TECHNOLOGY BRIEF

SOURCEFIRE RNA® (REAL-TIME NETWORK AWARENESS)

DEALING WITH DYNAMIC THREATS

KNOW MORE NETWORK RISKS
NO MORE GUESSING

SOURCEfire®

INTRODUCTION

The Maginot Line is considered to be one of the greatest failures of military history. It is a line of fortifications, machine gun posts, and other defenses that France built on its borders with Germany and Italy. The Line was a response to the brutal attrition of trench warfare used in World War I, where each side established fixed fortifications and sought to grind down the other side over many months. The considerable firepower of the Maginot Line was intended to prevent a surprise attack and hold the enemy until reinforcements could be called in.

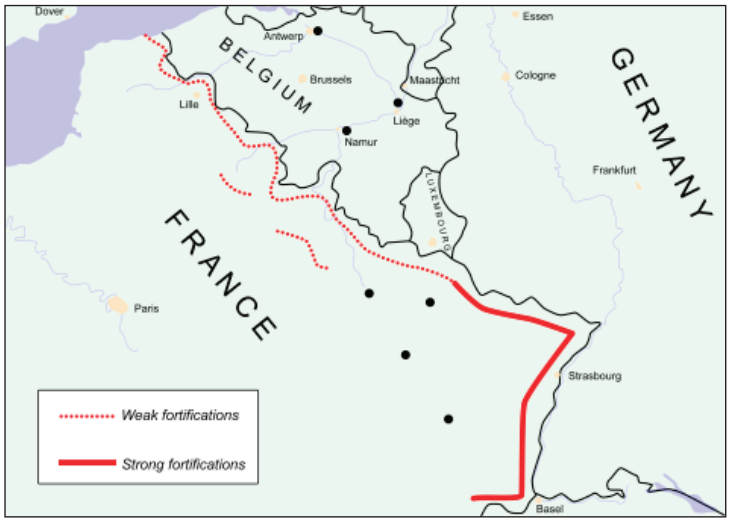

The French had adopted a World War I mentality when building the Maginot Line. They assumed that warfare was relatively static because munitions would be relatively immobile, and that the enemy would approach from known vectors. Both assumptions were proven wrong by the German Blitzkrieg attacks in 1940. Blitzkrieg literally means "lightning war," where Germany used mobile weapons such as tanks and planes in a coordinated fashion to bypass the Maginot Line entirely and overwhelm France's defenses. Ironically, the Maginot Line worked exactly as planned. It prevented any significant attack where the line was strongest. The Line, however, was made irrelevant with the advent of modern warfare, which prioritizes mobility and coordination.

Although the German Blitzkrieg took place nearly 70 years ago, it can still teach today's security professionals valuable lessons. Many enterprises rely on static security defenses, such as a firewall and an intrusion prevention system (IPS), as their sole form of security. These defense mechanisms are similar to the Maginot Line—they are required to prevent attacks from known vectors, but they are the product of last decade's security war. Today's network security requires technology that goes beyond static port blocking and signature detection.

THE ADVENT OF DYNAMIC THREATS AND IMPORTANCE OF CONTEXT

This section provides some background about modern security tools. It describes how many tools are static in nature, but the emergence of dynamic threats makes static tools increasingly irrelevant. Modern tools must respond to dynamic threats by automatically providing context that can be used to effectively respond to the threats.

The Prevalence of Static Security Tools

Security tools introduced in the last 10 to 15 years are mainly based on static mechanisms. For example, firewalls block traffic on specific TCP or UDP ports. Intrusion detection systems (IDSes) look for fixed signatures to identify malware. Static tools were originally designed with the assumption that human operators were knowledgeable about interpreting the output and would continue to modify the tools' configuration accordingly. Another critical assumption was that the operators understood the security threats they faced and could configure the security infrastructure correctly.

Both assumptions have not proven to be correct over time, however. Modern security analysts are usually pulled in many different directions, needing to maintain multiple network- and host-based technologies and monitor their output. Analysts must also respond to crises at any time, such as a virus outbreak on an executive's laptop or forensic research for a human resources investigation. The result is that security analysts often lack the time to monitor any one tool closely, let alone update the tool's configuration regularly. Analysts also are unable to become experts with any one tool in deployment or optimization.

What further adds to the static nature of security tools is that many organizations reinforce this nature in their processes. For example, blocking a new port or unblocking an existing port on a firewall usually involves a change request process and can often take days. These processes exist to avoid destabilization of the security infrastructure, but they also make it more difficult to quickly respond to emerging threats.

The Emergence of Dynamic Threats

Even as security tools have continued to be static, malware has evolved considerably in the last 10 to 15 years. Worms and viruses used to be written by curious hobbyists, but are now developed by dedicated research labs and funded by organized crime organizations that are financially motivated to steal personal data and sell it to the highest bidder.

Modern malware can easily circumvent the static defenses of a basic firewall or IPS. The malware can dynamically use any number of ports to connect with hosts, so the blocking of specific ports is

often ineffective. And malware can use various self-modifying techniques to hide its shellcode, so searching for static attack signatures may not yield useful results.

Fortunately, some security tools have evolved over time to deal with modern threats. For example, Snort® is an IPS that can detect attacks made by modern malware. Snort models specific network protocols and understands how they work so it can detect whether the protocols are being exploited. Snort does not solely rely on signature detection and can identify protocol exploits even if malware tries to obfuscate itself. Protocol modeling enables Snort to identify any potential exploit, regardless of what the malware code actually looks like.

Even with this type of intelligence, however, any IPS may generate so much noise that a security analyst may have difficulty properly assessing the situation. Even if an IPS sees a packet that could possibly be identified as containing malicious code, it may not have sufficient context to make the correct assessment.

The Importance of Context

Context refers to the ability to understand and analyze important environmental factors. For example, here is a picture of a man walking into a building. Is this man a threat? It is impossible to know for sure without considering the proper context. He could be carrying a gun into a bank, or he could be carrying an umbrella into an apartment building. One must consider contextual factors, such as the type of weather outside or exactly what type of building he is walking into.

In the same way, understanding the network environment and having context allow an intrusion analyst to prioritize IPS alerts and to decide whether some alerts warrant further investigation.

For example, if an IPS creates an alert because it detects an exploit packet using a specific vulnerability in Windows XP SP1, the security analyst does not need to take action if the packet is targeting a Linux host because the exploit will not work. There is also no need to respond if the packet is targeting a host running Windows XP SP2, if it is already patched against the vulnerability.

Unfortunately, most enterprises lack this type of context. They often lack sufficiently-trained personnel to gather accurate information about their networks and keep this information updated. Therefore, they are not properly equipped to quickly respond to incidents. Martin Roesch, Founder and CTO of Sourcefire, has mentioned numerous times that organizations suffer from a lack of visibility into their networks. They cannot properly defend and respond to what they do not know about. This problem is echoed by Marcus Ranum, a well-known security expert, when asked to name the weakest links in the network security chain. One of these links is network awareness:

> "All through the 1990s until today, organizations were building massive networks and many of them have no idea whatsoever what's actually out there, which systems are crucial, which systems hold sensitive data, etc. The 1990s were this period of irrational exuberance from a security standpoint - I think we are going to be paying the price for that, for a long time indeed. Not knowing what's on your network is going to continue to be the biggest problem for most security practitioners."[1]

These expert opinions are supported by known data. According to a 2009 Verizon Business study of nearly 600 data breaches taking place over five years, 71% of the breaches were not discovered by the breached organizations, but instead by a third party, the confessions of a perpetrator, or a third-party audit.[2] These breached organizations simply had no idea that their networks had been compromised.

Figure 32. Breach discovery methods by percent of breaches

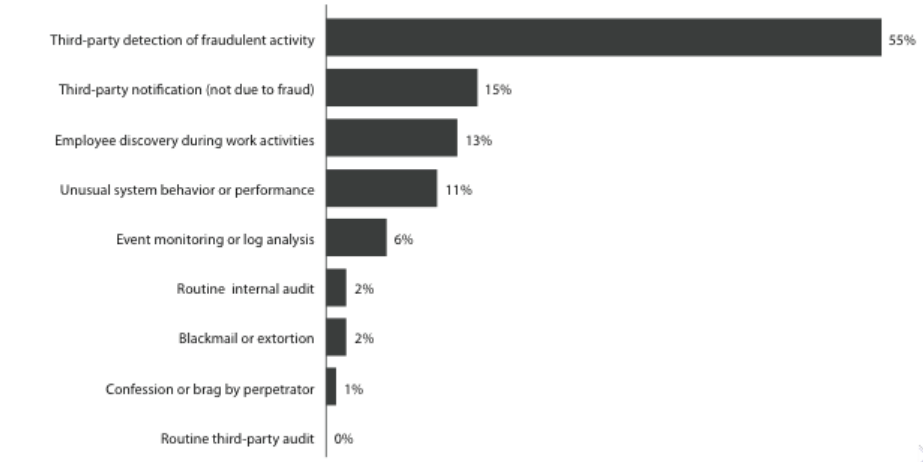

1 http://www.csoonline.com/article/461422/Marcus_Ranum_on_Network_Security
2 http://www.verizonbusiness.com/resources/security/reports/2009_databreach_rp.pdf

Another important finding from the Verizon Business study data is that only 6% of the breaches were discovered internally by event monitoring or log analysis technologies. Organizations may purchase tools and even have staff to monitor them, but unless the staff is given the correct context to interpret the data they are receiving, the tools and data have limited usefulness.

In today's age of dynamic threats, it is essential for organizations to better understand the composition of the hosts on their networks and the applications that these hosts are running. This intelligence provides the context needed to prioritize intrusion alerts and gauge their malicious intent.

REAL-TIME NETWORK AWARENESS

Sourcefire RNA® (Real-time Network Awareness) is a passive sensing technology meant to assist with the daily tasks of intrusion analysis. RNA monitors network traffic in real time and tracks the configuration changes and network behavior of hosts. Because RNA is passive, it provides a number of advantages over traditional network monitoring technologies that rely on active scanning or host-based agents:

- Generates no traffic that can either consume bandwidth or potentially harm network infrastructure
- Provides a real-time network view that is continually updated as devices produce traffic—data never becomes "stale" or dependent on periodic scans
- Does not rely on endpoint agents that require management—unknown or rogue devices are always visible

RNA has a number of features that are highly relevant for intrusion analysis. It helps to provide network context and substantially reduces the number of actionable intrusion events. RNA also leverages its knowledge of the network to provide auto-tuning for Sourcefire IPS™, reducing complexity and increasing performance.

Reducing Intrusion Alerts

RNA uses its passive technology to gather data about assets on the network. This data becomes the basis of a network map of all hosts along with their operating systems, applications, and potential vulnerabilities. The Sourcefire 3D® System leverages this data to prioritize the potential impact of intrusion events.

For example, Sourcefire IPS sees a packet in the network that exploits a vulnerability in Windows XP SP2. Because this packet is being sent to a host that has been classified as a Linux host, the 3D System rates the impact of the event as "Not Vulnerable." At a later time, however, Sourcefire IPS sees this same type of packet being sent to a host that has been classified as running Windows XP SP2. The 3D System rates the impact as "Vulnerable" because the host is potentially vulnerable to the exploit and may have been compromised by the packet. A security analyst can then focus on investigating only the "Vulnerable" events and disregard the "Not Vulnerable" events.

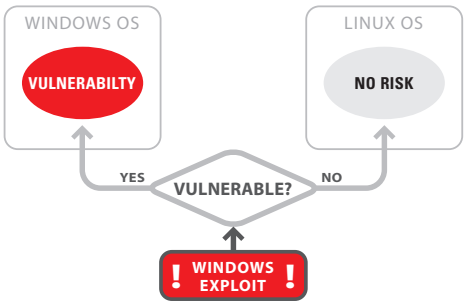

Sourcefire's Impact Analysis can reduce the volume of actionable events by a factor of 10 to 100. RNA gives the security analyst the necessary context to properly understand whether an intrusion event is truly worth investigating, allowing him to focus on other duties.

IPS Auto-tuning – "Adaptive IPS"

RNA uses a feature known as RNA-Recommended Rules (RRR) to help tune Sourcefire IPS. RNA knows what operating systems and services are running on the network, so RRR can recommend that only relevant Snort rules should be enabled. For example, if RNA determines that a protected network segment is only running Linux systems supporting Web services and NIS, then RNA can recommend disabling any rules pertaining to Windows hosts and services, such as IIS. RRR is designed to maximize protection and sensor performance and significantly reduce, or virtually eliminate, the manual and ongoing effort required to tune Sourcefire IPS. Rule recommendations made by RRR can be implemented with or without human intervention.

RNA also supports other auto-tuning technologies. Another Adaptive IPS tuning feature, non-standard port handling, helps to prevent possible IPS evasions by inspecting traffic on non-standard ports. For example, a user may attempt to conceal HTTP traffic by running it on TCP port 8888 instead of the standard TCP port 80. Traditional IPS implementations can detect this traffic, but the result is usually sub-optimal performance or requires manual configuration of rules for non-standard ports. RNA, on the other hand, can identify the ports and services on the hosts it is monitoring and configure the IPS to apply the correct rules for any non-standard ports.

In summary, RNA utilizes a number of auto-tuning technologies that provide multiple benefits:

1. Reduce manual work in tuning Sourcefire IPS
2. Help prevent IPS evasions
3. Maximize sensor resources

HOW DOES SOURCEFIRE RNA HELP ON A DAILY BASIS?

Besides features directly relevant for intrusion analysis, RNA also supports other features that are useful for security analysts. RNA creates a real-time inventory of all operating systems, services, applications, and protocols on the network. The product also tracks network traffic flows between hosts. All of this information is relevant for different functions including change management, policy enforcement, and network behavior analysis (NBA).

This section provides a look into the life of a typical security analyst to show how these features can be used on a daily basis.

Profile of a Security Analyst

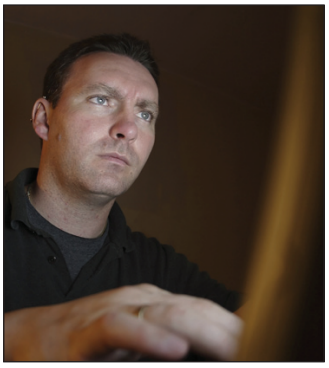

David is a Security Analyst for his company of approximately 1,000 employees. His company's security group is very small, consisting mainly of him and his manager, who functions as the IT Security Director and spends much of his time dealing with management and other departments. David focuses mainly on the technical aspects of security, such as administering the IPS, implementing wireless security, and performing vulnerability assessment scans.

David may be called upon at any time to deal with crises, such as incident handling after a malicious compromise or the forensic investigation of an employee who is suspected to be illegally accessing corporate resources. He is "spread extremely thin" because various groups demand his skills, expertise, and time. Due to financial constraints, their company is unable to hire additional information security personnel.

Deploying Sourcefire Network Security Equipment

David is a fan of Sourcefire because he has some experience using the open source version of Snort. When his company was initially purchasing Sourcefire network security equipment, David also advocated the purchase of RNA because he realized that RNA could save him considerable time by reducing the number of intrusion events he needs to investigate.

The company has installed Sourcefire 3D® Sensors running both IPS and RNA at various monitoring points, such as the DMZ, inside the firewall, and in front of the database servers, corporate finance systems, and the IT management network. There are a mix of Windows and Linux servers in the data center and mostly Windows desktops in the internal network.

Because David is overwhelmed by his many responsibilities, he has tried to automate his IPS infrastructure as much as possible. He would also like to use RNA to help him enforce a security policy he is developing to only permit certain applications and services on the hosts on his network.

Here are some of the steps he has taken:

- He configures the Sourcefire 3D System to send intrusion events rated as "Vulnerable" to his pager. All other events are simply logged to the 3D System. Every morning at 7 am, the system generates a new report for all events over the past 24 hours and sends it to him via e-mail. Since David reviews this report every day, it is easy to scan for anything out of the ordinary.
- He enables RNA-Recommended Rules (RRR) to provide auto-tuning so he does not need to spend much time hand-tuning the IPS intrusion policy. For example, RRR identifies that the internal network consists only of Windows systems, so it recommends that all rules related to non-Windows systems be disabled.
- RNA supports a feature known as host attributes to assign labels to certain hosts. David uses this feature to identify specific servers in the data center as "Critical." He works with the server team to determine what services and applications should run on these critical servers, and he uses RNA's white list feature to automatically notify him if any new services or applications are running on these servers. He coordinates with the server team so that when new applications are installed, he is notified so he can update his white lists.
- For certain critical subnets, such as the DMZ, the database server subnet, the finance server subnet, and the IT management network, David uses RNA's traffic profiles feature to analyze traffic in each subnet to search for anomalies. If network traffic patterns change due to a possible worm or DDoS attack, David can automatically be notified.

Because of his visibility into the network, David is in a position to help other IT groups in his company. He regularly integrates his vulnerability scan data into RNA so RNA can track the specific vulnerabilities of each host. Therefore, he can provide reports of all hosts showing their respective operating system, services with vendor/version, and potential vulnerabilities. The server group needs this data, but they cannot easily obtain it from the network management system because the system relies on agents. Agents cannot be deployed on all systems for performance and compatibility reasons.

Prevention Eventually Fails

David has created an automated system to help him defend his company's network infrastructure. He is a realist, however, and knows that no system is completely impervious to attack. He agrees with Richard Bejtlich, a well-known security blogger and analyst, who says, "Prevention eventually fails. [The enterprise environment] is too complex, staffed by

overworked, under-resourced administrators meeting 'business requirements.' [As a result,] every enterprise will eventually be compromised." David wants to be prepared for the eventuality that a malicious user will gain access to hosts or data within his network.

One day, David learns of a possible intrusion when the 3D System notifies him that some internal hosts are generating an unusually-high amount of network traffic to the Internet. From a quick visual scan, it looks like much of the excessive traffic consists of TCP connections to destination ports 135–139, 445, and 1433. David is also paged by a Desktop Administrator and is told that users have been complaining about very slow network performance. David suspects that at least one or more systems have been compromised by a worm.

David has collected flow data for the past seven days from each of the 3D Sensors. He applies a filter so he is only looking at TCP connections destined for ports 135–139, 445, and 1433. He then narrows down the number of fields to source IP address and number of packets per connection, and sorts the connections by the number of packets in descending order. The screenshot below displays the list of IP addresses that are sending out the greatest number of packets.

David deduces that these hosts may be infected and asks to examine them. While he is waiting for access, he looks up their operating systems and lists of possible vulnerabilities in the RNA network map to look for other hosts on the network that may also be vulnerable to the same exploit. If he finds them, he will check if they are also infected by inspecting traffic levels on their network segments. He will also give this list of hosts to the server operations group so the machines can be patched against the original exploit.

When he eventually gets access to the systems generating a high amount of traffic, David sees that they were indeed infected. He discovers that a USB flash drive containing infected files was connected to a server, and this started the infection. He makes a note to contact the server group and let them know of what happened, so they can ensure that USB drives are scanned before being connected to servers.

In summary, this narrative illustrates how Sourcefire RNA has multiple capabilities that can help a Security Analyst do his job more effectively:

- Provides recommendations on how to best tune the IPS
- Detects applications and services on critical servers that can potentially violate policy
- Monitors network segments for anomalous traffic patterns
- Tracks host and network flow data that can save time during an event investigation

CONCLUSION

The Maginot Line was intended to defend France against surprise attack by providing overwhelming firepower. Unfortunately, the Line was designed and operated in an environment where it was obsolete and could not properly defend against the speed and flexibility of the German Blitzkrieg. Similarly, many enterprises still operate static security systems that were designed in a prior era. They are not equipped to defend against the dynamic threats that are now commonplace.

Security analysts are overwhelmed by many responsibilities and need to rely on an automated system that can provide context and help them to differentiate the legitimate threats from all other traffic. Sourcefire RNA supplies this context to security analysts and helps to automate the process of securing the enterprise.

To learn more about Sourcefire RNA and the Sourcefire 3D System, visit our website at www.sourcefire.com or contact Sourcefire or a member of the Sourcefire Solutions Network™ today.

Initiator IP	Initiator Pkts	Responder Pkts	Initiator Bytes	Responder Bytes	Count
192.168.18.245	90420	89080	4338600	3565272	16492
192.168.1.129	276827	3100	14815874	439396	13004
192.168.16.253	7806	2688	455080	114996	2796
192.168.16.249	4824	3219	257888	142080	2415
192.168.16.254	5267	2568	210680	104164	2184
192.168.18.205	19661	7827	1179660	313080	2151
192.168.21.88	148757	138499	17826765	20434204	1955
192.168.30.32	18116	10652	1021620	426080	1537
192.168.16.243	3458	1112	138320	45376	1535