


SANS ANALYST PROGRAM

By Jerry Shenk & Dave Shackleford



**Sourcefire
Real-time Network
Awareness**



Sourcefire Real-time Network Awareness

CONTENTS

Executive Summary	1
Introduction - The Need for Endpoint Intelligence	1
Sourcefire 3D System - Discover, Determine, Defend	2
Sourcefire Real-time Network Awareness	4
Passive Discovery	4
Active Discovery	5
Targeted Scanning	5
Business Context	5
Anomaly Detection	6
Easily Deployable	6
Value	7
Benefits of Network Discovery	7
How Passive Discovery Works	7
<i>Frame Header</i>	8
<i>IP Header</i>	8
<i>TCP Header</i>	8
<i>Application Data</i>	8
<i>Passive Discovery Summary</i>	8
How Active Discovery Works	9
<i>Port Scanning</i>	9
<i>Service Banners</i>	9
<i>Malformed Packets</i>	10
True Integration	10
Conclusion	11
Author Bios	11

Executive Summary

In 'The Hound of the Baskervilles', Sherlock Holmes said, "The world is full of obvious things which nobody by any chance ever observes." The detective was famous for seeing obvious information in a way that most people did not notice. Network traffic in our offices and on the Internet has many peculiarities that are not readily apparent. Sourcefire's RNA (Real-time Network Awareness) product uses the minute details in network traffic to passively identify computers on the network and what they are doing. RNA's most recent version combines this passive discovery with active discovery and detection in a way much like Mr. Holmes used his powers of observation and coupled them with probing questions to get the full story.

Intrusion Detection Systems and Intrusion Prevention Systems (IDS/IPS) are prone to false alarms. They notice lots of questionable traffic on the network but don't have a way to know which of that traffic is important. Sourcefire RNA helps solve the problem of false alarms by giving the IDS/IPS situational awareness. RNA can passively monitor the network traffic to get a better understanding of what computers and services are running on it. In the authors' experience, traditional IDS systems are prone to copious alerting; in some real-life deployments of IDS, at least 90% of the alerts that were produced by IDS sensors turn out to be non-events - that is, the alarm tripped but after investigation, the event was determined to be irrelevant. The purpose of this paper is to examine the Sourcefire 3D product claims in an understandable manner while including a few details on how the magic of separating valid alarms from the noise is done.

Introduction - The Need for Endpoint Intelligence

IDS/IPS have typically been lacking in endpoint awareness. They do not know what types of endpoint systems they are protecting. If the IDS/IPS sees a web server attack for example, it does not know or care if it is protecting a web server that is even vulnerable to that attack. Many administrators of Apache web servers will see evidence of attacks that are specific to vulnerabilities in an IIS web server when they review their logs. If there is no IIS web server, then those attacks are not a valid threat against the hosts. This serves to underscore the general concept of "knowing your network" - often, as organizations grow and networks become more distributed, this gets more and more difficult. As a consequence, network and security engineers have a much more difficult job in ascertaining the validity of threats to electronic assets.

Valuable resources are wasted by alerting people to attacks that are irrelevant. An even worse byproduct of these false alarms is that the administrators begin to ignore the alerts. Once people get in the habit of ignoring alerts, the alerts are worthless. Think about the last time you heard a car alarm go off in a parking lot. Did you do anything? Did you even bother turning toward the alarm? Probably not...that alarm is worthless and so are many of the alerts from our IDS/IPS because they are ignored.

Sourcefire 3D System - Discover, Determine, Defend

Sourcefire has developed what they call their 3D System to provide a unified network defense system. The 3D's are Discover, Determine and Defend. The key components of the Discover phase are Threat Intelligence, Endpoint Intelligence and Network Intelligence. The 3D System automatically correlates these three types of intelligence, prioritizes the results and takes automated action to protect against relevant threats.

Martin Roesch, CTO and founder of Sourcefire said, "The Sourcefire 3D System provides users with the first system to integrate both active scanning and passive discovery capabilities in order to have accurate real-time endpoint intelligence. No other system provides this integrated combination to allow users to address the full threat spectrum."

Determination of the validity of a threat is done by the Sourcefire Defense Center. The Defense Center maintains a database containing information about the hosts on the network and builds profiles of traffic that is normal for them. This information comes from RNA sensors and Sourcefire Intrusion Sensors and Agents on the network. The Defense Center can then correlate network changes, anomalous traffic and suspicious activity with its knowledge of the network to provide real-time, prioritized notification and threat mitigation.

The 3D System defends the network according to the ABCs of Defense - Alert, Block, Correct. This includes sending alerts via SNMP traps, syslog alerts and e-mail notification; blocking malicious traffic or replacing the malicious content with benign; firewall responses to block malicious traffic; as well as integration with patch or configuration management systems to apply configuration or code changes to eliminate possible exploitation. The current version of Defense Center includes modules to actively block hostile communication in supported firewalls. Cisco PIX firewalls and any firewall that is compatible with Check Point's OPSEC framework will support active response.

Sourcefire's 3D System is a fully integrated solution allowing management of intrusion detection and prevention, network discovery, vulnerability scanning and event correlation from a single console. This integration allows high-value alerts to be sent to the incident response team and even allows for automated threat mitigation. Incident handlers can look at current data to determine what communication is going on at a given point in time including historical analysis. The data mining and event drilldown features are useful in forensic investigations.

The key differentiator for Sourcefire's system is the capability to perform continuous passive discovery. This is the system's primary method for gathering endpoint and network intelligence, versus more traditional methods of active scanning or distributed agent-based technology. Sourcefire's ability to gather information about hosts and network traffic passively, and build additional intelligence over time to use in correlating event data, makes the solution much more effective than a traditional IDS, or even a standalone IPS.

This passive discovery enables another of the 3D System's key differentiators, the notion of "Impact". When combined with passive system fingerprinting, active targeted scanning, and robust vulnerability signatures, the Sourcefire 3D System is able to label each alert with an Impact rating. The Impact calculation is a radical departure from the traditional arbitrary approach of allowing researchers to rank the relevance of an attack without context. The Defense Center calculates the relevance or 'impact' of an event based on the potential for successful exploitation of the intended target. For example, an alert on an Apache exploit would have a much lower Impact rating if attempted against servers known (via RNA data) to be running Microsoft IIS versus servers running Apache.

These Impact ratings are one of the key inputs that enable the Defense Center's powerful Policy and Response engine to automatically tailor the appropriate responses in real time. For example, an Apache exploit targeted at an IIS server could simply be logged for future analysis while the same exploit against an Apache server could trigger a remediation action to the firewall. The Policy & Response engine also enables enforcement of complex security policies, including compliance with regulations such as HIPAA, FISMA, PCI, etc. An example of this could be a policy that dictates peer-to-peer transfers are not allowed. In this case, a Policy and Response rule could be written to deny any P2P connections before any potentially damaging violations such as data leakage can occur.

Endpoint Intelligence is accomplished by the Real-time Network Awareness (RNA) sensor using passive discovery methods to define the entire network. Sourcefire has determined that the RNA sensor is 80-90% accurate in determining what devices are on the network simply by monitoring traffic as it passed through the network. In addition to passively monitoring network traffic, RNA has an active component that can make targeted scans of a device or specific service when new traffic patterns are detected on the network. This targeted scanning will substantially increase the accuracy of data in the RNA database. To test these claims, an RNA sensor was configured as a component of the overall 3D System, and placed into a test network with a variety of operating systems running. The following systems were detected and accurately identified:

- Windows 98, 2000, and XP
- Solaris 8 and 9
- FreeBSD 5.4
- Red Hat Enterprise Linux 3
- Fedora Core 3 (Linux)

In addition, RNA was able to accurately identify a variety of Apache and IIS Web servers with version information, as well as FTP and other network servers.

Sourcefire Real-time Network Awareness

Sourcefire uses four techniques to maintain its view of the network - passive discovery, active discovery, targeted scanning and business context information. This information is correlated to determine the criticality of any questionable network traffic.

Passive Discovery

Passive discovery is the primary function of the RNA sensor. The sensor must be placed in an area on the network where it can observe traffic on the wire. This placement would normally be near an Internet gateway or routers that connect sites. Another option is to set up a monitor port in a core switch.

Different operating systems and hardware devices have peculiarities in the way they transmit data. One simple example is the Time To Live (TTL) field in an IP packet. This field is decremented every time a packet crosses a router and is used to prevent routing loops. Most modern Windows computers have a starting TTL of 128 while Windows 3.11, some Windows 98 computers and some Windows Advanced Server 2003 have a starting TTL of 32¹. If a packet were picked up on a network from another computer on the local network and the TTL was 31, it would probably be from a Windows 2003 Advanced server because the other named operating systems are pretty old. There are also some other hardware devices that have a starting TTL of 32 so the TTL setting alone isn't foolproof, but there are also a lot more settings in the packet headers that can help us determine the operating system on the computer. In Michael Zalewski's book, *Silence on the Wire*, he dedicates chapter 9 to the many parameters in IP, TCP and UDP headers that can be used for passive fingerprinting. Some examples include Window size, the Don't Fragment (DF) bit being turned on, and the Type of Service (TOS) field.

Network administrators gain real-time knowledge about their network from RNA's passive discovery. The Defense Center can be configured to alert any time a new host shows up on the network. In a large network, it's too easy for a well-meaning end-user to install a device that compromises the security of the entire network. One example of this would be a wireless access point that enables a technically savvy user to roam about his office. As soon as a single packet from this rogue access point is seen by a sensor, the Defense Center can alert an administrator that a new system has been seen on the network. Most of the time, that first packet will be a DHCP request that happens as soon as the device is connected to the network. The network SWAT Team can then be deployed to contain and eradicate the situation.

Placement of the RNA sensors is critical to the accuracy of the endpoint identification. In the previous example, it is easy to see how a sensor that is 10 network hops away from a given endpoint would not be as accurate. This holds true for other identification vectors as well. RNA sensor placement is a balancing act. The sensors should be as close as possible to the endpoints while avoiding the cost of deploying too many sensors.

In smaller networks, locating sensors at network chokepoints could be as simple as a hub at a natural chokepoint such as a router. In most large networks, a switch will be configured with a mirrored port or SPAN port. The switch will be configured so that all the traffic that should be monitored gets forwarded to the port that RNA is monitoring. This could be monitoring all traffic on the network or traffic going into and out of the local network. It should be noted that this technique may not scale well in high-speed networks, as some packets may be lost when multiple ports are spanned. Initial deployments of RNA would often start by monitoring Internet traffic, traffic to directly-connected vendors or suppliers and critical servers.

¹ Pof documentation

Active Discovery

RNA couples the advantages of passive discovery with the advantages of active discovery. Active discovery has the advantage of soliciting responses from computers that may not normally send data across the network. A computer not sending data would be quite rare. Active discovery can also look for open ports on a computer that normally wouldn't be receiving data but could still be vulnerable. A recent example of this would be the vulnerability that was discovered on port 445 of many windows computers in mid-June 2005. An RNA appliance could be configured to scan the network looking for vulnerable computers so that the IT staff could be alerted. Initially, there was no worm to take advantage of this vulnerability but analysts expected that one would soon be coming. In a case like this, waiting for the bad traffic would be too late. Detecting a security problem after a system compromise is better than nothing but it is far better to know ahead of time and patch vulnerable systems to eliminate a widespread network attack. In some cases, RNA can even force patch management on systems that it finds that are not compliant with the current patch policy.

Targeted Scanning

In order to increase the accuracy of the network discovery, RNA can also leverage open source Nessus to actively scan devices to see how they respond to certain network stimuli. Douglas Hurd, director of product management at Sourcefire says, "If you listen passively, you can learn 90 percent of what you need to know, then fill in the gaps with surgical scans."

Nessus is rated by many as the best network vulnerability scanner and Sourcefire has chosen to use it as the vulnerability scanning part of the Sourcefire 3D System. These scans can run full system scans on a periodic basis or targeted scans based on pre-defined network activity criteria. For example, targeted scanning might be initiated whenever RNA detects a new server process or endpoint on the network.

RNA can also do full scans of a host or network, as a response or on a scheduled basis. The primary difference between having RNA run a scan and doing a stand-alone vulnerability test is that the RNA scan puts the information into Defense Center so that it can be used to target alert notification. As RNA by itself will not detect patch levels or specific vulnerabilities in running network-based applications (such as Apache or IIS), the ability of the system to perform targeted Nessus scans based on RNA passive detection is a very powerful combination of tools. This can be used to great effect, significantly reducing the rate of false positives seen with many traditional IDS and IPS systems.

Business Context

Hosts in the Sourcefire 3D database can be prioritized based on business criticality. This is done using the "host criticality" setting. This is a manually maintained setting that is included in Defense Center criticality calculations. By increasing or decreasing this value, it is possible to differentiate between an HR server with highly confidential data and a test server. This information can also be fed into Sourcefire's Policy & Response engine to build complex compliance rules to trigger a different response or remediation for hosts with different criticalities. This is best illustrated in terms of the threats against these systems - the exact same threat targeting a critical payroll or HR server can be treated with a higher priority than if it were targeting a test server. This can easily help an analyst to quickly prioritize tasks when responding to a potential incident.

This is where Sourcefire's solution really shines, in the authors' opinion. The ability to assign host importance ratings, in conjunction with the definition of Impact values, allows analysts to get the full value of Intrusion Sensor data and RNA data combined. This puts all of your network's detected assets into a "real-world" context that truly changes how Intrusion Detection and Prevention is done.

Anomaly Detection

With RNA, the Sourcefire 3D System integrates and correlates Network Based Anomaly Detection (NBAD) functionality. Traffic profiles can be created for any endpoint, group, subnet or network. These profiles can be monitored using absolute values or standard deviations. With this type of monitoring, it would be possible to raise an alarm if a user connects to an external server and starts transmitting a large amount of data. This could be perfectly legitimate activity but it could also be related to transmitting confidential data to an external host. In mid June 2005, a credit card processing company was found to have been transmitting a large amount of credit card information to hostile third parties. This was eventually discovered but it would have been better for the information leakage to have triggered an alert. Prevention would have been the best but if a break-in occurs, quick detection is necessary. The 3D System could have handled this in a number of ways. First, outbound rules could easily be triggered when patterns matching credit card data started traversing the network. Second, alerts could be triggered when large amounts of data began flowing from certain hosts to external systems. Finally, alerts could be triggered for any other detectable type of traffic pattern that was identified; for example, a Windows system mysteriously sending Syslog traffic outbound.

Sourcefire's NBAD system integrates the concept of traffic baselines into network monitoring. Based on defined policies and rules, NBAD can track specific network events or traffic flows that exceed thresholds. This can help to identify and follow Distributed Denial of Service (DDOS) attacks as they are happening, as well as malicious code events such as worm outbreaks. This is also a simple method for detecting zero-day exploits, as signatures will not detect these in most cases.

Easily Deployable

RNA sensors and Intrusion Sensors can be deployed throughout a network infrastructure easily. RNA is available as an appliance or as a software only product. The appliances can be pre-configured with specific configuration information, shipped to their eventual destination, and then simply connected when they arrive. There is no need for on-site configuration. Final modifications can be done remotely over an encrypted SSH session if needed. Most modifications can be done from the GUI of the Defense Center console.

There are no agents that need to be installed on workstations for either RNA sensors or Intrusion Sensors. Not having to install client software alone can save an IT department hundreds of man-hours in deployment time across even a moderate sized network. In addition to the costs associated with installing agents are the hidden costs related to agent maintenance and performance and compatibility problems that often are accompanied by installing any type of agent on a large number of workstations.

In large and distributed environments, the Sourcefire 3D System can also easily incorporate dynamic load balancing, with multiple Defense Centers managing sensors. In addition, existing Snort sensor deployments can be immediately leveraged by installing the Sourcefire Intrusion Agent software, allowing these sensors to report to a Defense Center along with other RNA devices and Sourcefire Intrusion Sensors. This allows many organizations to quickly gain the additional features and functionality of Sourcefire's commercial tools while having a minimal effect on current infrastructure.

Value

Because of the insight into the network that RNA provides, it adds value to any security technologies that may currently be deployed on a network. RNA can maintain a database of endpoints on the network, what services they are running and the layout of the network. This gives network operations staff the ability to see what's happening on the network and either update their policy or take corrective action to return the network to policy compliance.

The biggest value that RNA provides is an increase in efficiency of incident handler's time. RNA increases efficiency because it reduces the number of alerts that need to be analyzed by eliminating alerts that are not relevant.

Benefits of Network Discovery

Passive network discovery and integration with an intrusion detection and prevention system is a revolution in the computer security space. The fact that passive discovery can be running 24 hours a day, 7 days a week means that anything new that happens on the network will get picked up. Augmenting that passive discovery with the option of periodic and/or targeted active discovery gives RNA a view into the network that is complete. This enables network administrators and security personnel to be more fully aware of what's happening on the network.

Passive network discovery avoids the dangers of active discovery because RNA never sends any traffic to the endpoints while in passive mode. Passive network discovery avoids one of the risks of active network discovery; that of accidentally sending traffic to a host that causes an unexpected result such as a system slowdown or a crash. Another advantage of full-time passive network discovery is that people can't just turn a machine off during the vulnerability scan. At times, computers can be deliberately hidden from the scan because they know they aren't up to date.

How Passive Discovery Works

Passive network discovery is simply listening...listening very carefully. In Zalewski's book, chapter nine is subtitled, "Passive fingerprinting: subtle differences in how we behave can help others tell who we are."

Data packets are part of our everyday lives. We look up the weather, read news, communicate using e-mail, coordinate schedules, use VOIP to carry phone conversations...the list seems endless. All the data packets that carry the data we want to use (pictures, words, sounds, etc.) also have headers that the computer and routers on the Internet use. As a packet leaves a computer, it gets different headers added to it. It may have an application header and then a protocol header and a frame header. On the receiving end, the computer strips those various headers off as it processes the data packet. Those headers contain the fingerprints of the computer that put them together, the routers that last touched them and at times other fragments of fascinating information.

The reason that there are small differences between the ways different computers build their data packets is because there is some room for interpretation in the rules that govern Internet communication. These rules; perhaps guidelines is a better term, are defined in papers called RFCs or Request for Comment papers.

Frame Header

The first header contains the local source and local destination. This isn't the ultimate source or destination, just who touched it last and who will touch it next. In the case of TCP/IP, which is what we use the most, it also contains information about the protocol header. However, not all traffic on the wire is TCP/IP. Some packets will include other information. The most interesting thing we can get from a packet at this level is the hardware address of the sender and receiver. That hardware address is assigned to a specific manufacturer. If one of the addresses is assigned to Cisco, and the other is 3Com, we can start our information pool with that.

IP Header

If this is a packet that is part of an e-mail session across the Internet, the next header would be the IP header. Here we start to see quite a bit of information. The TTL field that was discussed earlier is in this header. The final source and destination IP addresses are in this header. This header also contains options, IP identification numbers and fragmentation information. Different endpoints will handle this information slightly differently. There is quite a bit of information here to help identify the originating host.

TCP Header

Inside the IP header of our mail session is a TCP header. This header contains source and destination ports, and a number of other pieces of data for transporting data between computers. One interesting field is the window size. This field controls how much data can be sent before an acknowledgement is expected. This setting has an impact on performance and responsiveness. For use in fingerprinting, different operating systems have different default window sizes. This is one of the fields that can help differentiate operating systems and versions. This can even identify different service pack levels on various Windows operating systems.

Application Data

The last grouping of data is the application data. In our e-mail session, a user is picking up mail from a mail server. When the user connects to the POP3 server, the server normally announces who they are and what software they are running...that is helpful for OS determination. Then the client announces who they are and provides proof of that in the form of a password. The next step is to transmit messages to the client, do some housekeeping and disconnect.

Passive Discovery Summary

In a very few seconds, we have gained a very good idea what the client OS is running. We know what the mail server says it is running. This information could have been changed by the administrator but is compared with other information about the server. We've seen a user log in to a mail server so now we know the user's name.

There are many other communication sessions that RNA could see to increase the reliability of its OS classification. NetBIOS information includes the machine name and gets broadcast on a regular basis. Cisco devices send Cisco discovery packets on a periodic basis that announce things like router name, IOS versions and more. Novell and Microsoft servers announce the services that they have available. New workstations send out broadcasts asking for an IP address and routing information. RNA keeps collecting new information that it sees about hosts on the network. As that information supports information that it has previously learned, the database in Defense Center is updated and the confidence factor is increased. If the information contradicts what's previously been learned, the confidence factor is decreased.

How Active Discovery Works

Active discovery is a feature that has been added to the latest version of RNA. This feature supplements the passive discovery that has been the core of RNA since its introduction in 2004. Active discovery has three major components - port scanning, banner grabbing and sending malformed packets. Port scanning is where RNA will send packets to every port on a computer to see if something answers. In the banner grabbing, RNA connects to computers and other endpoints and analyzes the service banners. Service banners are normal responses from services that are offered by a computer. These banners often include information about the manufacturer and version of the software that is listening on that port. In the third test, RNA sends strangely configured packets to an endpoint device and analyzes the response.

Port Scanning

Computers have a total of 65,535 TCP ports and just as many UDP ports. Portscanning is done by attempting to connect to each one of those ports to see if anything responds. In some cases, a sub-set of the most common ports will be tested to minimize traffic and get quicker but less complete results.

Service Banners

Service banners are the “welcome screens” that most computer services use to announce to other computers what they are running. In attempting to discover what a computer is running, Nessus will do banner grabbing. For example, if a person were to connect to port 25 on one of Microsoft’s mail servers, you might get something like this:

```
220 IGR-IMC-01.redmond.corp.microsoft.com <Inbound SMTP Virtual Server>  
Thu, 7 Jul 2005 14:53:57 -0700
```

For the purposes of identifying the operating system, this message does not tell us a whole lot because the most interesting information has been hidden. This suggests that the server is in Redmond, Washington and it tells us that what the time is and what time zone the server is reporting to be in.

If we connect to a web server from this same company and request some server information, we get the following:

```
HTTP/1.1 404 Not Found  
Content-Length: 103  
Content-Type: text/html  
Server: Microsoft-IIS/6.0  
Date: Thu, 07 Jul 2005 22:01:00 GMT  
Connection: close
```

Here, we can clearly see that this server is running Microsoft IIS version 6 and we can also see the current time but not the time zone as we saw on the mail server.

These are two simple pieces of information that Nessus will use in its active network discovery. Nessus can run thousands of tests.

Malformed Packets

Some of the active discovery tests are quite benign like banner grabbing, but other tests send malformed traffic to a host to see how it responds. Nessus uses a program called nmap for some of these “malformed packet” tests.

One example of this is to send a packet that claims to be the start of a communication session and also the end of a communication session. This is called a SYN/FIN packet and it really does not make sense in any TCP/IP stack. There is no legitimate reason a computer would ever send a packet like this, so the RFCs that govern IP communication do not specify how these packets are to be handled. For this reason, response to a packet of this type is left up to the operating system developers. Some operating systems will ignore this packet while others will respond as if it is a perfectly normal start of a connection and there are other possible responses as well. Active discovery can be used to evaluate responses to these types of packets and come up with a very good guess of what the operating system and version is. This information can be combined and correlated with what RNA has already discovered passively to increase the accuracy of the network discovery.

True Integration

Sourcefire has created a unique, truly integrated solution with its 3D System, which includes the Defense Center, RNA and Intrusion Sensors and Agents. The integration of threat intelligence (Intrusion Sensors and Snort sensors), endpoint intelligence (RNA) and network intelligence (RNA’s NBAD capabilities) provides a comprehensive approach to network security.

This “true integration” is achieved through:

- Natively and simply integrating with Cisco PIX and Checkpoint and other OPSEC-compliant firewall devices. Sourcefire has made active response a simple matter by including simple GUI-based interfaces in Defense Center that allow alerts to trigger firewall rule generation. Sourcefire also provides an open API for creating custom alerting and active response mechanisms within 3D.
- Integrating with Nessus for active scanning for vulnerabilities. The ability for passive discovery and active scanning to work seamlessly together adds an exponentially-powerful layer of security intelligence into this platform.
- By integrating the endpoint intelligence about the network hosts, the 3D System is able to decrease security events by over 90%, greatly increasing the security analysts’ productivity.
- Integrating with Business Context data to allow for security alerts to be tailored to each business’ needs.
- The ability to run Sourcefire Intrusion Sensors or Snort sensors with Intrusion Agents. This is a very attractive option for organizations with an existing Snort deployment, and allows for rapid scalability.
- Ease of integration. Sourcefire offers both passive and inline (IPS) capabilities with its sensors, and has a number of models to choose from depending on network speed, port density, and other aspects.

By integrating passive and active network discovery with intrusion sensors, a database of hosts, services and traffic, Sourcefire has developed a technology to make key security components and technologies interoperate.

Conclusion

With RNA, the Sourcefire Defense Center and 3D System has the potential to be at the center of the network. RNA adds situational awareness to Intrusion Detection and Intrusion Prevention. Using additional tools like RNA Network Visualizer, administrators can see graphical representations of the network, with real-time alerting and behavioral network intelligence. When new systems connect to the network, administrators will know about them quickly. When threats are detected, administrators can assess the actual risk impact to the organization and make informed decisions about incident handling. Sourcefire even includes some basic incident reporting capabilities as a built-in component of their 3D System.

By adding the ability to classify asset criticality and dynamically assess the impact of threats to those assets in short order, Sourcefire's system goes above and beyond the traditional taboo of IDS/IPS: too much data, not enough information. This sets them apart from other network-based detection and protection technologies in a number of ways, and truly represents an evolution in the way that security and network teams can proactively defend their information assets.

Author BIOS

Jerry Shenk: Jerry currently serves as Senior Analyst for the SANS Institute. Since 1984, he has consulted with companies, financial and educational institutions on issues of network design, security, forensic analysis and penetration testing. His experience spans small home-office systems to global networks. Along with some vendor-specific certifications, Jerry holds 5 GIAC GOLD certifications: GCIA, GCIH, GCFW, GSNA and GCFA: all completed with honors.

Dave Shackleford: Dave Shackleford has been involved in Information Technology, particularly the arenas of networking and security, for over nine years. Dave has worked as a security architect and manager for a number of large companies, and has also run his own consulting practice for several years. He currently works as the Solutions Engineering Manager for Vigilar, and his areas of specialty include incident handling and response, intrusion detection and traffic analysis, and penetration testing). Dave holds CISSP, GCIA, GSEC, GCIH, G7799, CCNA, MCSE, and MCIWA certifications, and is working on his MBA. He is also the co-author of *Hands-On Information Security* from Course Technology.