

IDS Correlation of VA Data and IDS Alerts

Neil Desai 2003-06-30

The sky is falling! The sky is falling! Okay, well the sky isn't really falling, but isn't that the way that we all felt the first time we installed a NIDS and turned it on? We watched the alerts fly by the screen quicker than we could determine what they were. If we were lucky we could just make out what colors the alerts were. Unfortunately that stigma has stuck with the intrusion detection industry. Some people who have NIDS installed have just ignored their screens and been happy with telling the auditors: "Why of course we have intrusion detection. We use <insert brand here>"

The intrusion detection industry has matured over the last few years and most IDS vendors have tried to address the issue of false positives. Between the vendors creating more intelligent NIDS, IDS analysts learning more about how NIDS work and IDS analysts tuning their NIDS, the amount of alerts that are generated are getting to a point where they no longer overwhelm the analyst that is monitoring the events. Yet one of the questions that still looms is "How do you classify and/or group your alerts?"

When signatures are created they are assigned a default alert level based on the severity of the exploit. If you have all HTTP signatures turned on even though you only run Apache should all those Code Red alerts be classified as a high? What if you run an all IIS shop and all your servers are patched, does Code Red still rate a high in your environment?

In the Beginning

When IDSs first came out they all did the same thing, for the most part, and that is match the current packet against a known set of exploits. Enter Dug Song and `fragrouter`. This tool was the first widely known NIDS evasion tool. While Thomas Ptacek and Timothy Newsham wrote about NIDS evasion almost a year earlier it was this tool that took one of their ideas and turned it into reality. Now anyone with a *NIX system could perform NIDS evasion. The evasion was pretty simple, and that was to use IP fragmented packets to break the exploit code up into multiple packets. Most NIDS did not do any type of reassembly [1] so when they attempted to match a signature against only part of the attack they missed the attack. Only a couple of IDS vendors were able to deal with this new (at that time) type of attack/evasion. All the vendors who were susceptible to `fragrouter`'s techniques quickly worked to patch their systems.

Almost three years later Dug released *fragroute*. This utility expanded on the paper from Ptacek and Newsham. Now simple IP reassembly was not enough. NIDS now had to also be aware of each stream and reassemble the packets at layer four. Seems simple enough to fix right? Not quite. RFCs are only a guideline as to how a protocol should be implemented. They are not written in stone so everyone who writes a protocol based on an RFC may have their own little nuances to their version. Suppose the NIDS sees a TCP packet that is being retransmitted for whatever reason. How should it be treated? Should the NIDS favor the new packet or the old packet? Under normal circumstances the data in both packets will be identical, but we are dealing with NIDS evasion so they most likely won't be. So it is more critical to reassemble the packets correctly. But what is the proper way to reassemble a TCP stream? It all depends on the implementation of TCP/IP. This is still one of the issues with NIDS today. Since the NIDS is sitting in front of the host, it is not fully aware of how the end system is going to handle the packets. If it handles them differently it could generate a false positive or even worse, miss the attack.

IDS vendors have been trying to find ways to be more aware of the hosts that they are protecting so that they can treat packets just as the host would treat them. To make this a viable feature manual configuration was not an option. In looking for new ways to be aware of host operating systems, a couple of vendors came up with a couple of new ideas.

The Perfect Blend

When it comes to technologies that are specific to information security most of them are point technologies. By that I mean that each of them work to secure the company's systems without knowledge of each other or the information that they have. With most vendors only offering a specific technology (i.e. firewall, IDS, virus protection) or product, interoperability has been almost non-existent [2]. SIMS (Security Information Management Systems) were suppose to be the technology that we were all waiting for, but they don't offer the robustness that is needed. These products are suppose to be able to correlate information across many different product lines giving the end user the ability to correlate information gained from disparate systems.

As information security professionals, two of the technologies that we use will look for the same information but from different perspectives. Vulnerability assessment tools look for vulnerabilities in a proactive manner while intrusion detection systems look for vulnerabilities as hackers try to exploit them. Until recently we scanned our systems with the vulnerability assessment tools to make sure that we were secure, and when we looked at

This works great if you are a small shop or are the guy who does both the vulnerability assessment and monitors the IDS events, but unfortunately most information security groups are not built this way. So the group that monitors the IDS events has no real way to assess the threat of a particular attack. All they can do is consider the level or priority that the signature is assigned by default.

Now there are a couple of vendors [3] that can correlate this information in near real time. Not only are the events from the IDSs correlated with information from the vulnerability assessment tools, but the alerts are categorized into a more manageable way. Before our IDS consoles were based on near real-time delivery of each individual event. This worked a couple of years ago when attacks were not as frequent and worms were not that prevalent. Today worms scour the Internet looking for vulnerable hosts. When a host is found an attack is launched and in a matter of seconds a worm has just lit up your IDS console with more events than you can decipher. By having more correlation of events the data that is presented is not as abundant as before, but it is more meaningful.

So you may be wondering: isn't this technology just a slimmed down version of SIMS? Not really. SIMS focus on correlating data from many different systems (IDS, firewalls, routers. etc.) which may make them slower in the correlation process due to the large amount of logs that need to be searched for an attack pattern. Alert management systems only worry about taking your IDS alerts and seeing if the host that is being attacked is vulnerable to that particular attack, and adjusting the priorities of the alerts and responses based on the outcome of the correlation. You will still have your false positives, but the alert management systems will help prioritize your events so that the false positives are put at the bottom.

Alert Management

Whether we like it or not information security is about managing risks not enforcing as much security as we can get away with. If we were to lockdown our systems as security purists we would end up in the same situation as the CIA [4]. Our job is to bring the risk to a level that is acceptable for the company. The correlation of vulnerability assessment data with IDS data can therefore be considered an alert management system.

Even though IDS vendors have improved their products there are still going to be false positives and false negatives. Alert management systems will help in reducing the false positives. The IDS will still detect an attack that is not really an attack (false positive), but the correlation system will validate the alert against the

information from the vulnerability assessment data and determine the level of threat that a particular event really is. There is still some configuration that needs to be done on the management system so that it can deal with the events in a manner that suits one's needs. The management systems have the ability to group attacks that correspond to the way an attacker would conduct a focused attack. Let's look at any one of the worms that have come out over the last few years and see how the alert management system would deal with it and compare it to how the regular IDS console would deal with it.

The first step of the worm is to scan for a particular port on which the vulnerable software would normally reside. In the case of Code Red and Nimda this was TCP port 80.

1. **Regular IDS console** The amount of data that is presented will depend on the following factors:
 - o Number of IDSs that are deployed. If an attack crosses multiple IDSs then a single packet will cause multiple alerts to be sent.
 - o Number of publicly available address that are being protected. A single attacker scanning a class B address can generate a lot of alerts.
2. **Alert Management Console** A single incident will be displayed to the analyst. The amount of IDSs that detect the scan or the amount of addresses are scanned does not determine the amount of data presented to the analyst. Since this is inline with a reconnaissance technique, some basic data about the scans gets stored for future correlation with other events. The alert management system will now wait for alerts that correspond to the attacking host, the hosts that were scanned and the ports that were scanned.

The second step of the worm is to attempt to exploit the hosts that it discovered during step one.

1. **Regular IDS Console** Again the amount of data that is presented to the analyst will be determined by the same factors as above. For every packet that contains exploit code an event will be sent to the console. With an attack line Nimda, where multiple UNICODE attackers were attempted against each host, the number of events will be numerous regardless of IDSs reporting on the attack.
2. **Alert Management Console** All the alerts from the attack will be consolidated into a single alert. The only thing that will change is the number of times that the alert has been tripped. The correlation system will match up the reconnaissance scan with the attacks. Depending on the information in the vulnerability database the threat level of the alert may change.

The next step will only occur if one of the hosts is compromised as a result of the worm:

1. **Regular IDS Console** The amount of data will be directly related to how many hosts are compromised. There may be so much data presented to the analyst that these attacks get lost in all the other data.
2. **Alert Management Console** The console will raise the alerts to the highest priority and respond as configured. All three events are correlated together. Instead of all the different events being show to the analyst a minimal amount of data is show to the analyst. The count for each event replaces the need for each event to be shown. At this point a real threat has been detected and the data is presented to the analyst in a manageable fashion. Below are screen shots of two different products, ISS's SiteProtector 2.0 with Fusion 2.0 and Tenable Network Security's Lightning Console.

In Figure 1 below, all similar alerts are organized into a single event shown to the analyst. Only the event, source, target, and object count are updated. The Fusion 2.0 module correlates the information between the NIDS (RealSecure) and the vulnerability assessment tool (Internet Scanner) and updates the "Status" column. The analyst can then prioritize the investigation of alerts based on the probability of the success of the attack. Also the event "HTTP_Code_Red_II" has been classified as a low as opposed to the default, high, because of the correlation.

Security Alerts						
Tag Name	Status	Severity	Event Count ▾	Source Count	Target Count	Object Count
HTTP_HTDIG_htsearch	Unknown impact (no correlation)	High	4	1	1	1
IIS_Reveal_Address	Unknown impact (not scanned recently)	Low	4	1	1	1
HTTP_Favorites_Icon_Overflow	Unknown impact (not scanned recently)	High	3	2	3	1
HTTP_URL_dotpath	Unknown impact (no correlation)	Low	3	1	1	1
Netbios_Session_Granted	Unknown impact (not scanned recently)	Medium	3	2	1	1
Email_Mime_Filename_Overflow	Unknown impact (not scanned recently)	High	3	2	2	1
Fragment_Differential_Overlap	Unknown impact (not scanned recently)	Medium	3	3	3	1
Disk_space_shortage	Unknown impact (no correlation)	Medium	3	1	1	1
HTTP_Put	Unknown impact (not scanned recently)	Medium	3	2	1	1
Y3K_UDP_Response	Success likely (target vulnerable)	High	2	1	1	1
HTTP_IIS_Unicode_Translation	Failed attack (blocked at host)	Low	2	2	1	1
HTTP_Code_Red_II	Failure likely (wrong OS)	Low	2	2	1	1
Email_Encap_Relay	Failure likely (wrong OS)	Low	2	2	1	1
Email_Vrfy	Unknown impact (not scanned recently)	Medium	2	1	1	1
POP_Filename_Overflow	Unknown impact (not scanned recently)	High	2	1	1	2

Figure 1. ISS SiteProtector 2.0 with Fusion 2.0.

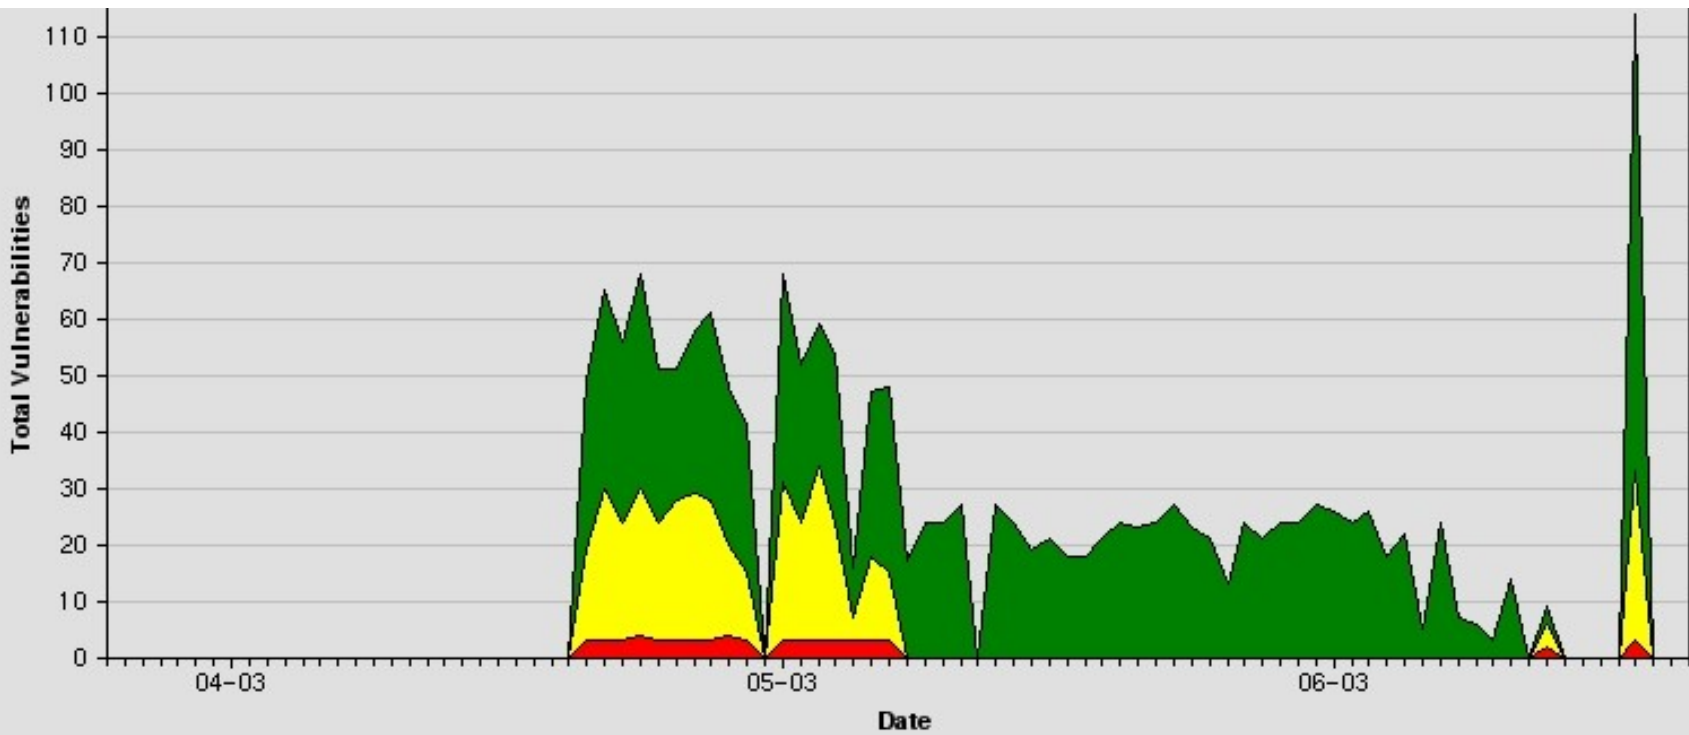
In the next example, Figure 2 shows that all the alerts from the IDS (Dragon, Snort, ISS, Bro) are correlated with information from the vulnerability assessment tool (Nessus). Any attacks that are more likely to be successful are flagged so that the analyst can readily identify them and take appropriate action.

Time	Event	Source	Destination	Vuln	Sensor	Type
03:01	(spp_portscan2)	192.168.0.4:3347	192.168.0.5:21 [6]	no	Destroyer	Snort
03:02	(spp_stream4)	192.168.0.4:59117	192.168.0.5:21 [6]	no	Destroyer	Snort
03:02	(spp_stream4)	192.168.0.4:59117	192.168.0.5:21 [6]	no	Destroyer	Snort
03:02	TCP-FLAGS	192.168.0.4:59115	192.168.0.5:21 [6]	no	Rodan	Dragon
03:03	(spp_stream4)	192.168.0.4:59117	192.168.0.5:21 [6]	no	Destroyer	Snort
03:03	TCP-FLAGS	192.168.0.4:59115	192.168.0.5:21 [6]	no	Rodan	Dragon
03:03	FTP:USER-ANON	192.168.0.4:2062	192.168.0.5:21 [8]	YES	Rodan	Dragon
03:04	(spp_stream4)	192.168.0.4:21228	192.168.0.5:21 [6]	no	Destroyer	Snort
03:04	(spp_stream4)	192.168.0.4:21230	192.168.0.5:21 [6]	no	Destroyer	Snort
03:04	(spp_stream4)	192.168.0.4:21231	192.168.0.5:21 [6]	no	Destroyer	Snort
03:04	TCP-FLAGS	192.168.0.4:21230	192.168.0.5:21 [6]	no	Rodan	Dragon
03:04	FTP:BAD-LOGIN	192.168.0.5:21 [6]	192.168.0.4:2070	no	Rodan	Dragon
03:04	FTP:USER-ANON	192.168.0.4:2084	192.168.0.5:21 [8]	YES	Rodan	Dragon
03:04	FTP:LINUX-NUL-USER	192.168.0.4:2106	192.168.0.5:21 [6]	no	Rodan	Dragon
03:04	FTP:LINUX-NUL-PASS	192.168.0.4:2106	192.168.0.5:21 [6]	no	Rodan	Dragon
03:04	FTP:BAD-LOGIN	192.168.0.5:21 [6]	192.168.0.4:2106	no	Rodan	Dragon
03:04	FTP:BAD-LOGIN	192.168.0.5:21 [6]	192.168.0.4:2107	no	Rodan	Dragon

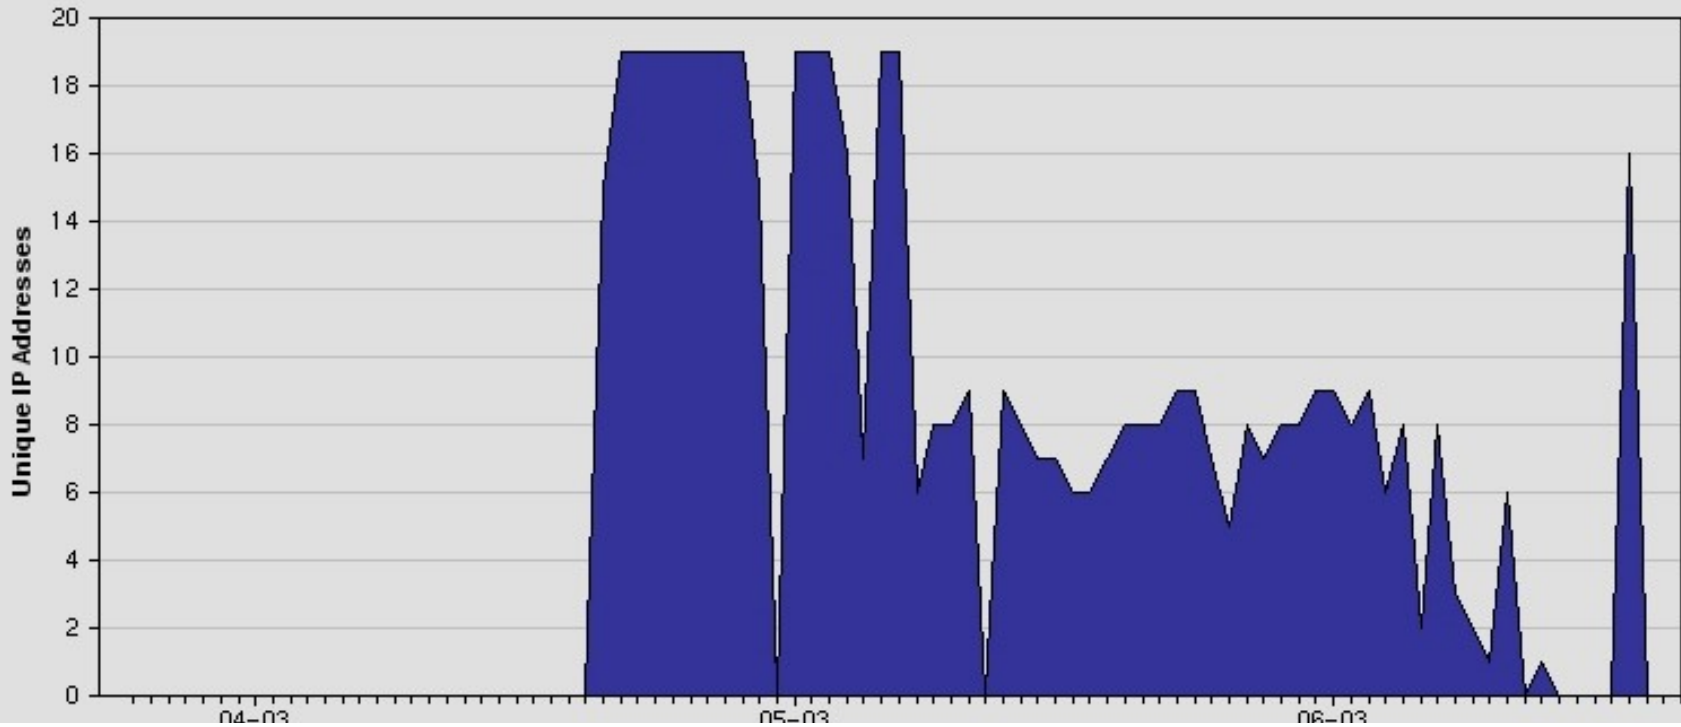
Figure 2. Tenable Network Security Lightning Console.

Next in Figure 3, below, you see the analyst trending information. This can be used to determine if an major attack is imminent or a pattern of attacks.





Total Unique IP Addresses for last 90 days



Date

Figure 3. Tenable Network Security Lightning Console.

The next example, Figure 4, shows a different type of trending chart based on attacks that are considered high.

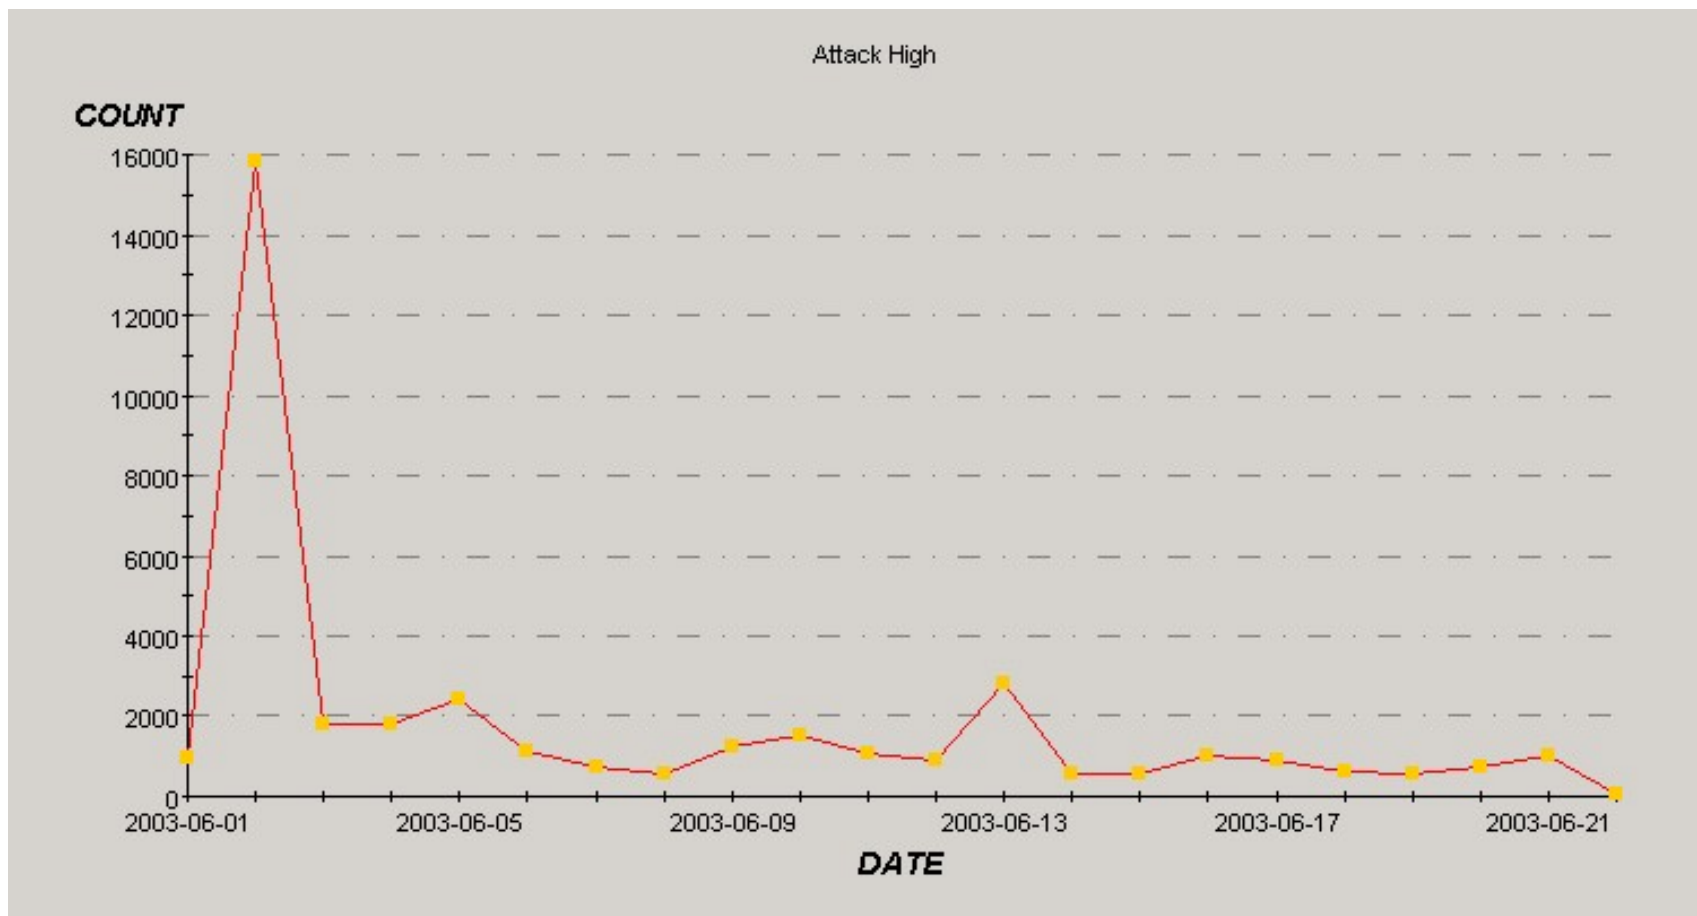


Figure 4. ISS SiteProtector Enterprise Dashboard.

The value of this type of system is apparent when dealing with hundreds of events as part of a normal day.

Currently only the IDS and vulnerability assessment data is correlated, but soon other logs will be part of the correlation. This is when the alert management systems will deliver what SIMS have not done.

Eventually information security products will have the integration that network management tools currently have.

Conclusion

The number of attacks that are taking place is increasing while the budget to deal with them may either be staying the same or decreasing. Adding additional staff to manage the security needs of a company may not be an option. As usual you have to do more with less. Alert management systems will cut down the amount of time that it takes to sift through data. More false positives are weeded out and the data that is presented to the analyst has been correlated with other information and the priorities have been adjusted accordingly. Now the analyst can properly prioritize the events that need to be investigated. The initial configuration is crucial and can be daunting, but the payoff is great.

Relevant Links

[1] <http://archives.neohapsis.com/archives/ids/1999-q4/0189.html>

[2] Vendors like CheckPoint offer integration with their products via APIs that are released.

[3] Currently the only two vendors that the author knows of that do with with system-specific data and provide near real-time events are Tenable Network Security and ISS.

[4] Report: Too Much Cyber Security at the CIA <http://www.securityfocus.com/news/5201>

[Privacy Statement](#)

Copyright 2006, SecurityFocus