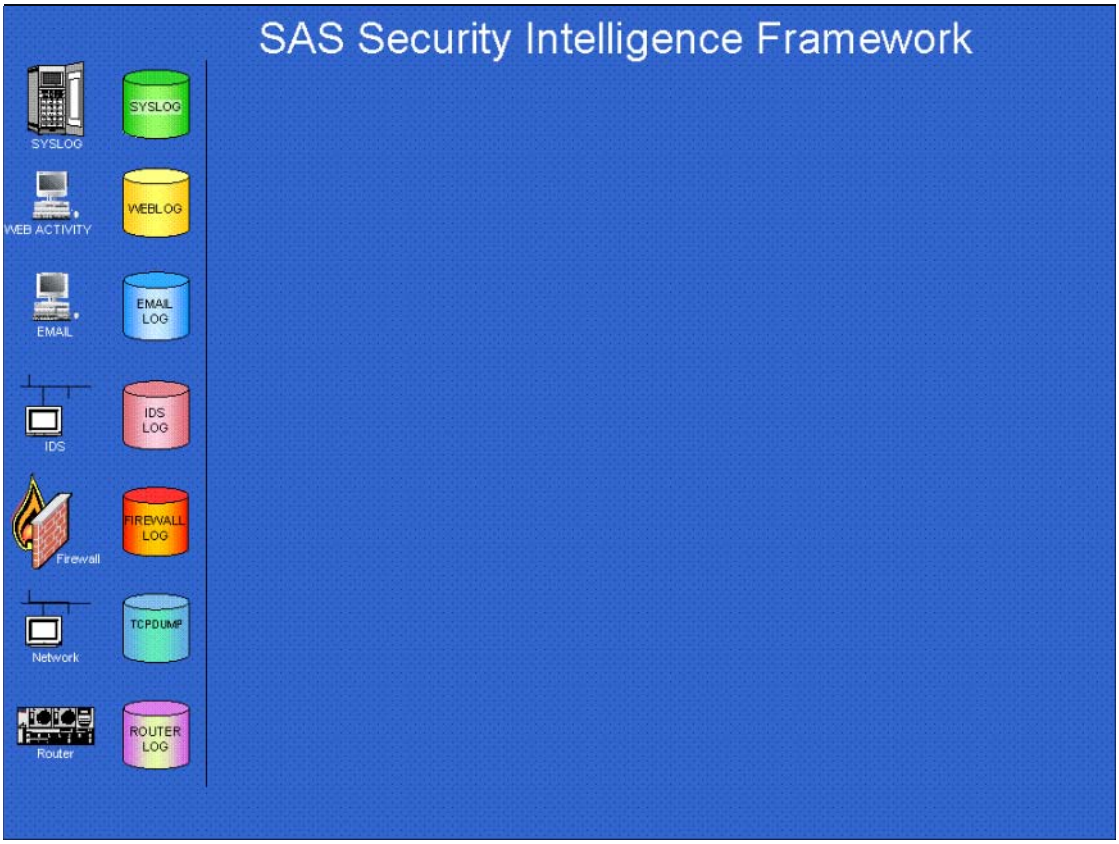

Figure 1. Point Products

Security devices and applications produce numerous events, spanning informative notifications through solution faults and attack alerts. Security-oriented events come from system and application sources. Useful security information can be found in log files. Managing the breadth and volume of security events, along with maintaining sufficient information for forensic and auditing purposes, is a major challenge. Correlating this wealth of security data to provide useful information is essential for timely and effective management. As a result, organizations have turned to real-time event correlation engines (Figure 2) or Security Information Managers (SIMs), to help consolidate the information flow from the vast number of point products and data sources in the environment.

To make intrusion detection suitable for the enterprise, dispersed sensors report events to a central console. An event correlation tool or SIM facilitates this task by placing agents on the various IDS sensors, firewalls, and routers, and consolidating the information to a single management console. Correlation rules identify when activity related to a specific event or attack is detected across multiple devices. As a result, correlation rules reduce false positives. SIMs capture feeds from vulnerability scanners and anti-virus tools to assess the weakness of an enterprise or a specific IT asset against an identified threat.

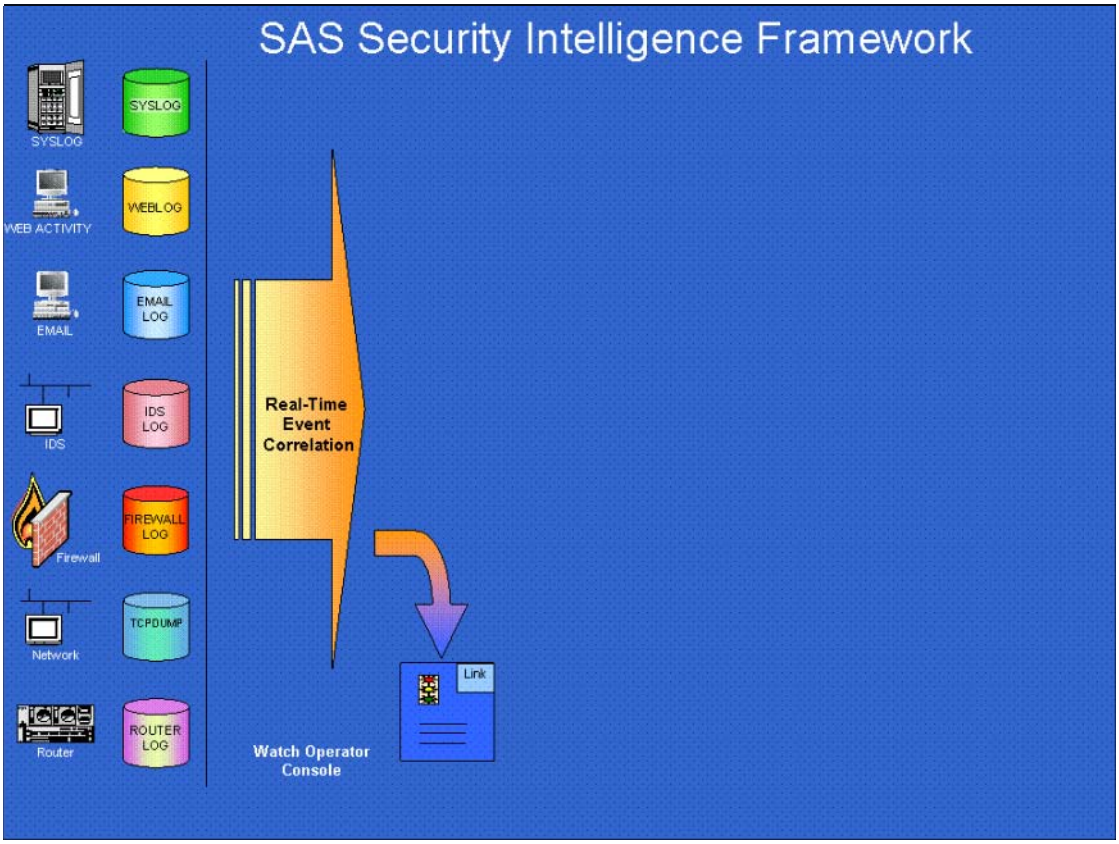

Figure 2. Point Products with Real-Time Event Correlation Engine

EXPANDING THE CONCEPT OF CND

Point products, real-time event correlation engines, and SIMs are the foundation of any organization's layered CND, but they are not the only pieces in the puzzle. Logs produced by point products contain all types of information, depending on the type of sensor (firewall, router, IDS, etc.) and the sensor vendor. Although real-time event correlation engines and SIMs are designed to overcome the variety of point product sensors and perform their functions extremely well, they are limited in the amount of data they can store and are not effective repositories of historical sensor data. Additionally, event-correlation engine and SIM data provided through the application console are filtered to ensure that only the most critical information makes it to the user to maximize response time.

More important, sensor-based data is only one facet of a fully integrated CND framework. Other referential data sources provide complementary information and are integral to building the complete picture of an organization's network health. These referential data sources provide information such as:

- incidents stored in systems such as REMEDY or other trouble ticket-tracking applications
- vulnerability assessment and remediation
- network topography or topology
- asset criticality

PRIMARY INFORMATION CONSUMERS

To achieve a comprehensive CND framework for security intelligence, it is important to understand the four primary users of CND data:

- watch operators
- incident handlers
- incident analysts
- executive-level decision makers

Each user has a specific need for CND data to perform daily tasks and requires an interface targeted to a specific skill set and job function.

Watch operators are the frontline human sentries of CND. They are responsible for monitoring point products and SIMs and responding appropriately to alerts as they are displayed on the real-time console. To effectively respond to these alerts, ready access to all facets of CND data, real-time and historical, through a user-friendly interface is critical.

Incident handlers gather information associated with suspected malicious activity. Whether the activity is identified by a watch operator or information is submitted by an organization, incident handlers make decisions about how to resolve the incident. An organization uses trouble ticket-tracking and resolution software to manage the incident information and workflow. Some of this functionality is found in the SIM software. As with watch operators, incident handlers benefit from access to historical data because it assists them in making appropriate decisions regarding an incident.

A single incident or pattern of incidents requires in-depth investigation to understand the full life cycle and how effectively the various layers of the CND infrastructure reacted or, more important, failed to react. *Incident analysts* are the forensic investigators within an organization's CND infrastructure who are responsible for this activity. Of the primary CND data users, incident analysts are the heaviest users of historical data. They are power users that require heavy-duty toolsets to perform various levels of detailed analyses. Incident analyst requirements are never static, so data and tool flexibility is extremely important.

Executive-level decision makers are concerned with having the most current information about the performance of the networks they monitor and, most important, the effectiveness of the CND infrastructure in preventing malicious activity against those networks. They are not concerned with details, but they want the option to review the details when necessary. Frequently, they want a metrics-based, situational awareness view into their CND operations that enables them to respond to customer inquiries.

INTRODUCING THE SAS FRAMEWORK

The primary CND data users require a framework that provides a fully integrated, historical view of real-time event data and complementary referential data. They require targeted tools and capabilities to perform different levels of flexible reporting and analysis. This includes basic "Top 10" and baseline analyses, as well as advanced statistical analysis of data over longer periods of time that looks for activity that can go undetected by signature-based IDS sensors and correlation engines.

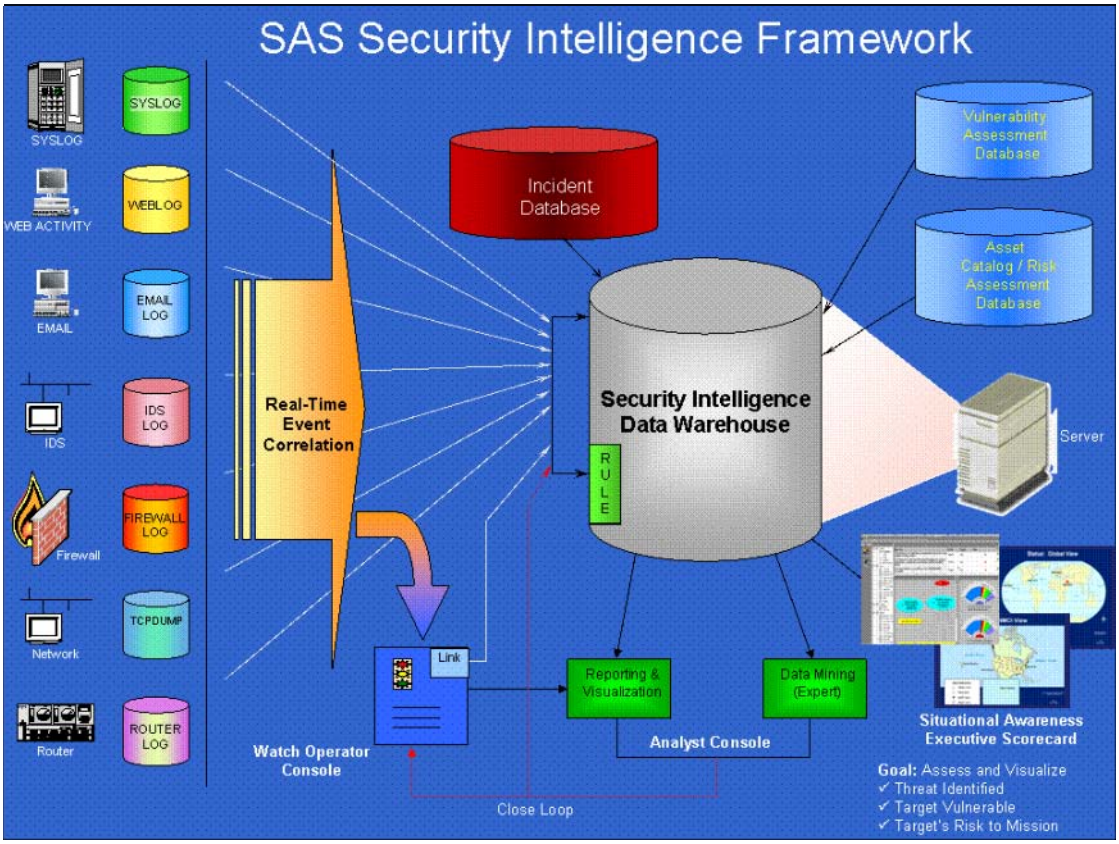

Figure 3. SAS Framework for Network Security Intelligence

The SAS Intelligence Platform has all the building blocks for this framework (Figure 3): access, manipulation, and management of CND data inherent in the components of the SAS Foundation and SAS ETL Server; storage of integrated CND data in structures designed specifically for reporting and analysis through SAS Intelligence Storage; and interfaces targeted for the primary CND data users through the SAS Enterprise BI Server.

SAS FRAMEWORK COMPONENTS

There are several types of data that contribute to the effective operation of CND and are supported by the SAS Framework for Network Security Intelligence. Collectors gather information from various points on the network, while hosts log activity related to Web, e-mail, and authentication. Vulnerability scanners inspect the network and hosts for known vulnerabilities and record this information for action and compliance. Records of patches applied and vulnerabilities remediated help assess future vulnerabilities. Information about an organization's assets, the criticality of the assets to operations, and the sensitivity of data can create an asset database that provides the necessary data for assessing the impact of a compromised or corrupted host. A compilation of past incidents generates a historical context that provides insight into current incidents.

COLLECTOR DATA

Collectors are devices that gather relevant and often large volumes of information about activity from a network or host. *Routers*, specifically border routers, are the connection point of an organization's network to the outside world, also known as the Internet. Because routers pass packets into and out of the internal network, they record valuable statistics about network activity, monitor for specific attacks, and block IP addresses from coming in or going out. Other programs that "sniff" network traffic, such as tcpdump, create immense amounts of data and provide a detailed look into what is happening on the network.

Network Intrusion Detection System (NIDS) sensors listen to network traffic and are designed to protect a LAN. *Host-based Intrusion Detection System (HIDS)* sensors monitor traffic and logs to protect the host on which they are installed. Both trigger an alert when nefarious activity occurs, and they can drop questionable packets. The sensors accomplish this by analyzing the packets and comparing them to known attack signatures. They employ anomaly-

detection and signature-detection techniques to perform their functions. HIDS sensors can also monitor file systems, logs, and user-specific activity. *Event logs* gather information relevant to applications, security, and system events used by HIDS sensors. These event logs can be used as detail source data for the host on which they reside. Alarms can be raised, sent to a console, and used by people who monitor these hosts or networks.

Firewalls control the network traffic by employing packet-filtering, proxy service, or stateful inspection. Traffic that fails to meet the rules specified by the firewall can be dropped. Firewall information is logged and provides another useful vantage point for monitoring network activity.

Web logs provide basic information about the requests sent to the Web server. This information can help determine who is on the Web site, where they came from, and what they are doing. *E-mail server logs* maintain useful records of e-mail activity in an organization, help verify existence of viruses and worms, and identify infected hosts.

VULNERABILITY ASSESSMENT DATA

Some of the most useful tools in proactive network defense are vulnerability assessment scanners. Scanners can be configured to scan a network or set of networks for known vulnerabilities. Information produced during the scan includes the identity of the scanning host, the time range of the scan, and the number of hosts scanned. Scanners search hosts for existing services or user IDs that might have an associated vulnerability. Information about the host, for example, the version of the operating system and patch level, is recorded along with vulnerability attributes such as relative severity, full descriptions, and recommended fixes. By itself, the data helps determine whether systems have the current recommended patch levels and can prevent possible configurations inconsistent with the organization's security policy. This data, coupled with data regarding vulnerability compliance, provides information about relative strength against known exploits. If there is no system to address compliance, a repeated scan of the networks indicates whether the vulnerability has been addressed. The severity of a vulnerability, the number of machines vulnerable, and the speed in which they are remediated provide useful metrics on the health of a network or how effectively threats are addressed.

ASSET CATALOG / RISK ASSESSMENT DATA

Some of the most valuable, yet most difficult, data to obtain are the location of assets and how critical an asset is to the organization. Organizations have problems identifying what equipment they own, where it resides, what department it belongs to, or who uses it. By implementing tools to inventory, map, and assess the criticality of assets, an organization can more effectively determine the impact of a vulnerability or an active threat against the business activity or mission of the organization. When used in conjunction with collector data, vulnerability data, and incident data, an organization is provided with statistics on network alarms, security holes, compromised assets, or attempts at compromising an asset.

INCIDENT DATA

Historical incident data can assess incoming threats or targets of a threat. Incident databases record contact information for the target, summaries, current status, comments, and remediation steps. The number of times a target has been compromised, how it was compromised, the suspected culprit, and who is responsible for the machine can help thwart future attacks. When used with collector data, vulnerabilities, and asset criticality, information can be ascertained about the attempt on other hosts from the same source IP, the types of vulnerabilities that existed during past compromises, and what the effect was on the organization.

SAS ETL SERVER AND SAS INTELLIGENCE STORAGE

The multitude of sources involved in the framework and the large data volumes generated by the components require an enterprise-level Extraction, Transformation, and Loading (ETL) tool. In addition, a data storage strategy that can support terabyte data stores and allow efficient and effective access to the information by applications is needed. Capabilities provided by SAS technology packages such as SAS ETL Server and SAS Intelligence Storage provide the key components to perform effective ETL and create a storage strategy for querying, reporting, and analyzing large volumes of data. Key components include SAS Management Console, SAS ETL Studio, and the SAS Scalable Performance Data Server.

SECURITY INTELLIGENCE DATA WAREHOUSE ORGANIZATION AND MAINTENANCE

Data used in the Security Intelligence Data Warehouse (SIDW) originates from collectors via a real-time event correlation engine, incident handling databases, vulnerability assessment databases, asset catalog/risk assessment databases, flat-files, and more. SIDW data is grouped by its function: input sources, incremental, summary, and

application target data stores. This data can be further organized by subject, which maps it to entities in the framework. Examples of framework entities include intruders, events, non-events, assets, networks, collectors, incidents, and vulnerabilities.

The real-time event correlation engine uses agents to monitor activity on the collection devices on which they reside. Collector information is normalized, meaning that severity levels are mapped among the collector types and vendors to provide a simplified view of network activity. This data is the input source for the collector subject area within the SIDW. The information (non-events) and alerts (events) are fed into several data structures. A table of network activity contains the information that is common to all collector types, which gives users a historical view of all network activity. The detail data is split into data structures specific to each collector type because the information available for the different types of collectors varies greatly. The end result is a complete view of all data captured by the real-time event correlation engine for each collector.

An organization might employ multiple incident databases or have different collection methods that capture information in differing formats, particularly when the organization merges with another organization or takes over responsibility for an additional monitoring system. Whether an off-the-shelf trouble ticket system is used and/or a collaboration of home-grown systems is used, incidents need to be consolidated into a structure that provides applications and power users the ability to quickly search incidents. Mapping common information and transforming data into standard formats facilitate efficient searching. Incidents can be tracked separately in more than one system. However, this creates the question of which system is the authoritative source. If used, data elements for performance metrics need to be from the authoritative source or from the next source in the progression. This can be accomplished by a reliable cross-reference table that provides a key or ID number for a corresponding incident in the alternate tracking system.

The vulnerability assessment data provides the SIDW information about security holes and target hosts. Scanner software pushes results into a data table or set of data tables and might provide a basic reporting tool. Scanner data can be incorporated into the SIDW for more advanced reporting capabilities. Detail tables consist of hosts, known vulnerabilities information, jobs, and vulnerabilities found, which provides a structure that can be easily merged with other information groups. The flexible framework accommodates scanner tools that use highly normalized storage structures or simple flat-file results.

The current SIDW framework incorporates data-driven metrics. The relative health of a network over time can be derived from the data sources. Network activity trending above the norm might indicate worm activity, while minimal or no activity on a specific device might indicate a problem with the asset. Normal activity can be measured separately from events. Resolution time for open incidents can be monitored to assure problems are addressed quickly. Vulnerability scans are compared to previous scans to assure compliance with remediation directives. These lower-level metrics provide an assessment of the overall health of the networks monitored by the SOC or CIRT, and deserve their own SIDW warehouse subject area because they are a layer above the detail and summarized data of the input sources.

A number of processes, many dependent on another, are generated to extract the data from the source structure, transform it using the required business rules, and load it into the appropriate detail or summarized data store. SAS ETL Server effectively manages the loading of the SIDW. Job flows are required for the diverse data sources. Job dependencies, which range from detail to summary and OLAP to metrics, are required for the layered data structures. Impact analysis of source and target data elements is necessary to quickly adjust to the schemas and data layouts of the vendor products that provide the source data and to the dynamic requirements of the user community.

MANAGEMENT OF LARGE DATA VOLUMES

The amount of data created and stored by collectors and sensors can be tens of gigabytes a day for a network that employs 50 to 75 IDS sensors, two or three firewalls, and just one border router. This can total over a terabyte of information in less than three months and does not include sniffer-type sources that generate even larger volumes of data. Data generated by vulnerability assessments, asset criticality, and incidents are 1 to 2 magnitudes less than collector information.

This high volume of data requires the following:

- user and application needs must be assessed to plan the data stores
- software and hardware must be configured to accommodate storage and processing
- refresh strategy and time window for processing must be determined
- backup and archival strategy must be fully defined

While implementing the framework at a CIRT, we asked an incident analyst what data makes his job easier. His response was “All of it!” It became apparent that incident analysts want to query large volumes of detail data, preferably with a single tool against one data repository. Other users want to produce basic “Top 10” reports and baseline reports and perform simple searches through a Web-based interface. Response time has to be fast, which dictates the use of summary data stores. Viewing metrics in detail and using drill-down interfaces require another set of summary and OLAP-based structures.

A properly configured hardware and software system is essential when processing, storing, and serving large volumes of data. An improperly configured or tuned system might lead to unacceptable response times or large processing windows. Storage Area Networks (SANs), external storage arrays, RAID technology, cache, and CPU power need to be managed for an effective solution. Software components that serve the data, such as the SAS Scalable Performance Data Server, enable parallelization of data storage processes and require correctly set parameters and input/output configuration to maximize efficiency.

A data refresh strategy is driven by processing and user requirements. Collector and vulnerability data use append processing and incremental updates. Re-processing captured information is not needed because it is static in nature, which maintains the smallest processing window. Incident and asset criticality ETL jobs might re-process all records to capture updates to the respective sources and provide the most current view. Because these jobs are small, refreshes have little impact on processing time.

Large amounts of data require large backups. An organization cannot afford to lose a terabyte data store that required hours of resource time to assemble—this plunges the user community into an information blackout. Data can be backed up incrementally by using SAS Scalable Performance Data Server utilities. Large data sets are reconstructed from incremental backups if necessary, which creates a desirable situation in the event of a disaster. Data can be “rolled off” based on a time dimension to keep data sets online and available to users at a reasonable size. This keeps the processing window within limits and provides efficient response time. Using the SAS Scalable Performance Data Server **Partition By Value** capability provides this functionality. Data rolled off can be archived to media to provide near-line or offline storage, which users can access with a little more effort.

EXPLOITATION OF THE DATA BY USERS

To effectively and efficiently deliver large volumes of data, a variety of user interfaces is required for querying, reporting, and analyzing by users who have different skill sets. Flexible choices in the SAS Enterprise BI Server package make it the ideal choice. The center of the framework is the SAS Information Delivery Portal, which is the wrapper for all Web-based content, and is complemented by the SAS Add-In for Microsoft Office and SAS Enterprise Guide. Each of these interfaces provides a specific set of capabilities to users.

For the different SIDW consumers, there is overlap in the information that each of the users requires for their jobs and individual information requirements. The reports and queries that provide this information can be both dynamic and static, but regardless of the flexibility, reports and queries must perform well against large volumes of data and results must be generated quickly. A set of Web-based content providers meets these parameters and fulfills the common and individual requirements.

The first Web content provider is an IP address query that provides a single snapshot, displayed on one screen, of all CND data about a specific IP address. This is useful to watch operators, incident handlers, and incident analysts. The interface enables users to specify a source or destination IP address, a time period for the query, and the type of content they want to view: basic WHOIS information, Collector Events, Collector Non-Events (Information Alerts), Incidents, and Vulnerabilities. Output consists of a textual summary of information with the option to view manageable detail information about collector information and incidents.

The second Web content provider is a “Top 10” report-generating query that provides views of activity seen on a network through the different types of collectors. The user specifies a time period for the report, a collector to filter by (if desired), and a “Top 10” report to create. The “Top 10” reports available include the following:

- Top 10 Source IP Addresses
- Top 10 Destination IP Addresses
- Top 10 Source-Destination Pairs
- Top 10 Destination IP-Port Pairs
- Top 10 Network Resource Services

Output consists of graphical and tabular summaries that can be useful to watch operators, incidents handlers, and incident analysts. Executive-level decision makers also use this information because it provides a high-level summary of what is happening on the network.

The next Web content provider is a Baseline Report of network activity. Opinions on the value of baseline reporting vary. Limitations of baseline reporting need to be kept in mind, especially that a baseline report might not distinctly show malicious activity mixed in with "normal" noise. Baseline information should be used as a starting point for identifying possible candidates for further, in-depth analysis. Because of the limitations, baseline information needs to be used exclusively by highly skilled incident analysts.

The Baseline Report enables the user to specify a time period of interest (for example "Yesterday" or "Previous 7 Days"), a collector type to filter by, and whether output should be broken down by severity. The output contains a line graph showings the specified time period against a four-week, same-day, running average. The running average eliminates time and seasonal variances in network activity. An example of a Baseline Report is shown in Figure 4.

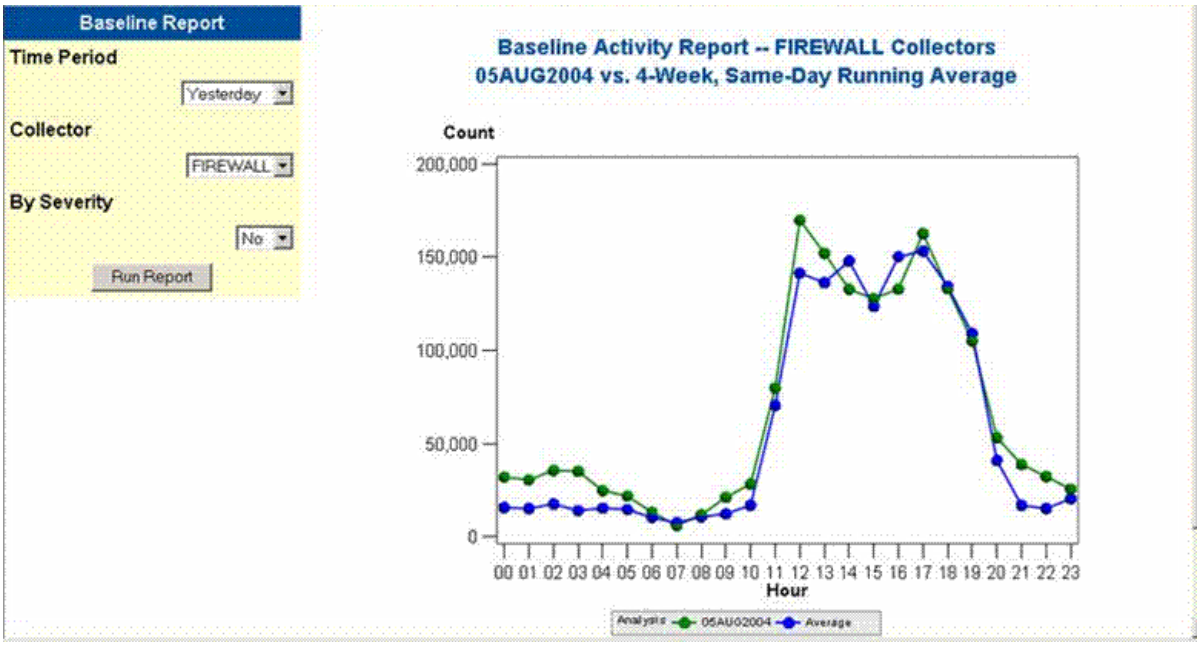

Figure 4 – Sample Baseline Report

Each of the three Web content providers generates useful information that is critical to effective CND operations. By going one step further and linking providers together, interaction with content can be greatly enhanced. For example, output from the "Top 10" report is interactive. Clicking on an IP address in the graphical or tabular output takes the user to an IP address query screen that is populated with the parameters of interest. Data points on the line graph of the Baseline Report take the user to a "Top 10" report of source IPs and subsequently to an IP address query. CND data users move easily between the providers as they progress through their analysis.

WEB-BASED SITUATIONAL AWARENESS

The last of the Web content providers is a console-like component designed specifically for the executive level of a CND organization. This component contains key CND information to provide decision makers with an overall snapshot of their network security operations. A streamlined set of reports and data-driven metrics provide operational status. The data subject areas of the metrics include, but are not limited to, collector events, collector non-events, incident identification and remediation, vulnerability identification and remediation, criticality of network assets, and network topology.

Collector events and non-events are used in a similar manner. As situational awareness is updated over an established time period (12 to 24 hours), significant alert and information activity is analyzed against an established baseline. Activity falls within standard deviation boundaries and results in a situational status score of "Normal/Expected," "Marginal," or "Critical."

Incident identification and remediation is assigned a situational status score using parameters specific to the subject area. Parameters are the number of new tracks opened and their corresponding severity and the number of open tracks not remedied and their corresponding severity. The latter is a measure of incident response time. Vulnerability identification and remediation is assigned a score in a similar manner. Analysis is based on the identification and severity of new vulnerabilities, the existence of recurring vulnerabilities, and the average remediation time for a vulnerability.

For incidents and vulnerabilities, the situational status score follows the same taxonomy used for collectors, but incorporates another factor—asset criticality. Asset criticality assigns a score (usually ranging from 0 to 1) to identify the importance of that asset to the organization's successful operations. Critical data and application servers have a score closer to 1, while less critical servers and desktop machines have a lower score. When assigning situational status scores to incidents and vulnerabilities, host criticality is considered. The greater the number of critical hosts affected, the worse the status score.

After situational status scores are determined, they are summarized for higher levels of detail. This higher level of detail uses network topology to summarize data by areas such as business unit, department, and overall organization. As metric subject areas are combined, weighting factors are used, based on the confidence in the data in the various metric subject areas. Summarizations provide the snapshot view of overall network security health and are displayed in a simple table or in an organization chart or map. Regardless of the display, drill-downs into individual metric levels are included. The executive level can get greater detail about why the status is at its current assigned level and prepare questions for staff regarding efforts to address the problems.

CND PORTAL

After Web content providers have been developed, using tools such as SAS AppDev Studio and SAS Web Report Studio, a wrapper is created to host all of the provider content. In the SAS framework, the SAS Information Delivery Portal provides the wrapper that creates the CND portal structure where all users perform reporting and analysis. Different views can be created for different users: one view for watch operators and incident handlers, one view for incident analysts, and one view for executive-level decision makers who want to view content related to situational awareness. A public kiosk can provide content to users outside an organization. In the CND community, related organizations share information and create an alliance against hackers who practice information warfare. An example of a portal view for situational awareness is shown in Figure 5.

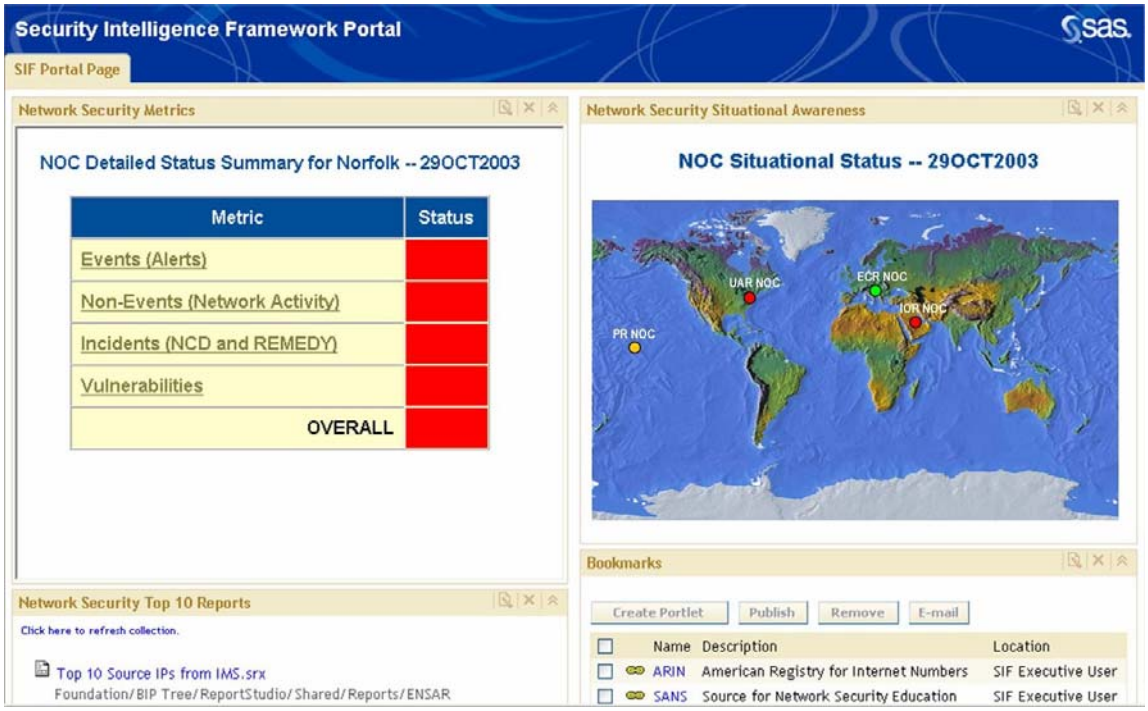

Figure 5. Portal View Showing Situational Awareness

SAS ENTERPRISE GUIDE – WORKBENCH TOOL

Even though Web content providers produce much information for different CND data users, they are limited in the detail they provide and the user-guided analysis they enable. Incident analysts require a powerful workbench tool to perform reporting, querying, and analyzing, from simple listing reports and queries, to more advanced statistical analyses that look for patterns and perform data visualization. SAS Enterprise Guide is an ideal choice for forensic investigators. All CND data in the SIDW is at their fingertips. They define exactly how they want to use data from the different subject areas, how to combine or integrate data according to their needs, and when to produce reports that can be pushed to the Web or provided to users through various publish and subscribe channels of the SAS Intelligence Platform. Because of the advanced analysis capabilities of SAS Enterprise Guide, incident handlers have a powerful set of tools to effectively perform their CND responsibilities.

REPEATABLE TASKS THROUGH SAS STORED PROCESSES

Every CND organization has unique processes based on business rules. For example, many CND organizations generate IP block lists that are used to configure and maintain the point products that monitor the network. The basis for these IP block lists might come from the analysis of information provided by correlation engines or SIMs, incident information, or both. This type of analysis lends itself to the creation of SAS stored processes that can be executed through interfaces such as SAS Enterprise Guide or the SAS Add-In for Microsoft Office. A SAS stored process can be created to generate the IP block list and then executed by an incident analyst in Microsoft Excel to view the output. The incident analyst can manipulate the output in Microsoft Excel to create a report that is pushed to the team responsible for sensor configuration.

The wide variety of interfaces available through SAS Enterprise BI Server and the data management and statistical analysis strength of the SAS Foundation are extremely powerful components in this framework. The technology packages provide the right interfaces to the right users so they can access the right data at the right time.

THE NEXT LEVEL – “SLOW AND LOW” DETECTION

By now, large organizations have deployed various real-time point products as part of their CND strategy. However, new capabilities are needed to protect against the sophisticated attacker who goes undetected and causes the most damage. This type of network penetration is referred to as the “slow and low” attack. Current network security devices are not effective in detecting these attacks because these “under-the-radar” penetrations usually span longer periods of time, making it difficult to uncover any type of pattern or signature. Examples include zero-day exploits, IDS-evasion techniques, and stealth network reconnaissance. They originate from customized Trojan horses and compromised kernels or programs. Defense requires a new approach to identify and alert organizations to this type of attack.

Currently, research is underway between SAS and the Department of Defense to study and apply existing fraud detection methods to detect “slow and low” activity. Fraud detection concentrates on the behavior of events, rather than the events. Ultimately, all network traffic reflects human actions and, therefore, displays innate behavioral characteristics. These characteristics, such as temporal, peer, or volumetric patterns, are not examined by any commercial approach to network security. However, they can pinpoint anomalous properties of traffic that would normally go undetected. The ultimate goal is to identify long-term systematic network penetration, by a dedicated aggressor, by using a series of self-tuning threat indications and early warning notifications. The result is a safety net for real-time point products (IDS sensors and firewalls) and an improvement in overall CND.

REAL-WORLD EXAMPLE – NAVCIRT

The threat to Navy cyber security and shipboard LAN systems is real. In addition to sea-skimming, anti-ship cruise missiles and other conventional attacks, Navy commanders face the threat of offensive information warfare by rogue governments and terrorist groups.

Information warfare is the new operational battle space. Through information warfare and cyber terrorism, terrorists and foreign governments accomplish political objectives without firing physical bullets. Hackers can gain access to a government or military classified network. They can plant information bombs or computer viruses designed to destroy or render computer networks that control weapon systems, financial transactions, and other communications traffic inoperable. The Navy Computer Incident Response Team (NAVCIRT) is responsible for CND for the United States Navy. The NAVCIRT, based at Little Creek Amphibious Base in Norfolk, Virginia, watches for virus attacks against naval networks worldwide, as well as other types of intrusions or disruptions of service. The NAVCIRT understands the impacts of malicious activity on Navy networks, and it analyzes how intrusions happened and how to keep them from being repeated. The NAVCIRT mission statement is “to coordinate and direct the defense of all Navy computer systems and networks in support of Navy forces ashore and afloat; to constantly improve the Navy's

defensive network posture; to maintain a secure and interoperable network through policy, guidance, education, incident detection, analysis, and incident response; to prevent computer incidents through network surveillance and proactive measures against potential threats."

Before the system eventually named MOBIUS, the NAVCIRT was faced with many challenges. As a basic requirement, the NAVCIRT protects Navy networks and provides a comprehensive view of network conditions to the Commander. This includes the fused operational view and protection of four primary classes of networks, including the afloat and ashore networks, the major Navy network outside the United States, and various Navy specialty networks. To meet this challenge, the NAVCIRT must provide the following at an enterprise level:

- View of network topology, software configuration management, and traffic flow
- Fused view of IDS data
- Fused view of router and firewall logs
- Fused view of malicious code activity
- View of system audit logs
- Status of defensive measures

Before MOBIUS, a number of factors made these efforts a challenge. The Navy initially deployed many point products, each of which captured and produced data in different formats. The data volumes are enormous and continue to grow, exponentially. The NAVCIRT lacked event correlation capabilities, which hampered efforts to gain real intelligence on a potential threat. In fact, the NAVCIRT did a study and determined that its watch operators could spend only 2.9 seconds on every IDS event or alarm to assess its potential threat. At the time, the NAVCIRT could not store cyber security data for historical analysis, trending, data visualization, and reporting. They could not provide the commanders with situational metrics on the status of the network. The solution was not to increase staff but to build MOBIUS. This was a two-part effort that would provide event correlation capabilities and create a data warehouse to integrate and store all CND data. In addition, it would provide the appropriate interfaces to make that data available to consumers. The SAS Intelligence Platform was selected as the backbone for the data warehouse and the user interfaces.

NAVCIRT SOLUTION AND IMPLEMENTATION

Based on the deficiencies, a SAS solution was architected for the MOBIUS Security Data Warehouse. This solution consisted of SAS 8.2 software and three-to-four months of services to implement the first iteration of the proposed framework. The software components included Base SAS, SAS/ACCESS, SAS/Warehouse Administrator, SAS Scalable Performance Data Server, SAS/SHARE, SAS Integration Technologies, SAS AppDev Studio, and SAS Enterprise Guide.

Implementation began in March, 2003, with detailed requirements gathering. One of the key activities was an interview with each NAVCIRT data user, watch floor personnel, incident handlers, incident analysts, and management/command level personnel. Interviews identified the goal of the engagement and critical factors that would ensure its success. Interviews identified application requirements for each data user:

- A Web-based interface with limited dynamic reporting and querying capabilities such as time-based IP querying of integrated CND data, "Top 10" reporting, and baseline reporting
- A desktop workbench tool for more involved querying, reporting, and complex analyses
- A situational awareness console that provides daily, metric-oriented information regarding the overall health or status of monitored networks

Following the interviews, a final Statement of Work was generated, detailing the project and providing a timeline. The rollout date for the first iteration of the MOBIUS Security Data Warehouse was targeted for the end of October, 2003. Using the information collected during requirements gathering and user interviews, a design document that outlined application interfaces and back-end data structures was presented to NAVCIRT in April. This design document was reviewed by the NAVCIRT project team and the individual data users who participated in the interview process. Interfaces included SAS Enterprise Guide for use by the incident analysts as a heavy duty analysis workbench tool, and the following Web-based modules:

- Historical Analysis Module (HAM), which provides limited dynamic reporting and querying capabilities
- Situational Awareness Module (SAM), which provides a metrics-based view of monitored networks
- Data warehouse documentation, which provides technical and business metadata to MOBIUS users

On-site development began in May, 2003, with the installation of software and the development of the warehouse to retrieve CND subject data, transform and cleanse it, and load the physical data structures. The extract process

included sensor-based data from an event correlation engine implemented by NAVCIRT earlier that year. After installation and development were performed, testing was completed and the data warehouse was production the last week in August. Web interface development was started and completed at the end of September. Testing of the Web interfaces and SAS Enterprise Guide was completed in October and the full system was production by October 31, 2003. The concept of MOBIUS was a reality.

CURRENT STATUS AND FUTURE ENHANCEMENTS

By the summer of 2004, the data warehouse and associated applications had been operational for almost 10 months with a less than 1% failure rate. During this period, the SAS®9 Intelligence Platform was released. It included key improvements that could significantly enhance the MOBIUS Security Data Warehouse. These improvements included scalability enhancements that would greatly improve the performance of the data warehouse processes, and usability enhancements that would enable greatly flexibility of interfaces available to the user community, and most important, the SAS Information Delivery Portal that would fulfill the NAVCIRT concept of a full-service CND portal interface. Three SAS technology packages are the foundation of these improvements: SAS ETL Server, SAS Intelligence Storage, and SAS Enterprise BI Server.

Recognizing the advantages available in SAS®9, installation and configuration of the newer release was completed in December, 2004. The migration to a SAS®9 foundation is in progress and incorporation of SAS®9 interfaces should be completed in spring of 2005.

LESSONS FOR SUCCESS

A number of factors contributed, and will continue to contribute, to the overall success of the NAVCIRT implementation. These factors include a phased implementation approach and iterative development life cycle. However, two factors are more significant than the others—a truly collaborative environment and a full integration of the stakeholders. Ongoing user interviews ensure that the data warehouse is meeting the requirements of day-to-day activities. More important, interviews enable every user to guide the direction of the project. Everyone is invested in its success.

CONCLUSION

The threats to computer systems and networks will never be stagnant. They continue to grow in volume and constantly adjust and re-invent themselves as different measures are put in place to defend against them. The goal of this paper is to emphasize that CND is more than sensors, real-time SIMs, and event correlation engines. Complete CND is a strategy that consists of a framework that includes these real-time sentinels with a historical facet of all integrated CND information as the backbone. The SAS Intelligence Platform can serve as that backbone in any organization's CND infrastructure.

RESOURCES

Kazienko, Przemyslaw and Piotr Dorosz, 2004, "Intrusion Detection Systems (IDS) Part 2 - Classification; methods; techniques". Available http://www.windowsecurity.com/articles/IDS-Part2-Classification-methods-techniques.html.

Magalhaes, Ricky M., 2004, "Host-based IDS vs Network-based IDS (Part 1)". Available http://www.windowsecurity.com/articles/Hids_vs_Nids_Part1.html.

Northcutt, Stephen, et al., 2003, *Inside Network Perimeter Security: The Definitive Guide to Firewalls, VPNs, Routers, and Intrusion Detection Systems*, Indianapolis, IN: New Riders Publishing.

Schweitzer, Douglas, 2002, *Securing the Network from Malicious Code: A Complete Guide to Defending Against Viruses, Worms, and Trojans,* New York: Wiley.

ACKNOWLEDGMENTS

The authors want to acknowledge the highly skilled cyber security professionals at both the NAVCIRT and Naval Research Laboratories. The framework outlined in this paper would not exist without the true collaborative relationship that is the foundation behind the success of the MOBIUS Security Information Data Warehouse.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors:

Mike Protz
SAS Institute Inc.
6501 South Fiddler's Green Circle, Suite 600
Greenwood Village, CO 80111
Work Phone: 303-290-9112 x1763
Fax: 303-290-9195
E-mail: Mike.Protz@sas.com

Joe Zilka
SAS Institute Inc.
6100 Oak Tree Boulevard, Suite 400
Independence, OH 44131
Work Phone: 216-642-0641 x1117
Fax: 216-642-4226
E-mail: Joe.Zilka@sas.com

Jeff Mudd
SAS Institute Inc.
SAS Campus Drive, U-3058
Cary, NC 27513
Work Phone: 919-531-3960
Fax: 919-677-4444
E-mail: Jeff.Mudd@sas.com