

Threat Management Systems
The State of Intrusion Detection

By Steven J. Scott
sjscott007@yahoo.com

August 9, 2002

Table of Contents

- Threat Management System – The State of Intrusion Detection 3
 - Devices 4
 - Aggregation 5
 - Correlation 5
 - Analysis 5
 - Alerting 7
 - Reporting 8
 - Evaluation Criteria 8
 - Concerns 11
 - Conclusion 12
 - References 12

Threat Management System – The State of Intrusion Detection

The current state of intrusion detection can be quite limiting and given the fact that most concentrate on signature based detection we are limited in gauging the success and susceptibleness of an attack. The process of validating an alert is ineptly manual and can become unmanageable when thousands of events are occurring. Most Intrusion Detection Systems (IDS) do not include the intelligence to reliably validate an attack, identify critical assets, correlate information from other devices and then make a response based upon validity and severity of the intrusion.

As a Security Analyst for large fortune 500 company, this researcher is overwhelmed with large amounts of log information and bombarded with alerts. Our network contains some 200 externally accessible servers that span the globe, and are monitored by 30 network based Intrusion Detection Systems (NIDS) from two different vendors. You can probably imagine how much information is generated, and the problems with having two different NIDS management consoles. On top of that, when a suspicious event does occur it is useful to have the firewall, router and systems logs at hand. Currently, this is a process of contacting the firewall admin, server admin and the network admin, or in some cases utilizing yet another management console to gain access to the logs. This whole process is very inefficient and takes time away from really identifying an intrusion. What we need is a unified system that can calculate an attack based upon multiple inputs.

This brings us to the concept of Threat Management and the prospect of having an automated and manageable enterprise security system. This researcher defines Threat Management as; a centralized system that incorporates system logs from vendor independent devices and provides a correlated and calculated representation of security events that may pose a threat. Before we delve into the details of what makes up a Threat Management System (TMS) we first need to define a couple of terms. The first term that will be used is Aggregation. Aggregation is simply the term used to describe collecting and normalizing diverse information from multiple sources. An example would be collecting log information from a firewall and IDS system to a central server. Even though the information collected from the devices is different, it is now accessible from one

database. The other part of aggregation is normalizing the data. This entails taking different vendor brands for the same device type and creating a unified data format. What we then have is one data format for all devices in that category. Correlation is the process of grouping like information together. At the most basic level it would be matching a firewall packet with IDS alert, but correlation can be much more complex than that. Now that we have defined our key terms take a look at figure 1.

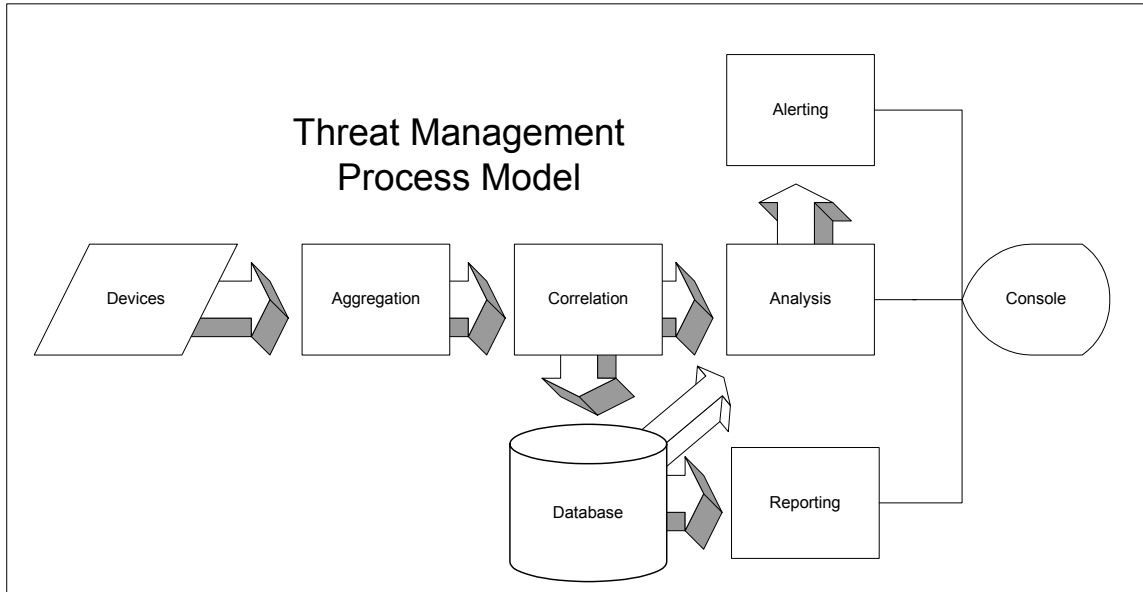


Figure 1

The diagram depicts the information flow through a Threat Management System. We are going to breakdown each process and establish the relationships between the different processes.

Devices

The devices represent the logging inputs to the system. These devices can be firewalls, routers, NIDS, HINDS, hosts and application logs. All the devices talk directly to the aggregation module of the system. A good Threat Management system will support numerous device types and a plethora of vendors.

Aggregation

This is where all your devices log their information. Depending on the device there are different ways in which the information is pulled. Information can be collected via SNMP, Syslog, Agents and other proprietary protocols. This is a critical module for any enterprise deployment. Depending on the amount of logging information collected, and the number of devices polled you will want to insure the Threat Management system can scale properly.

Correlation

There is a close relationship between correlation and analysis. The main difference is that the analysis process provides a rule-based interface and the ability to include non-aggregated information in the calculation process. The root responsibility of the correlation engine is to process events and find relationships between them. The actual correlation can be based on a combination of time stamps, source or Destination addresses, and other logical relationships established between the logging information. Again, this can be as simple as matching a firewall, router and IDS alert and translating it to a correlated event. Correlated events are logged events that are group logically by a relationship. This function alone will provide you with great insight as to where the packet came from and where it actually ended up. On the more complex side of correlation it can recognize patterns between devices. For example, Host A has failed to login to 6 different machines which generated six different syslog alerts. Now unless you had all this information in one place you would have mostly likely ignored the one failed login attempt, but when this information is presented in a correlated fashion it is obviously apparent that there is a problem.

Analysis

Here is where the real pieces of the puzzle are put together. The analysis process will take the correlated information and compare it to non-aggregated data. This non-aggregated could be any of the following:

- Asset database
 - Criticality of host
 - Identify network services
 - Relationship to other hosts

- Contact information
- Vulnerability information
 - Host vulnerability

You might be thinking why would I want an asset database mixed in with my Intrusion Detection System. First and foremost it will help you in eliminating false positives, which will enable you to focus real threats. This can be easily achieved by including what services a specific host is running in the asset database. So for example, let's say you get a, "WEB-IIS _mem_bin access" reported by one of your IDS sensors and then the analysis module checks the asset database to find out what services are running on the box. The analysis module reports back that the system is a Lotus Domino web server. This may be an attack, but the system is immune to the exploit and will not alert you to the event. The event is recorded for trending though. This example is rather simple, but demonstrates the concept. In reality the asset database would have to include application and operating system version numbers to distinguish between vulnerabilities. Figure 2 visually demonstrates the logic.

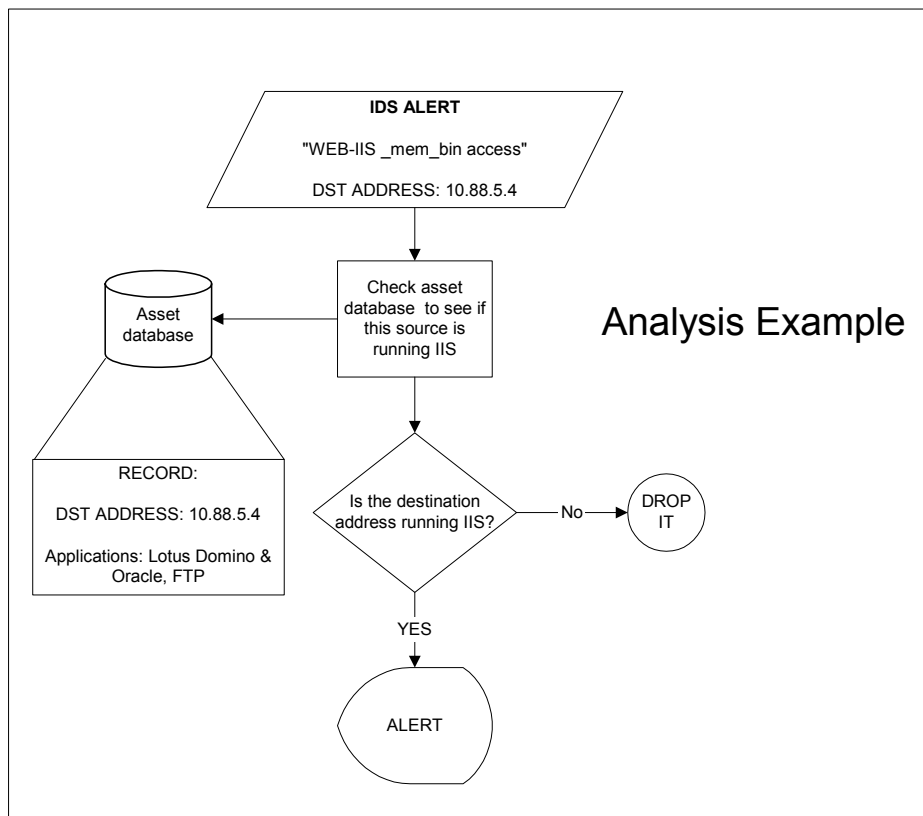


Figure 2

After reviewing the diagram it should be clear how intelligent your system can become by including just the services that a particular host is running. You may also be thinking I can accomplish this by putting filters on my IDS system. You could do this, but it can become unmanageable and the IDS administrator is forced to maintain the filters. Depending on how the asset database is implemented the work of maintaining what services are running will be the server administrators or an automated process. Another great field that can be added to the asset database is criticality. Criticality describes how critical this asset is to your operation. The same logic as above could be used to rank events by criticality, thus allowing you to focus on most important hosts first.

Now let's expand the asset database to include contact information and how this host is related to other hosts. Wow! Now at the click of a button you have all the information necessary to address a threat. The options as to what to include with the asset database are endless and provide a great wealth of information in a central location.

The last piece of critical information that can also be included in the asset database is vulnerability information. Many security departments have the responsibility of performing regular vulnerability assessments. Since this is normally a separate function wouldn't it be nice to include this information in the asset database. This information could be used to validate a specific vulnerability targeted towards a specific host. Once again allowing you to focus on real threats.

The analysis engine also contains a very powerful rule language to allow the system to be customized to the environment. Rule building is nothing more then creating custom pattern recognition as discussed in correlation, or they can be simple rules to monitor critical assets. The rules that you create analyze information from a building top view or just watch for specific activity down on the street.

Alerting

The alerting process provides the mechanism for bringing attention to certain events. This can be provided through a visual display highlighting troubled areas with the color red, sending an email or paging the analyst on duty, or in some cases the system can take corrective action.

Reporting

Reporting is critical to identify long term trends and provides management with palatable information. Reporting on these systems comes in many shapes and sizes. Some systems concentrate on the graphical approach which is easy on the eye, while another's provide more textually detailed reporting. Most of the systems allow you to create customizable reporting using a built-in reporting system or third party report writing software like Crystal Report. The key is to find a good balance between management reports and technical reports.

Evaluation Criteria

When evaluating vendors in this space its critical to have grading criteria that is generic to the products in this space. This is important because each vendor takes a different approach in providing a Threat Management Service. The following are 8 categories of questions that can be used when evaluating the offerings of each vendor.

I. Management Interface

Does the product allow various operating systems to access and manage the product? E.g. Windows, Unix, Linux, Web
Can the status views be customized to meet our needs? The system will be placed in a NOC with multiple heads up displays.
Does the product provide object orientated access controls and to what extent?
Does the product provide descriptors for the various sensor devices (Firewall, IDS, routers) in our network? This could be information on where the device is and what its monitoring.
Does the product contain or have the ability to include forensic tools? E.g. NMAP, whois, etc.
Does the product provide multiple ways of alerting personnel to events? E.g. Email, ticketing systems, page, etc.
Does the product provide a way of ranking security events?
Does the product include a knowledge base to help define an event, and provide solutions for correcting it?
Does the product display real-time events?
Does the product provide the ability to segment groups of devices? E.g. We have a special group of devices it does not want included with its normal operations. The devices report to the same system, but access to the information is limited and the devices have their own metrics. This segmentation could be based on geography, responsibility, or functional group.

II. System Architecture

	Does the product provide the ability to collect information from multiple event collectors while keeping the information in a centralized location?
	Does the product provide a caching mechanism when part of the system is unavailable?
	Will the prospective company provide our company with the database schema to their product?
	Does the product scale to a multiprocessor architecture? E.g. Multi-threaded
	Does the product provide any API's for customization?
	Does the product keep intact the original log information?
	Does the product provide a method to prevent data tampering? E.g. event signing.
	Does the product provide any tools for managing the database? E.g. Retention policy, purging and merging. Are they scripts or GUI based tools?
	Does the product provide the ability to filter certain events from the console and the database?
	Does your company provide benchmarking information on the product?
	Does your product scale to an enterprise class environment? Describe the way in which your product can be scaled.
	Does your product incorporate a self-monitoring capability such as status of the sensors, agents, server, and database?
	Describe the products auditing and logging capabilities?
	Describe any response mechanisms the product may have?
	Describe the recommended configuration for a high availability operation, ie (24x7 NOC) in a geographically disperse organization.

III. Third Party Vendor Support

	Does the product aggregate information from firewall devices? Which ones?
	Does the product aggregate information from Intrusion Detection devices? Which ones?
	Does the product aggregate information from routers? Which ones?
	Does the product support aggregate form mainframe logs?
	Does the product aggregate information from host based intrusion devices? Which ones?
	Does the product aggregate information from host and application logs? Which ones?
	Does the product provide an API for developing custom aggregation agents?
	Does the product aggregate information from anti-virus systems?
	Does your product incorporate IDS payload information in to the database? E.g.

	Snort
	Describe how you attain and maintain interoperability with the various IDS, firewall, IDS, and router vendors?

IV. Correlation Engine

	Does the product correlate related events in to a single event? E.g. This IDS alert is related to this firewall packet or multiple IDS events consolidated in to a single event.
	Does the product correlate on time stamps, source and destination address, event types and device relations? Explain the products correlation engine.
	Does the product keep all fields of the original devices logging information during the normalization process?

V. Information Analysis

	Does your product include the incorporation of any non-aggregated information sources? E.g. Asset database or vulnerability knowledge bases, etc.
	If your product does include an asset database, what type of fields does the system track? E.g. Criticality, running services, applications, host owners, etc...
	Does the product provide the ability to calculate a threat based upon non-aggregated information?
	Does the product provide the ability to create custom rules for alert notification?
	Does the product include any type of intelligent pattern recognition logic?

VI. Reporting / Information Gathering

	Does the product provide canned reports for technical and management level personnel?
	Does the product provide the ability to create custom reports? Describe the report creation process?
	Does the product provide any type of visualization of the information collected?
	Can all data fields be queried against?

VII. Documentation and Integration

	Does your company provide documentation of the product? Electronic or paper?
	Does the company provide integration services? Does the product vendor or a third party provide this service?

VIII. Misc

	Due to the immaturity of the space will your company provide source code escrow?
	Does company have the ability to provide long-term maintenance agreements with fixed pricing?
	Does your company have any enterprise licensing strategies?

Concerns

There are few issues associated with these systems that you should be aware of that can easily be overlooked by the hype and glitz that surround these products. The first item that you need to be aware of is the storage requirements for operating one of these systems. Depending on the number of devices, and your companies retention policy this can easily grow into the terabyte range. The cost of the storage is area of concern, but more importantly you have to take in to account the cost of backing up and maintaining a database of this size. The second area of concern is the hardware requirements. These systems require lots of memory and processing power for accomplishing the expected result. This again is not just the cost of hardware, but maintaining a complete TMS infrastructure. Most of these systems are implemented in a tiered architecture of event collectors, consoles, and database servers that can span multiple locations. Then if your offering this as a 24x7 service you better have some redundancy capabilities.

The last area of concern is the viability of the companies that occupy this space. This is a major concern for any company that plans to use one of these systems as the heart beat of the operation. Most of the companies that offer products are venture funded. It is absolutely critical that you get references and investigate the financial health of the vendor. As part of any

deal that you sign I would recommend having the source code escrowed. This will protect you somewhat in the case the company goes bankrupt and if you're a large enough of company it may be cheaper for your company to maintain code yourself.

Conclusion

Everything this researcher has discussed is based upon meetings with the vendors that occupy this space, and this researcher's opinion on what constitutes a Threat Management System. Some of the functional details I discussed are not available in any products at the time of this writing, but should be in the near future. The market is still very immature, by means of products and the viability of the companies, but is definitely on the right track of addressing a very important business problem. What you will find when searching for one of these vendors is that each one takes a different approach to solving the same problem. Some vendors are very strong in one area and weak in another. When evaluating these systems take your time in defining exactly what functional issues are going to be addressed by the solution. This way, you'll get exactly what you paid for and probably a little more.

References

ArcSight, <http://www.arcsight.com/>

NetForensics, <http://www.netforensics.com/>

e-Security, <http://www.esecurityinc.com>

Intellitactics, <http://www.intellitactics.com/>

GuardedNet, <http://www.guarded.net/>

Tier-3, <http://www.tier-3.com>