# Intrusion Detection Systems:

## A Feature and Capability Analysis


Sig Myers
University of California, Santa Cruz
Research Assistant
Computer Science
sig@soe.ucsc.edu

John Musacchio
University of California, Santa Cruz
Research Advisor
Technology & Information Management
johnm@soe.ucsc.edu

Ning Bao
University of California, Santa Cruz
Research Assistant
Computer Science
nbao@soe.ucsc.edu



## ABSTRACT

Network security is an ongoing concern for many businesses, governments and individuals looking to protect their information assets. This work aims to address the capabilities of current generation intrusion detection/prevention systems with a specific focus on metrics of interest to the "Game Theoretic Approaches to Cyber Defense" research being headed by Dr. Musacchio at the University of California, Santa Cruz. Further, it is also a hope that this paper will serve as a reference to those interested in understanding what knowledge can be extracted from and evaluated by intrusion detection/prevention systems.



## Categories and Subject Descriptors

D.3.3 [**Programming Languages**]: Language Contructs and Features – *abstract data types, polymorphism, control structures.* This is just an example, please use the correct category and subject descriptors for your submission. The ACM Computing Classification Scheme: http://www.acm.org/class/1998/

## General Terms

Your general terms must be any of the following 16 designated terms: Networks, Security, Intrusion Detection Systems, and Intrusion Prevention Systems


## 1. INTRODUCTION

### 1.1 What is an Intrusion Detection System?

An Intrusion Detection System (IDS) is tasked with delivering, in real time, a reliable analysis of a given networks traffic and deeming whether a computer system, network or information asset is being (or has been) attacked. An Intrusion Prevention System (IPS) is, at its core, an IDS that automatically attempts to block or subvert an attack (or sequence of attacks) on a system. IPS' can protect systems by blocking connections, disabling services, or other pre-defined activities in order to minimize human intervention and system/network damage while simultaneously allowing any legitimate traffic. Both IPS and IDS solutions generally offer a wide range of reporting features to their users so that a visualization of a network's activity and health can be assessed. Additionally, when alerts are triggered, automated e-mails, SMS text messages and phone calls are able to be sent out by the IDS/IPS systems and events of the attack logged to a database. This functionality runs standard in open-source software and commercial appliances alike.

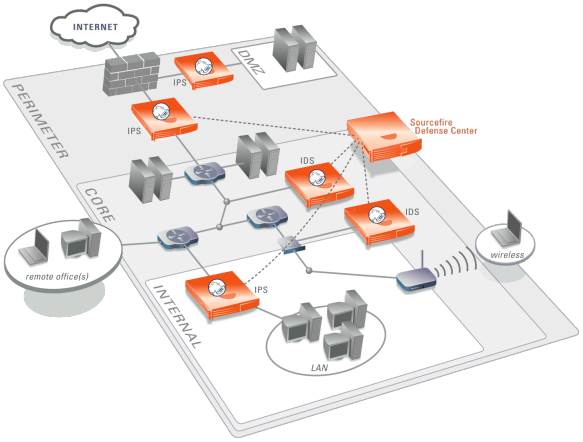


**Figure 1 – An example of placing IPS nodes throughout a network's infrastructure in order to heighten security.**

Originally, IPS and IDS systems were placed on the edge of a company's network and monitored all incoming and outgoing connections to the Internet. Later, due to security issues, it was realized that a network edge-based IDS/IPS solution was not enough. Traffic flowing within an organization was just as important to monitor due to insider attacks and worms, which can unknowingly wreck havoc on a network and lead to serious security concerns. This led to the adoption of "nodes" which are placed strategically throughout a network to identify traffic flowing in and out of definable sub-networks in an organization. A consortium of nodes which report to an edge-based IDS seem to be the popular setup of large entities such as corporations and governments, and they are well supplied by nearly every IPS security vendor in the U.S.

### 1.2 Why do we need Intrusion Detection/Prevention Systems?

In order to understand why such systems are necessary it is pertinent to understand what kind information a successful attack can gather and how it can harm an organization. Computer worms, viruses and trojans are generally well-known terms by the public, but their capabilities often misunderstood. These malware programs have the capability to compromise a system to such a state that data can be destroyed or modified, services can be disabled and multiple computers can become infected with or without participation from a user. In some cases when a computer is compromised, an attacker can remotely utilize system resources (such as in a botnet) or upload more malware (such as keyloggers) in an attempt to steal passwords and other valuable data from the compromised system.

One can quickly understand why an IDS is necessary—the possibilities of an attacker finding a software vulnerability, outdated operating system, or luring a user into executing malicious code are quite plausible scenarios in many respects.

```
Bro Report                                                    Organizatio
================================================================================
Summary                          July 28, 2004 17:01 to July 29, 2004
================================================================================
 Incident       Likely Successful          1
 Summary        Unknown                    0
                Likely Unsuccessful        0
                Scans                     10

 System         Bro disk space:   <% at time of report generation>
 Statistics     Bro Process cpu:  <time>
                Bro restarts:     <date/time>
                System reboots:   <date/time>

 Traffic        Number of packets:       <count>
 Statistics     Number of valid packets: <count>  <% of total>
                Protocol summary
                Http: <count>   <% of total>
                SSH : <count>   <% of total>
                SMTP: <count>   <% of total>
                Etc.
                Average bandwidth:
                Peak bandwidth:
================================================================================
```

**Figure 2 - A Bro incident report summarizing metrics and system information for a 1-day period.**

From a network security administrator's point of view, securing a network can be quite daunting. However, with either a state-of-the-art security appliance or open source software solution, the challenges posed to the security administrator can become much more transparent and manageable. Figure 1 shows an example of how Sourcefire, a leading IPS vendor, aims to deliver a more secure network by placing IDS "nodes" throughout a network's infrastructure.

## 1.3 Metrics and Capabilities of Intrusion Detection/Prevention Systems

In order to evaluate an IDS/IPS we looked for certain metrics that were of particular interest to creating a game theoretic approach to the network and cyber security research being conducted. Some capabilities of IDS/IPS systems were irrelevant, such as lines speeds (often referred to as throughput and ranging anywhere from 50Mbps to 10Gbps or more) and the number of user sessions able to be monitored (generally ranging in the millions). Of more interest is to understand what sort of information can be collected about a particular attack or attacker. We wanted to know if it is possible to identify:

- How long an attacker/attackers were in the system
- How long it takes to detect an attack from it's onset
- Is there is a return rate of the attacker(s)
- If there is a confidence interval to let us know that there is an attack occurring, even if no signature exists for the given attack yet
- How close the system is to crossing a threshold and alerting us to an attack
- What is the overall strength of the access control system
- What is the IDS/IPS overall strength

We will identify via screen shots, packet captures and analysis reports how each of these metrics are obtained within different IDS/IPS systems, and then provide a comparison chart of all the systems reviewed for maximum ease-of-use and understanding.

# 2. Intrusion Detection Systems

## 2.1 Bro Intrusion Detection System

Bro is an open-source network intrusion detection system, which lends itself particularly well to forensic tasks due to its great data collection and analysis capabilities. Bro is a signature-based IDS, meaning that it attempts to match a signature to network traffic in order to find the 'attack.' Bro is unique in that it utilizes regular expressions, rather than fixed strings, to understand network activity. The creators of Bro translate this to mean that a lower false-positive rate is achieved due to Bro being able to understand the context of the traffic, rather than merely matching a static signature.

Bro also comes with it's own language which advanced users can utilize to program policy scripts. Policy scripts allow network administrators to fine-tune their Bro installation in order to specifically search out certain types and patterns of traffic, and define them as malicious. Further, developers can extend Bro's capabilities by having scripts execute in certain events to block, alert or log information about certain network traffic. Some of Bro's biggest shortfalls (or selling points in some scenarios) are that it only reports information to log files and does not have a graphical user interface (GUI). Log files are designed in such a way that humans can understand them and computers can easily parse them. While the option to report events to a database might be nice in some cases, especially for long-term storage of data, it is not an absolute necessity for Bro to be a worthwhile network security investment. The lack of a GUI is understandable given Bro's preference towards forensics and analysis rather than intrusion prevention techniques.

To further understand some of the capabilities and metrics captured by Bro, we refer to Figures 2, 3 and 4. Figure 2 offers us information regarding the report analysis time-period, in this case from July 28, 2004 to July 29, 2004. We also are able to see that ten scans (i.e. port scans) occurred and one incident (attack) was likely successful. The traffic statistics section displays the total number of packets and breaks that amount down by the traffic type, while also indicating the average and peak bandwidth that occurred in the given time period. Figure 4 elaborates on the summarized information and gives us a more detailed transcript of the incident events. We are able to see in this portion of the transcript the attacker's IP address, the target and what alarms and attacks were used.

```
Remote Host Connection History (all successful/unsuccessful to site)
      24 hrs     |      3 days      |      7 days      |      30 days
-------------------------------------------------------------------------------
       14/10     |        0/0       |        0/0       |        0/0
-------------------------------------------------------------------------------
  Total since remote host first seen on 07/29/04: 14/10

================================================================================
Scans
================================================================================
==
Date Dropped            Host                                    Port Scanned
--------------------------------------------------------------------------------
Jul 29 13:14 n219077002119.netvigator.com                        (3128/tcp)
Jul 29 13:23 node1.lbnl.nodes.planet-lab.org                     (49702/tcp)
Jul 29 13:30 213-145-189-50.dd.nextgentel.com                    (4899/tcp)
Jul 29 13:32 211.55.52.67                                        (1034/tcp)
Jul 29 13:52 user-69-1-11-116.knology.net                        (3128/tcp)

********************************************************************************
```

**Figure 3 - Bro's remote host connection history and port scan list is the last information displayed in the incident report.**

```
Incident     ORGCODE-000002                                LIKELY SUCCESSFUL
---------------------
Remote Host: 84.136.138.21   p54877614.dip.hacker.net
 Local Host: 124.333.183.162 pooroljoe.dhcp.org.com

Alarm(s) 1 MS-SQL xp_cmdshell - program execution
           Jul 29 12:43 84.135.118.20 -> 128.3.183.62
         2 TFTP Get Runtime.exe
           Jul 29 12:43 128.3.183.62 -> 84.135.118.20

Connections (only first 25 after alarm are listed)
-----------
                  time      byte   remote       local    byte
 date   time    duration  transfer  port  type   port  transfer  protocol
----- -------- -------- --------- -----  ---- ------ --------- ----------
07/29 12:43:31        ?     566 b  4634  1  >   1433      467 b  tcp/MSSQL
07/29 12:43:31        0         ?  2318  2 <     69       20 b  udp/tftp
07/29 12:43:32    265.7       4 b  4638  * <   2318      3.0kb  udp
07/29 12:48:56        ?         ?  4640    >   2362          ?  tcp
07/29 12:50:05        ?    11.4kb  4639  * <   3333      8.6kb  tcp
07/29 12:53:00        0         ?  4684  *  >  2362          ?  tcp
07/29 12:53:07        ?         ?  4685  *  >  2362          ?  tcp
07/29 12:53:59        ?         ?  4689  *  >  2362          ?  tcp
07/29 12:54:14      6.1         0  4693  * <   2380     94.2kb  tcp
07/29 12:54:21       .5      50 b  4694    >   2381          0  tcp
07/29 12:54:23       .7         ?  4695   <    2382          0  tcp
07/29 12:54:25       .5      51 b  4696  *  >  2383          0  tcp
07/29 12:54:27       .5      61 b  4697  *  >  2384          0  tcp
07/29 12:54:28       .7      39 b  4698    >   2385          0  tcp
07/29 12:54:31       .5      41 b  4699  *  >  2386          0  tcp
07/29 12:54:33      1.2    4.9 kb  4700    >   2387          0  tcp
07/29 12:54:35     12.8  195.0 kb  4701  * <   2388          0  tcp
07/29 12:54:53       .2         ?  4703   <    2390          0  tcp
07/29 12:54:54       .5      37 b  4704    >   2391          0  tcp
07/29 12:54:56      3.4      23 b  4705  *  >  2392          0  tcp
07/29 12:55:04     21.4  308.7 kb  4706    >   2393          0  tcp
07/29 12:55:27     50.7         ?  4707    >   2394          ?  tcp
07/29 12:59:23        ?         ?  4775    >   1433          ?  tcp
07/29 12:59:25        ?         ?  4774  *  >  3333          ?  tcp
```

**Figure 4 - A continuation of a Bro incident report defining specific alarms raised and connections that occurred.**

The connection history displays the date, time and duration of all the remote and local host interactions. With a bit of footwork, and an archive of past logs from Bro, it is conceivable that we would be able to detect if the particular remote host (attacker) was in our system prior to the attack, assuming the same originating IP address was used. It would be extremely difficult to mine past logs and determine if the same attacker was connecting to the system via different IP addresses (locations) unless we had a specific means of identifying a given attacker (such as a particular signature or rule the attacker regularly violated, however unlikely this event may be).

Figure 3 elaborates on the connection history, telling us that in the past day, 14 successful connections occurred by the remote host, while 10 were unsuccessful. The listing of port scans is also displayed indicating the date, host, and actual port scanned. It is notable that the incident reports do not depict any amount of timeframe that it took for the rule to be violated, but Bro does offer us some form of a confidence interval in it's incident summary by telling us the likelihood of successful attacks in the timeframe.

## 2.2 Sourcefire Intrusion Prevention System

Sourcefire is an IPS/IDS vendor, founded by the creators of Snort, which is an open-source IDS platform. Sourcefire utilizes many of Snort's features in the backend of its security appliances. Snort's popularity is widespread and is considered the "most widely deployed IDS/IPS worldwide," according to its website *(http://www.snort.org)*. Snort has many capabilities that make it

**Figure 5 - Sourcefire's Network Behavior Analysis tool displays charts regarding current network usage.**

effective—logging to databases, free updates to the rule sets (in addition to zero-day updates on a subscription basis), a discussion forum for its large user base, and great documentation provided by both Sourcefire technicians and community members.

Sourcefire successfully built upon the Snort engine by focusing on enterprise network management—a task that requires Sourcefire to become self-aware of a network and identify persons, not just IP addresses. Sourcefire's "Real-Time Network Awareness," or RNA, tries to identify machines, printers and other devices, alerting the administrator to any missing patches and updates that need to be added to networked devices for maximum protection. Further, Sourcefire's "Real-Time User Awareness" allows for Active Directory and LDAP usernames to be associated with one another. If a worm has entered the network, it would be possible to track every user/machine it has come into contact with, via RUA, which would reduce the time to detect and eliminate the treat.

Figure 5 shows Sourcefire's Network Behavior Analysis tool (NBA) which baselines normal traffic so that anomalies, outages and bandwidth consumption across the network can be visualized. For instance, in the event a zero-day exploit has compromised multiple systems, the NBA tool can display a confidence interval of n-standard deviations, shown in Figure 6, that the traffic is abnormal and potentially malicious.

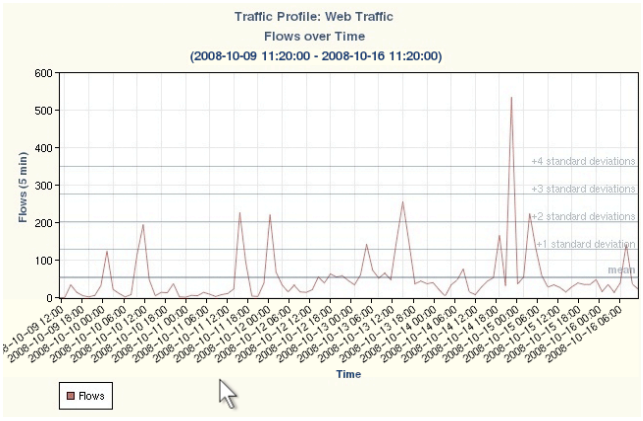


**Figure 6 - The Network Behavior Analysis tool can display how abnormal traffic is as compared to the baseline, normal traffic.**

## 2.3 Radware Intrusion Prevention System

Radware implements it's IPS systems in a unique way by analyzing multiple types of traffic in order to diagnose and defend against attacks before a formal attack signature exists. The IPS does this by analyzing network, server and client traffic patterns. These patterns consist of "rate-based anomalies" (large amounts of traffic), "rate invariant anomalies" (abnormal traffic) and the "attack degree" (actual harmful attack traffic detected). All of these patterns are ranked between one and ten to determine the attack area, which defines the severity of a probable attack as shown in Figure 7. Once a harmful attack pattern is detected, Radware's IPS appliances attempt to generate a blocking rule for that specific attack. If the attack later mutates, the IPS is able to "dynamical modify the signatures characteristics as the attack unfolds." The capability to see an attack unfold and understand that an attack is occurring without a formal signature or definition of the attack is unique to Radware's IPS.

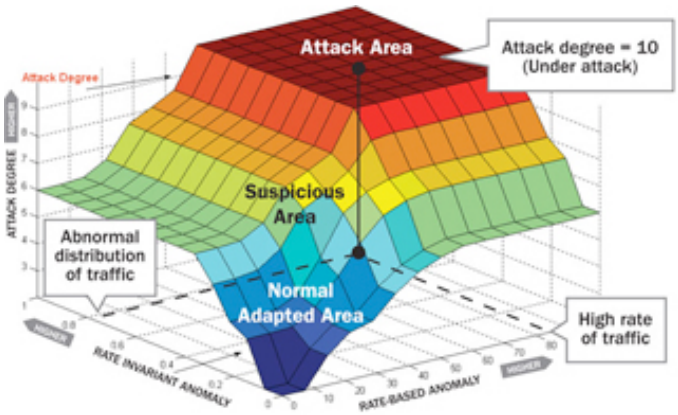


**Figure 7 - The Attack area of Radware's IPS is defined by the rate and abnormal distribution of traffic, as well as the attack degree in order to determine the attack area or severity of a probable attack.**

| Report | Description |
|---|---|
| Top 100 Attacks (last 24 hours) | Those attacks that are detected m |
| Top 100 Attacks Prevented (last 24 hours) | Those attacks that are prevented r |
| Top 20 Attackers (All Attacks - last 24 hours) | IP addresses that have most frequ last 24 hours. |
| Top 20 Attackers Prevented (All Attacks - last 24 hours) | IP addresses that have most frequer during the last 24 hours. |
| Top 20 Targets (last 24 hours) | IP addresses that have most frequ last 24 hours. |
| Top 20 Targets Prevented (last 24 hours) | IP addresses that have most freque |
| All Attacks by Severity (last 24 hours) | Number of attacks by severity lev |
| All Attacks Prevented by Severity (last 24 hours) | Number of attacks prevented by s |
| All Attacks Over Time (last 7 days) | All attacks detected during the las |
| All Attacks Prevented Over Time (last 7 days) | All attacks prevented during the la |

**Figure 8 - A small list of available reports able to be generated by the Juniper IPS solutions.**

## 2.4 Juniper Intrusion Prevention System

Juniper's IPS solutions handle reporting tasks quite well, offering many standardized reports and the ability to visually create custom reports in their security manager appliances. Figure 8 shows a list of the standard reports available to be generated. Additionally, the IPS' offer great functionality in handling user access control. The IPS is able to understand application layer traffic—a key selling point as described in promotional videos—which it also integrates into its access control policies. For instance, if an administrator has to block instant messaging (IM) traffic for a certain set of employees, but allow particular IM clients for select employees, rules can be created accept application from Google Chat for users X, Y, Z and no other users. The IPS being able to understand application-layer traffic and apply network and application rules to its access control system makes it unique as compared to some of the competition's access control systems. Figure 9 displays the interactions between Juniper's IPS nodes and how security policy information can be aggregated between multiple sources in an enterprise setting.

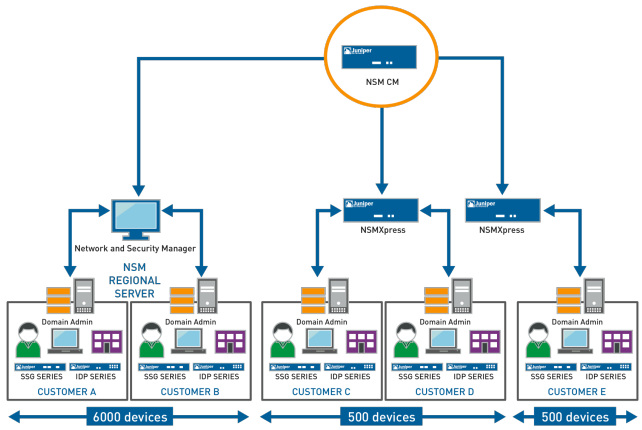


**Figure 9 - Juniper's Access control system is able to aggregate security policy information from multiple centralized sources.**

## 3. Conclusion

Current generation IDS and IPS systems have a vast amount of capability in terms of analyzing, detecting and preventing attacks. In trying to understand specific metrics we found exactly how IDS and IPS systems analyze, report and visualize information to security and network administrators. For brevity, we did not detail each metrics implementation on every IDS/IPS, but opted to include the following charts in Figures 10, 11, and 12. These charts depicts what features and metrics the different IDS and IPS solutions are capable of producing.

It is worthwhile to note that most IPS and IDS solutions can report almost any metric we are interested in finding, but the sheer amount of data could be overwhelming to analyze. Sourcefire engineer Todd Whiting said, in regards to finding an attackers time in the system, that "there would not be packet data provided by Real-Time Network Awareness… this step would require running an RNA report to pull it out and would most likely be a manual process through the GUI in the Defense Center." This indicates that it is possible to store and display such information, but might require some unusual tactics in order to do so. Overall, this paper should shed light on some of the lesser-

known traits and capabilities of both industry-leading and open-source IDS/IPS solutions.

| IDS Comparison Chart | Open-Source | Hardware-Based | IDS/IPS-Type | Attacker Time-in-System |
|---|---|---|---|---|
| Bro | ✓ | ✗ | Signature & Expression-based | ✓ |
| Snort | ✓ | ✗ | Signature-based | ✓ |
| SourceFire | ✗ | ✓ | Signature-based | ✓ |
| Radware | ✗ | ✓ | Signature and Expression-based | ✓ |
| McAfee | ✗ | ✓ | Signature & Expression-based | ✓ |
| Juniper | ✗ | ✓ | Expression-based | ✓ |

**Figure 10 - A feature analysis of IDS/IPS systems. The checkmark indicates the existence of a feature whereas an 'X' denotes the lack of a capability.**

| IDS Comparison Chart | Confidence Interval of Attacker Activity | Return Rate of Attackers | Average Time to Detect an Attack |
|---|---|---|---|
| Bro | ✓ ✗ | ✓ ✗ | ✓ ✗ |
| Snort | ✓ ✗ | ✓ ✗ | ✓ ✗ |
| SourceFire | ✓ | ✓ ✗ | ✓ ✗ |
| Radware | ✓ | ✓ ✗ | ✓ ✗ |
| McAfee | ✓ | ✓ ✗ | ✓ ✗ |
| Juniper | ✓ | ✓ ✗ | ✓ ✗ |

**Figure 11 - The return rate of attackers is dependent on the IP address, whereas the average time to detect an attack relies on a logging packet data for a significant period of time prior to the attack. Therefore, these metrics are considered possible to find, yet unlikely to be done in practice, which is denoted by the checkmark/'X' combination.**

| IDS Comparison Chart | Estimated time to "tripped" alarm | Strength of Access Control System | Overall IDS Strength |
|---|---|---|---|
| Bro | ✘ | N/A | ★★ |
| Snort | ✘ | N/A | ★★ |
| SourceFire | ✘ | ★★★★★ | ★★★★★ |
| Radware | ✓ | ★★★★ | ★★★★ |
| McAfee | ✘ | ★★★ | ★★★ |
| Juniper | ✘ | ★★★★ | ★★★ |

**Figure 12 - SourceFire's 3D IPS was top-rated by SC magazine since 2006, a leading magazine publisher of IT security content. We gave a lower overall rating to the open-source systems because of their lack of reporting tools, GUI's and access control systems. The rest of the rankings are subjective after reading evaluations from SC magazine.**

## 4. REFERENCES


[1] V. Paxon. Bro: A system for detecting network intruders in real-time. In *Proceedings of the 7th USENIX Security Symposium*, San Antonio, TX, 1998.

[2] Bro Intrusion Detection System. http://www.bro-ids.org

[3] Sourcefire Intrusion Prevention System. http://www.sourcefire.com

[4] Snort Intrusion Detection System. http://www.snort.org

[5] Radware Intrusion Detection System. http://www.radware.com/Solutions/Enterprise/Security/IDSTrafficManagement.aspx

[6] McAfee Network Threat Behavioral Analysis. http://www.mcafee.com/us/enterprise/products/network_security/network_threat_behavior_analysis.html

[7] Juniper Network Security Products. http://www.juniper.net/us/en/products-services/security/