



Interested in learning more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Distributed Security Management for the Enterprise

Managed security is the next step in the lifecycle of the network security industry. The information flow within an infrastructure today is unmanageable. Information comes from so many different sources and in such large quantities that identifying a potential security risk in real time is near impossible. The focus of this paper is on managed security, specifically one product that has been on the market for almost a year, Spectrum Security Manager. There has not been much mention of products like these in the SANS co...

Copyright SANS Institute
Author Retains Full Rights

AD

An advertisement banner for Watchfire. On the left, there is a graphic of a globe and a login form with fields for "lo" and "passw". In the center, a dark blue box contains the text "Testing Web applications for vulnerabilities?". On the right, the Watchfire logo (a red flame) and the word "watchfire" are displayed.

Testing Web applications for vulnerabilities?

Distributed Security Management for the Enterprise

GSEC Practical V1.2f

William DiProffio
Jan 3 2002

© SANS Institute 2002, Author retains full rights.

ABSTRACT

Managed security is the next step in the lifecycle of the network security industry. The information flow within an infrastructure today is unmanageable. Information comes from so many different sources and in such large quantities that identifying a potential security risk in real time is near impossible. The focus of this paper is on managed security, specifically one product that has been on the market for almost a year, Spectrum Security Manager. There has not been much mention of products like these in the SANS conferences that I have attended. It would be of great benefit for people in the security industry to know that there are some products that will actually help them with managing the piles of information they are forced to handle.

I briefly describe the existing problem in the industry and then discuss the product, its architecture and how it is implemented. Spectrum Security Manager receives information from your existing security infrastructure devices/software in the form of but not limited to SNMP traps, Syslogs, SMTP, and TCP/IP streams. SSM takes this information and stores it in one centralized database which allows the user to do forensic analysis to track events through their system. SSM has a rules based correlation engine, which then allows the user to correlate information and events across the entire infrastructure in real time. Using these rules the user can be alerted when a potential security violation exists. Also using something called dynamic mapping the user can see graphically the layout of attacks happening in their infrastructure. The benefits of a package like this are obvious, centralized data management, forensic capabilities, and real time monitoring of the security infrastructure.

© SANS Institute 2002, All Rights Reserved.

Distributed Security Solution for the Enterprise

Introduction

Network infrastructures are the lifeblood of today's large-scale enterprises. Securing those networks is now more than ever a full time job for a team of network professionals. Attacks from the outside as well as from within have time and time again crippled some of the largest companies in the world. Denial of Service (DoS) attacks coupled with numerous and always-present virus attacks limiting services and compromising sensitive information have cost companies billions. A single firewall can no longer be considered an adequate defense from the outside. Now network and host security is comprised of a complex infrastructure of security products.

In today's complex security infrastructures it can be difficult to stay current and keep on top of day-to-day operations. "One of the toughest jobs for information security professionals is getting real-time information about what's actually happening across their company networks."(Hulme) The security system administrator who often also wears the hat of the network system administrator and his/her team is given this daunting task. It is becoming increasingly necessary to consolidate and correlate information from many different sources to identify security risks and/or breaches. A product that can bring all of your security information under one pane of glass is now an essential tool for the security professional.

Defining the need.

What defines an enterprise's security infrastructure? The infrastructure is often times made up of many disparate systems. These disparate systems often fall into one or more of the areas described in the security process flow chart seen below in *Figure 1*.

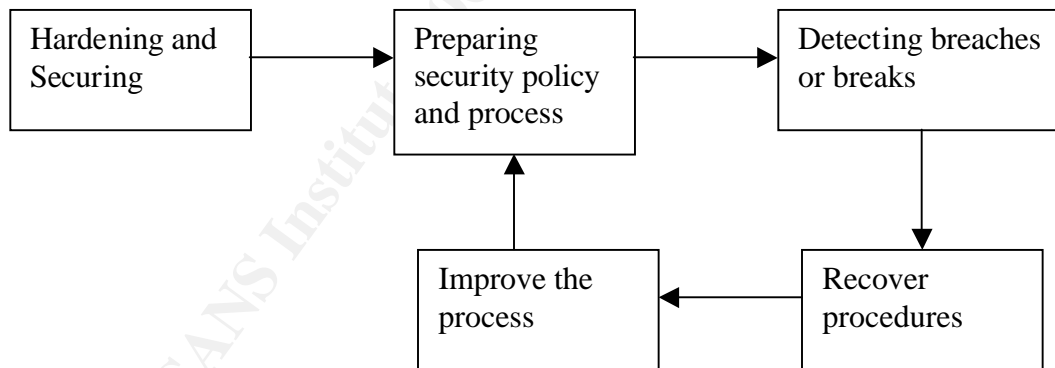


Figure 1: Security process flow chart. When a product is added to the security infrastructure this process is ongoing.

Products that you would expect to see falling into this flow chart include firewalls, Intrusion Detection Systems (IDS), operations systems, security scanners, packet sniffers, anti-virus applications and PKI and directory infrastructure packages.(Carnegie Mellon) All of these products work separately to help protect the network. In order to do so they must all produce some form of usable output. Most times this is in the form of logs or

event streams. If the security infrastructure is substantial enough, being able to keep track of output from each of your disparate systems is a mammoth task.

There are also components of your security infrastructure that are native to the network and do not require any additional financial investment. These are more passive means of defending your network and also come in the form of logs. Most if not all of your network devices, from routers and switches to hubs, have logs that can be enabled to track everything from statistics to traffic flow to log-on attempts. In addition, most if not all operating systems within the infrastructure have the ability to produce logs, whether they are in the form of event and security logs on an NT/Win2000 machine or Syslogs on a Unix box.

At this point the real problem becomes apparent. With all of these systems each producing output whether it is passive or active the ability to stay on top of your security infrastructure in real time becomes nearly impossible. "Managing a security infrastructure is people-intensive, hence error-prone." (Walker) In order to identify problems or breaches within the network there is, many times, a team of employees gathering and parsing these logs on a daily basis (That is if they have time to actually review this data). It is not only important to parse the logs and pull out pertinent data, but to be able to correlate that data across the system. If there is a team of engineers each responsible for some security component or components, how do you make sure that some incident that is minor individually but significant when distributed across the system is not missed.

For example, an enterprise has an extensive network with numerous links to the internet each being monitored by an IDS. Each IDS reports multiple port scans daily, which is not uncommon. In and of themselves the port scans are most likely insignificant. However if there were a port scan identified from the same source ip directed at the same port on each IDS across the entire network that may be very significant. It could signify a very focused attempt to find a way into the corporate net rather than some random port scans across a range of IPs. Would a standard security-monitoring program catch an attempt like this? No. Each team monitoring their individual IDS would see the attack as a typical random port scan and disregard it, not mentioning it to the security infrastructure as a whole. As a result there is no action taken to block this ip and the attacker is free to continue to distribute his attack until he eventually breaks through.

The enterprise needs a system that will take event streams and logs and bring that information under one pane of glass. "Service providers maintain large distributed networks to serve their customers. This creates a need for distributed security-management" (Gartner) With distributed management information can be collected in a single database event and log information can be organized, filtered and correlated. The event correlation will recognize the distributed attacks like the one mentioned above and give the enterprise a chance to defend itself in real time. SPECTRUM Security Manager from Aprisma Management Technologies is one product that does specifically that.

SSM Product is one answer

Aprisma has recognized the need to bring the enterprise security infrastructure under control. They released SPECTRUM Security Manager in the winter of 2000 to assist large and small enterprises manage their security infrastructure. In a nutshell

SPECTRUM Security Manager or SSM is a software product that ties together all of the disparate systems within the infrastructure under one pane of glass allowing administrators to make educated decisions based on real time security information. There are a couple of other companies on the market that attempt to tie the security infrastructure together like Net Forensics and Esecurity, but their attempts are focused mainly on certain parts of the infrastructure such as Cisco Syslogs or Snmp traps only and don't span the whole infrastructure like SSM.

SSM Architecture (How it works)

The SPECTRUM Security Manager architecture is not overly complicated. It is simply a matter of collecting the data, manipulating the data into a usable format and then correlating that data to identify security threats to the infrastructure.

There are two major components in the SSM architecture. The Master Server or MS and the Consolidator or CS. Each run on either Solaris or WinNT machines. The MS acts as the central console for SSM. From the MS the administrator can access the GUI interface for SSM. Policies and rules are written at the MS and then pushed out to the Consolidators. The Master Server and Consolidators contain the exact same intelligence but are constrained by licensing. There is no GUI access at the Consolidator level.

The MS and CS format gives SSM the ability to handle large complex environments. In some environments the security infrastructure can be responsible for millions of security events per day, ranging anywhere from a firewall's 'denied access' event to an intrusion detection system's 'attack signature recognized' event. It would be unreasonable to expect a single workstation to take in, filter and correlate this flood of events in real time. The CS allows for these events to be filtered at a lower level. With this architecture the product becomes scalable from the smallest business environment to the largest enterprise infrastructure. An example of this architecture can be seen in *Figure 2*.




Figure 2: The SSM architecture lends it self to be very scalable. This figure shows three geographically separated sites. The site with the MS and the two sites with CS's. Information is filtered and correlated at **site A** and **site B** separately and then sent up to the MS station where it is filtered and correlated again and alarms and alerts are generated. This way we limit the traffic and load for the MS station.

Network traffic and data come in many forms; SNMP, SMTP, TCP/IP Streams, syslogs to name some of the major protocols. SSM is capable of receiving information from the security infrastructure in each one of these forms. The first step in the architecture is *Data Collection*. Getting the data to the SSM Master Server can be as easy as pointing the IDS snmp trap policy at the SSM MS. In some instances however there must be some customization done to get data to the MS. This can be in the form of a Perl script which grabs data as it comes into a log file and configures it to the correct form and forwards it on to SSM or in the form of some other data conduit between the MS and the data collection tool.

Once the data is flowing to SSM we move to the second phase of the architecture which is *Normalization*. This is one of the most powerful components of the SSM architecture. Normalization happens within the protocol components as seen in *Figure 3*. It is here that the data will have the vendor specific tags stripped from it. What this means is that independent of what vendor's firewall or IDS system you have messages will be normalized to something the user can understand. For example each vendor's IDS system has a code that represents a certain event, say a port scan. SSM will take that code and transform that message into the form `ids.detect.probe.portscan` regardless of the vendor IDS that produced the event. Decoders that are pieces of java code, dissect the message and produce the normalized event in the proper SSM format. There are presently decoders written for all of the major security packages on the market a list of these is included in Appendix A. Now if the infrastructure is utilizing three different Fire Wall vendors the user will get the same event message for a "packet denied" alert regardless of vendor.(Aprisma)

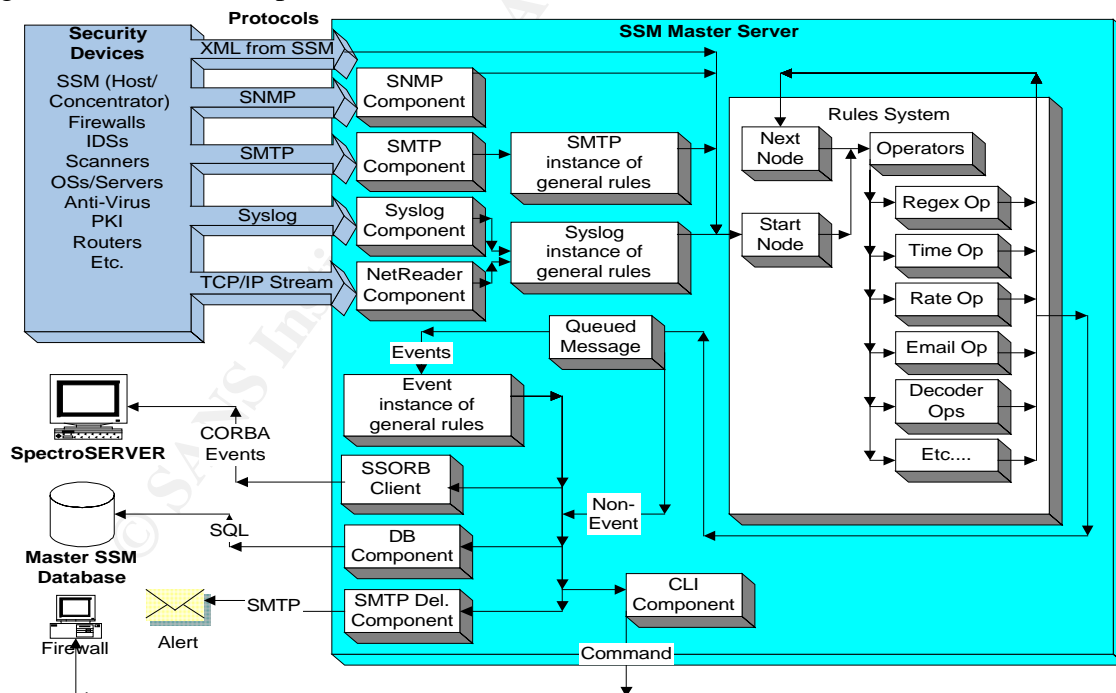


Figure 3: SSM architecture flow chart showing supported protocols and flow of event message within the product.

Not only do the decoders normalize the data coming in, but they also manipulate the data into the predefined database fields so the data can be stored if desired. Each event that comes into the system has specific information associated with it. Some of this information may be Target IP, Source IP, Protocol or Time and Date Stamp. The decoder peels this information from the event message and populates a table in the SSM database. Once the event is properly formatted the information can be correlated and used in forensic analysis at a later date. A full list of the default database values can be found in Appendix B. If there is some part of the event message that SSM doesn't have by default it can easily be added.

If the user chooses the data can be stored in a database. The current release of SSM supports both Oracle 8i and SQL7.0. The ability to store these events in a single database out of the box is a revolution in the security management space. Prior to this the security events within an infrastructure had to be gathered individually by scripts or some home built solution. With SSM we now have centralized data storage for all of the security events within the enterprise.

Immediately we see the advantage of centralizing the data. This information can now be correlated across the enterprise as it comes into the system in real time. Spectrum Security Manager has a rules based correlation engine. Most professionals in the network management space are familiar with this type of engine. The vast majority of rules based correlation engines on the market are programmatic in nature, meaning that there is some manual coding involved in the creation of rules. The SSM engine is primarily GUI based. What this means is that by using a set of very adaptive GUI operators the user can correlate the events as they enter the system in real time. The rule set also uses what are referred to as Boolean edges. These edges represent either a path of true or false. The edges are used to connect the operators so in essence you end up with a flow chart of operators that eventually reach some desired result. *Figure 4* shows examples of some typical operators.



Figure 4: Eight standard operators used in the creation of rules for the correlation engine. Operators manipulate, compare, modify and combine information from the events that are streaming into the SSM product. There are significantly more operators in the full rules set.

To show the simplicity with which rules are created let's take a look at a very simple case. A small enterprise has three critical servers within its infrastructure. The Security/Network Administrator for each system would like to know when any event is generated against the server they are responsible for. Their IP's are 10.4.0.5, 10.4.0.6, and 10.4.0.7. When an event is generated against one of these systems the person responsible for the server should receive an email. In the event that multiple events occur across the servers then the department manager should be emailed. Also the events should be stored in the database for future forensics. The operators are placed by choosing an operator and left clicking in the workspace. The Node Attributes of each operator can then be set to accomplish the intended purpose. An example of this rule as seen in the SSM graph space can be seen below in *Figure 5*.

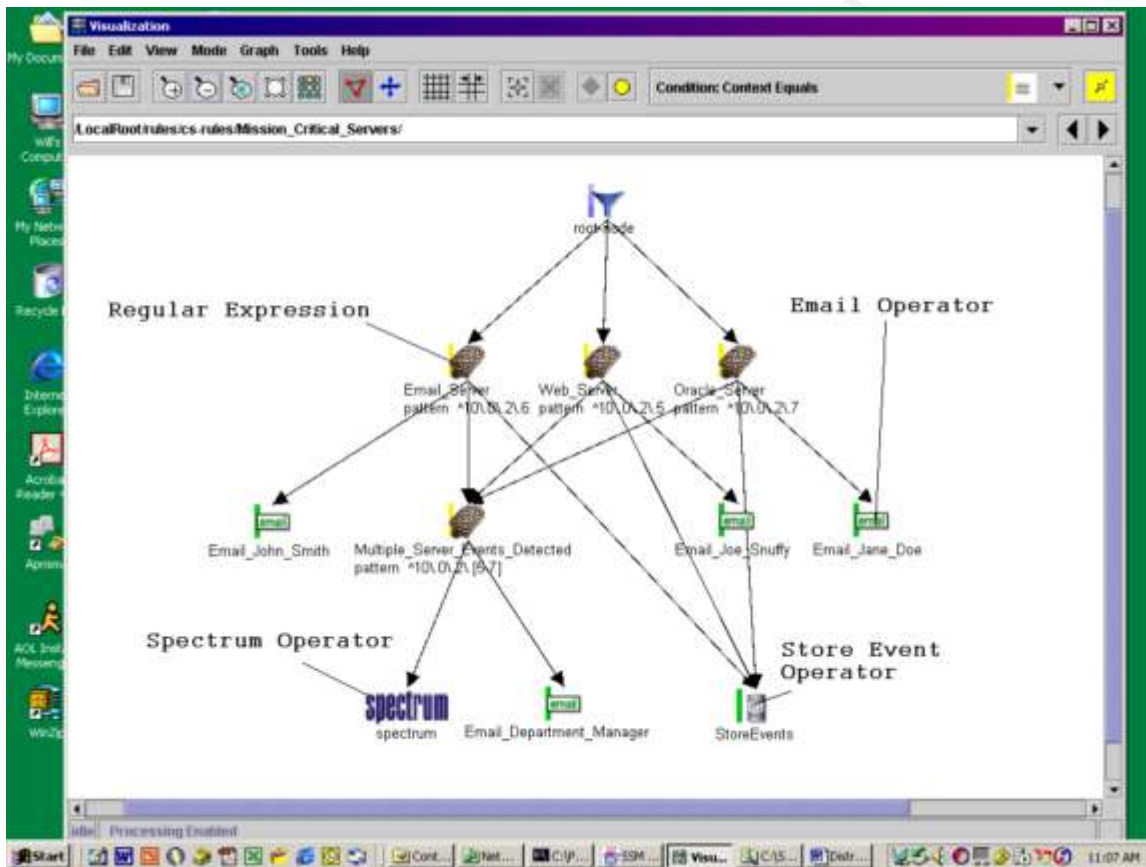


Figure 5: Mission critical rule space showing how the operators can be used to inform the personnel responsible for each of the servers.

Of course in a real world environment the administrators would not want to be informed of every event that concerns the servers. There would be many events filtered out to the database and other correlation operators to make sure that only significant events draw attention for human intervention.

Using some very simple rule combinations it is possible to get a pretty good grip on your security infrastructure in a very short period of time. The above example takes into account some specific devices, but what about the rest of the enterprise. Using the asset classification and prioritization capabilities of the SSM product it is easy to get a

handle on the events coming into the infrastructure. With an asset classification rule it is possible to assign a zone to each of your monitored devices, whether they are in a range of IP's or random across the enterprise. An example of a zone using the SSM convention would be ny.dmz.webserver. This means that when an event comes into SSM that concerns the webserver in the dmz in NY, the zone(Zone is a field in an SSM message left blank by default) will be set to ny.dmz.webserver. Using the zone attribute to classify assets allows the user to group events by asset and apply priority to them.

For example if the events concerning the ny.dmz(Notice how you can move up the zone tree to include all events in the entire dmz and not just specific devices like ny.dmz.webserver) are more critical than the events concerning the ny.3rdfloorlab then as the events come into the system the priority field can be set to some rating convention. It can be broken down even further by filtering on the type of event also. If IDS events in a zone are more important to the user than Firewall events then reset the priority. The user is only limited by imagination and the security policy that they develop.

Now with priorities set on all events coming into the system the user can be notified based on priority of events.

Below are some example of rules that are possible with the SSM system. These are some of the most basic rules.(how to develop some of these rules have not been discussed and further research by the reader into the product would be necessary)

- Advise me of all events with a priority greater than 90
 - If the dmz sees more than 30 events of greater than 60 priority in a specified time advise me
 - Store all events with priority 30 in the database but do not advise me
 - If my critical servers see events greater than 50 advise me.
 - If people who have left the enterprise try to log into any system advise the NOC
 - If more the 3 firewall configurations are changed across my 6 firewall perimeter in a specified time advise me.
 - If there are more than a specified number of dos type events in a specified time advise me.
 - If there is traffic across ports on a server that should not be active advise me. i.e. the email server is chattering on port 80
 - If critical files on a development machine are changed or moved advise me.
- (Intellitactics)

The last rules concept that is important to be familiar with is dynamic graphing. With dynamic graphing it is possible to see the graphical structure of events as they come into the infrastructure. For example events produced by IDS systems have a target IP, a source IP and a type associated with them. It is possible within SSM to show these attributes graphically. The view would show what systems are doing what attacks to what other systems. Given enough events, attacks like DOS can be seen graphically as they happen. A single attack event populated in the dynamic graph is shown in *Figure 6*.




Figure 6: One attack(an event shown in a dynamic graph. In a real world environment there would be numerous attacks shown like this on the screen giving the user a new level of situational awareness.

Dynamic graphing of events as they happen can also give the administrator a feel for what is happening internally to the network. Do you have a box broadcasting when it shouldn't be causing a near DOS attack on yourself? There is much more to dynamic graphing than can be mentioned in the scope of this paper.

The last part of the SSM product that is significant to the Security administrator is it's ability to allow forensic analysis. Since all of the data coming into the infrastructure is being stored in a central location it can all be accessed for analysis. The *Database Query Engine* is a web-based tool that allows the administrator to perform queries against the database.

The queries can be based on a time/date range, a keyword search or a combination of both. *Figure 7* shows a screen shot of the query engine. If a server that is behaving out of character and the rules that the administrator developed have not alerted SSM the Database query engine could be used to show all of the events against that server in the last 20 minutes. Maybe there is something suspicious happening. If an attack has been discovered and the administrator wants to find out the order of events the Database query engine could be used to track the events through the infrastructure.

The results of the queries are displayed as a list of events that match the request. Each one of the returned events can be analyzed individually by drilling down into the event to see all of the specific information associated with it.

As long as the events coming into the infrastructure are being stored to the database then the administrator can perform some level of forensic analysis. This is invaluable in today's fast paced environments where significant events take place only to be discovered later. With the database query engine we have the means to uncover information about those events after they have happened. (Unless you have the rules in place to help you discover the attacks in real time!!)

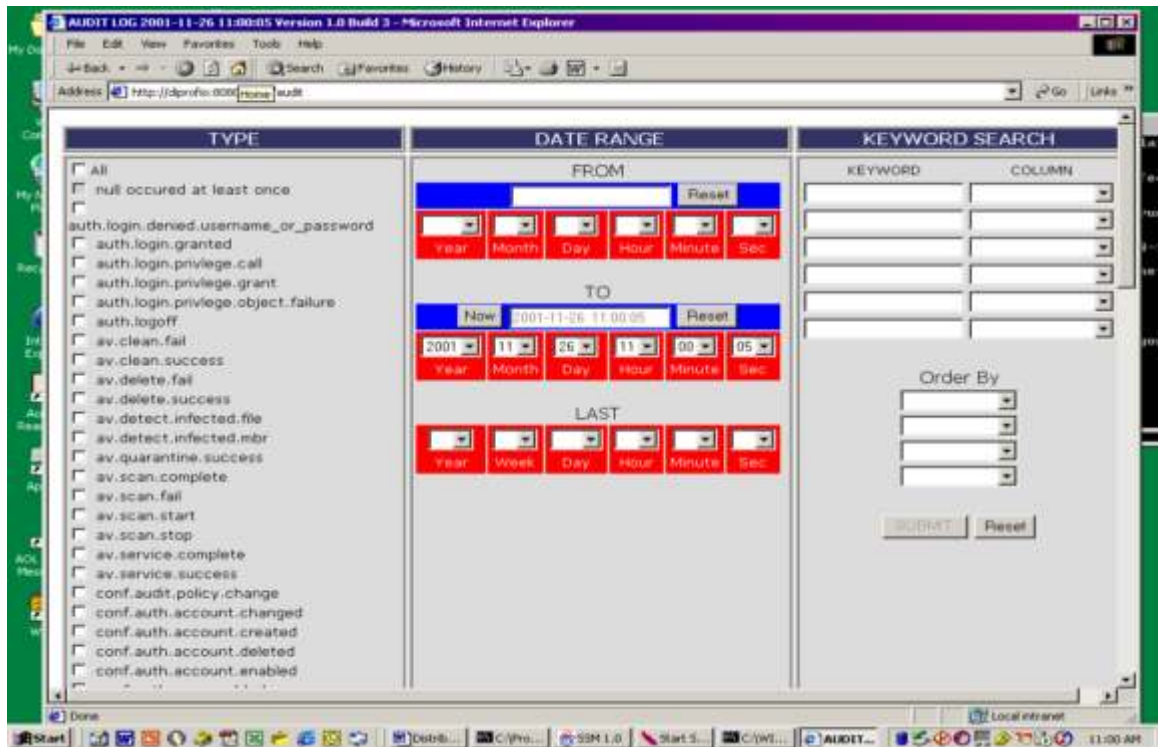


Figure 7: The Database Query Engine allows the administrator to query the central database and gain valuable forensic information about the events coming into the infrastructure.

Conclusion

In this paper we have discussed the need for some kind of order being brought to the chaos that is the security information coming into most enterprise security infrastructures. That need cannot be denied. It is now up to the administrator to discover the best way to manage that information. I am not saying that the SSM product is the best solution for all environments, but it is one solution.

Collecting and centralizing the security events across an infrastructure independent of vendor and type is invaluable. Allowing for one place to look in case of an emergency, one central location for forensic analysis is huge. Many enterprises spend huge amounts of man-hours developing custom solutions for just this. Bring together the centralized data storage with the ability to correlate the events as they enter the system in real time and the possibilities are exponential.

As the security industry grows and matures managed security will come to the forefront. In the near future many products will be developed with these ends in mind. The SSM product is only one of a handful that are already out there and in my opinion it is already ahead of its time.

References

1. Walker, John, Ph.D , NetIQ Corporation, 2001
http://www.netiq.com/Downloads/Library/white_papers/Security_Event_Correlation-Where_Are_We_Now.pdf
2. Hulme, George. "Centralized Security Management on the Way." Information Week. 28 May. 2001 <http://www.informationweek.com/839/security.htm>
3. Gartner Inc. 2001 <http://www.gartner.com/webletter/aprisma/article2.html>
4. Carnegie Mellon Software Engineering Institute. Oct 18, 2001. "Cert Security Improvement Modules." <http://www.cert.org/security-improvement/practices/p094.html> and p095.html
5. Intellitactics. 2001. "Network Security Manager Administrators Guide" Version 3.2. pages 65-71.
6. Intellitactics 2001. "Advanced Rule Training for NSM version 3.2" pages 41-60.
7. Aprisma Management Technology, 2001, "Managing What Matters" presentation pages 1-21

© SANS Institute 2002, Author retains full rights.



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS London 2009	London, United Kingdom	Nov 28, 2009 - Dec 06, 2009	Live Event
SANS WhatWorks in Incident Detection Summit 2009	Washington, DC	Dec 09, 2009 - Dec 10, 2009	Live Event
SANS CDI East 2009	Washington, DC	Dec 11, 2009 - Dec 18, 2009	Live Event
SANS WhatWorks in Data Leakage Prevention and Encryption Summit 2010	New Orleans, LA	Jan 07, 2010 - Jan 12, 2010	Live Event
SANS Security East 2010	New Orleans, LA	Jan 10, 2010 - Jan 18, 2010	Live Event
SANS AppSec 2010 and WhatWorks in AppSec Summit	San Francisco, CA	Jan 29, 2010 - Feb 05, 2010	Live Event
SANS Phoenix 2010	Phoenix, AZ	Feb 14, 2010 - Feb 20, 2010	Live Event
SANS Tokyo 2010 Spring	Tokyo, Japan	Feb 15, 2010 - Feb 20, 2010	Live Event
SANS Geneva CISSP at HEG 2009 Autumn	OnlineSwitzerland	Nov 23, 2009 - Nov 28, 2009	Live Event
SANS OnDemand	Books & MP3s Only	Anytime	Self Paced