

Network Monitoring for Intrusion Detection

Karen Kent Frederick 2001-08-28

Network Monitoring for Intrusion Detection

by *Karen Kent Frederick*

last updated August 28, 2001

In the world of intrusion detection, we tend to focus on detecting attacks and clearly anomalous activity. However, another important component of a complete intrusion detection solution is basic network monitoring and traffic analysis. Network monitoring collects information on connections, while traffic analysis allows us to see what services are being used on a network and to compare that against the activity that we should be seeing. This allows us to identify unauthorized services being used within a network, as well as gaps in network perimeter defenses. By combining basic network monitoring and traffic analysis with other intrusion detection methods, you can establish better overall security. In this article, I will present an introduction to network monitoring and traffic analysis.

An intrusion detection solution that simply looks for attacks is missing a key component: the identification of unauthorized and undesirable traffic that is not obviously malicious. Such traffic could be generated by misconfigured equipment, or a host that is providing or using unauthorized services, either accidentally or purposely. Perhaps a mistake in a firewall rule set is allowing certain types of external traffic to enter your internal network. Perhaps trojans are communicating through IRC. If you are unaware that these problems exist, you are providing easy targets for attackers. This may lead to more intrusion attempts and a higher chance of a successful intrusion occurring. So identifying unauthorized traffic can be critical to providing and maintaining a good level of security for your environment.

Network Monitoring

In order to perform basic network monitoring, you need to collect information on traffic at various points within your network. Although you definitely want to pay attention to your network borders, you should also look at purely internal traffic. If you have internal hosts providing unauthorized services for other internal hosts, you will miss this traffic if you only look at your borders. Various tools, including sniffers and packet capture utilities (such as tcpdump) and some intrusion detection systems (such as NFR Security's NID) and Internet Security Systems (ISS's RealSecure), can be used to gather the appropriate information on traffic.

When you are preparing to collect data on network traffic, it's very important to

get only the minimum information needed. Unless you are in an environment with a very low volume of traffic, trying to store the headers and contents of every packet will be far too resource-intensive. Besides, you can do initial traffic analysis simply by looking at a handful of the characteristics of packets and ignoring the packets' payloads altogether. Later, you can gather detailed information on particular services or hosts that require further investigation. One note of caution on this - by ignoring payloads, you will be unable to verify that the communications occurring on that port match the expected service. For example, a trojan could use TCP port 21, making it appear during your analysis that FTP activity is occurring.

In most environments, you will want to focus your analysis on TCP, UDP and ICMP traffic. Of course, you may also be interested in identifying protocols other than these that are in use on your network. But for the sake of our discussion, we will stick to these three protocols, as we are primarily interested in what TCP and UDP-based services are being used. The most fundamental elements you should examine during traffic analysis are:

- Source and destination IP addresses.
- For TCP or UDP traffic, the source and destination ports.
- For ICMP traffic, only the contents of Destination Unreachable (ICMP type 3) messages. These will be useful in identifying failed and blocked connection attempts.

You'll want to collect these traffic elements for a period of time. Depending on the volume of traffic that you see, you might start by collecting traffic for a few minutes, hours or days. Remember that the collection can be very resource-intensive, so it's best to start with a small sample and then collect larger samples later as resources permit. It's a good idea to either collect information over a long period of time or to do short collections repeatedly and on different days and times to get a more accurate sampling of data. This also allows you to create a baseline of sorts. It is very important to realize that by doing monitoring periodically, you cannot be sure that you are catching everything going on. You will only catch activity that occurs during the monitoring period.

Depending on what tools you have in your environment, you may be able to capture information on TCP traffic by packet or by connection. Ideally, you'd like to do both. By looking at TCP connections, you'll see what traffic is being permitted. However, if you only look at successful connections, you are missing valuable information on failed connections. Capturing data on TCP packets and querying it to reveal unsuccessful connections (those that lack the full TCP three-way handshake), as well as examining the contents of ICMP Destination Unreachable messages, can provide additional valuable data for analysis.

There are four different types of TCP activity that we should consider:

- Successful connection: three-way handshake is completed successfully.
- Failed connection: client gets no response to a connection attempt. The client often does two or three retries with brief delays between each.
- Blocked connection: client gets a negative response to a connection attempt, such as a TCP RST packet or an ICMP host unreachable or port unreachable packet.
- Aborted connection: three-way handshake is started but never completed (client crashed, network connection was lost, etc.) In most cases, you should ignore these, unless you're seeing large numbers of these.

Traffic Analysis

Once you've collected information from a particular point on your network for a period of time, the real fun begins - performing traffic analysis on the data. You should approach this differently depending on what your environment is like. If you permit everything that isn't explicitly denied, then you should look for those items that are explicitly denied. If you deny everything that isn't explicitly permitted, then you'll need to look for those items that aren't explicitly permitted. Of course, in many environments, no single person will know what activity is really unauthorized, particularly on a server-by-server or host-by-host basis. In those cases, your best approach may be to create a report that shows all types of activity occurring; then consult with the appropriate people to determine which activity is unauthorized.

When you are starting your traffic analysis, you should use a statistical approach. You certainly don't want to look at each packet; instead, you should focus on discovering which servers and ports are being used. So we want to know that there was an IRC connection from client X to your primary DNS server, but during the initial analysis, we don't care how many times that connection occurred, or how many packets were sent in that connection. Likewise, with a connectionless protocol such as UDP, we are initially interested in knowing what server ports were used, but not how many packets each received. By summarizing the activity and ignoring the finer details, we can reduce the amount of material for analysis to a manageable amount. Once we find the most significant details, we can return to the original traffic data to get more information.

At this stage, you're going to want to do some data validation on the IP addresses. Any traffic that has suspicious or improper IP addresses should be examined more closely. For example, packets coming from the Internet should not have reserved IP addresses. You should also check to confirm that all addresses for internal hosts are in your defined internal IP ranges. This sounds simple, but in some environments, there are many different IP ranges in use, and

it may be difficult to develop an accurate list of valid internal IP addresses. Checking internal addresses can help you in identifying misconfigured or unauthorized equipment and other items that need to be investigated. The IP address information may also be added to your network intrusion detection system so that it will alert when improper IP addresses are used.

There are many different ways that we can look at the data. One way that we can use the data is to find port scans and host scans against your internal hosts, particularly those that might occur over a long period of time. This can be done by sorting the data by client, then looking at how many server ports it attempts to contact, or how many different servers it attempts to contact. Those with the highest values should be investigated further to explain the activity. When using data for this purpose or others described here, you're certainly not going to be able to view it all manually. Instead, you should put the data into a database and use queries to crunch the numbers and produce a report for you.

One difficulty you will encounter is how to handle UDP packets. Because UDP is connectionless, it is often unclear which host is the server and which is the client. You can try to query the data in order to find out which host sent the first UDP packet to the other, but this may be difficult, and it will produce some inaccurate results because it will miss packets sent before the monitoring begins. You may achieve better results by sorting the data by source host and by destination host. If you find a host that always uses the same UDP port number when communicating with several other hosts, it's a fairly safe bet that it's acting as a server. You'll need to indicate in your database which entries are client to server and which are server to client.

When you're ready to analyze the traffic, there are many different ways to begin. One method would be to separate the UDP traffic and TCP connections and connection attempts into three groups:

- External client and internal server. In most cases, the source addresses and ports are irrelevant; all you will want to examine initially are the destination addresses and ports.
- Internal client and internal server. Although you'll want to focus on destination addresses and ports, source addresses may also be quite important.
- Internal client and external server. If outgoing connections are unrestricted in your environment, you can skip this data when performing your analysis. If outgoing connections are restricted, you will probably want to focus most on the destination ports.

The ICMP traffic should be placed into the same groups, except that the destination address of the ICMP packets should be aligned with the TCP client

address.

Within each group, sort the TCP information by server port or by server address and server port. This will show you what services are being used and what servers are running them. If you aren't familiar with any of the port numbers, there are many good references on the Internet that can tell you what typically listens at that port number. One of the best is the port list from Neohapsis (<http://www.neohapsis.com/neolabs/neo-ports>). If you still can't identify the port number, you may need to check the host in question or gather more details on the traffic in question in order to determine what it is.

Next Steps

Once you determine that suspicious or unauthorized activity is occurring, obviously you will want to determine why the activity is occurring. You may want to use the original traffic logs to get more information involving the services and hosts in question. If your perimeter security devices are permitting traffic through that they shouldn't, you'll want to review firewall rule sets, router ACL's and the like. If hosts are providing unauthorized services, you'll want to check those hosts for misconfiguration or signs of compromise, depending on the situation.

Based on the results of the traffic analysis and subsequent investigations, you should tune your intrusion detection sensors. This might include items such as:

- Alerting on traffic from internal hosts that use unauthorized addresses
- Alerting when a host uses or attempts to use an unauthorized service
- Alerting when a host attempts to connect to particular internal or external hosts
- Optimizing by enabling or disabling code that performs full analysis on particular protocols or services, in order to save IDS sensor resources.

By adding basic network monitoring and traffic analysis to your existing intrusion detection framework, you can improve the overall security of your environment. Network monitoring and traffic analysis are certainly useful as an auditing method, and they can permit you to reduce the chance that a successful intrusion will occur. They can also assist you in tuning your intrusion detection sensors more effectively.

Karen Kent Frederick is a senior security engineer for NFR Security's Rapid Response Team. She is currently completing her master's in computer security through the University of Idaho's Engineering Outreach program. Karen holds several certifications, including Microsoft Certified Systems Engineer + Internet, Check Point Certified Security Administrator, SANS GIAC Certified Intrusion Analyst, GIAC Certified Unix Security Administrator, and GIAC Certified Incident Handler. She is one of the authors and editors of "Intrusion

Signatures and Analysis", and she is a contributing author to the upcoming "Handbook of Computer Crime Investigation".

[Privacy Statement](#)

Copyright 2006, SecurityFocus