

Evaluating Network Intrusion Detection Signatures, Part One

Karen Kent Frederick 2002-09-10

Over the past several years, a number of academic and commercial entities have conducted evaluations of various network intrusion detection (NID) software, to determine the overall effectiveness of each product and to compare the products to each other. Many system administrators and security analysts are also responsible for conducting their own evaluations of NID products, in order to choose a solution for deployment in their environments. NID evaluations typically include some rough indication of the relative quality of each product's signatures. However, high signature quality is critical to achieving a good NID solution, so the importance of accurately evaluating signature quality cannot be stressed strongly enough.

In this series of articles, we will present recommendations that will help you to evaluate NID signatures. As you shall see, properly testing NID signatures is a surprisingly complex topic. We will begin by discussing some of the basics of evaluating NID signature quality, and then look at issues relating to selecting attacks to be used in testing. Although you may not necessarily perform hands-on NID testing and evaluations, the information presented in this series of articles will give you the knowledge and the facts to get the most out of published reviews and comparisons of NID signatures. Note that we assume that the reader is already familiar with the basic concepts and principles of network intrusion detection.

NID Evaluation Basics

You've just been asked to choose a network intrusion detection solution for your environment. To find out which products are the best for your organization, you may acquire demo copies of the products and test them yourself. If you are going to do your own testing, you might utilize an isolated test environment, which contains "attacker" hosts that you launch attacks from, "target" hosts that the attacks are directed at, and hosts that run the NID software demos and monitor the test network segment. This is by far the best way to evaluate products, because you can customize all the testing to focus on the features that are most relevant in your environment. Unfortunately, such testing may require you to have good knowledge of not only network intrusion detection principles, but also networking, various operating systems, exploit and vulnerability characteristics, and other areas of security and computing.

What can you do if you lack knowledge in these areas or, more likely, the time and resources to perform such testing? You are likely to instead read published reviews and comparisons of NID

products, and to use the data and conclusions from them as a basis for your evaluation. Reviews can be a great source of information, and can certainly save you a lot of time in your product selection process. However, such reviews are often focused on the overall usability of the product and give only a brief and limited look at signatures. Unlike the anti-virus industry, where you can expect products to have similar signature capabilities and coverage, network intrusion detection vendors have vastly different signature approaches, methodologies and coverage. The quality and completeness of signatures is one of the most important aspects of a NID solution, if not the most important; it is also the most difficult to determine.

Little academic and commercial work has been done on accurately evaluating the quality of intrusion detection signatures. In NID software evaluations, typically a number of known attacks are executed on a test network that the NID sensors monitor, and the tester counts how many of the attacks the IDS detected. Although this can sometimes provide a general sense for the quality of the signatures, it does not give us an accurate evaluation of the signatures, nor does it often provide valid data for comparing the quality of various signature sets. To expand NID software evaluations to include comprehensive testing of signature quality, several aspects of the testing must be changed. We shall examine each factor in depth in the following sections.

Signature Testing Approach

Network intrusion detection vendors currently say that their products contain hundreds or even thousands of signatures. Yet NID evaluations are still occurring (and are being published) that utilize fewer than five attacks, which tests only a few signatures. We really cannot say much about the quality of a signature set if only three signatures in a thousand are tested. Making this situation even worse is that there are often many ways in which a single signature can be triggered. Because there are so many signatures in a single NID, it is impractical to exhaustively test all signatures; the time and resources required would be far too great for nearly any organization. Further complicating this problem is the fact that most vendors have closed-source NID signatures, and the NID documentation often provides little specific information on the characteristics of the attacks that the signatures are looking for.

So how can we test the NID signatures? The trick lies not in focusing on the signatures, but instead in focusing on the potential attacks. We may not have much information on the characteristics of signatures, but there is a wealth of information available on attacks, vulnerabilities and exploits. We should determine what threats are important for our

environment, and then design and execute NID tests that address those threats. Since there are many potential threats, we must still define the scope so that it is practical to perform the testing. In order to do this, we can look to accepted software testing principles, and good old common sense.

Because it is not feasible for us to test using every possible attack that could occur, we should instead focus on choosing a representative sample of attacks and testing the NID's signatures with just those attacks. This should give us a very good idea of the quality of the signatures - if two conditions are met. First, we must be very careful to choose appropriate attacks and a sufficient number of attacks. Second, we must perform the testing properly. These points sound simple, but in reality they are very complicated. Many NID evaluations, including commercial reviews and product comparisons, often violate one or both of these conditions; consequently, their test results may be biased; therefore, their conclusions may be incorrect. We'll start our examination of this by looking at how attacks should be chosen.

Attack Selection Criteria

Generally, the most complicated aspect of NID signature testing is choosing and generating the attacks that we hope will trigger alerts by matching NID signatures. Traditionally NID evaluators have chosen a number of exploits - anywhere from a few to a few hundred - to incorporate in their testing. These exploits are often chosen rather arbitrarily: it's common for the tester to look for well-known exploits and just test the NID using those. A more structured way of choosing attacks is to establish a way of classifying attacks and then select one or more attacks that fit into each category.

Although this approach is somewhat effective, attack selection needs to take other factors into account: the threats and risks inherent in your environment. An attack that was released five years ago is more than likely not really a potential threat to your systems. If your organization utilizes seven protocols (such as FTP, HTTP and SMTP) on the network segments that the NID will monitor, the attacks you choose should address all of these protocols, and should perhaps exclude all other protocols. If fifty percent of the attacks that are attempted against your systems are targeting Web servers, approximately half of the attacks you choose to include in your testing should be HTTP-based. To provide an accurate assessment of NID signatures' quality, the attack suite must accurately reflect the nature of the threats and risks.

By combining your knowledge of the threats in your environment with a methodology for

classifying attacks, you can select a well-balanced attack suite. Most attacks (and most NID signatures, for that matter) are tied to a particular network protocol (IP, TCP, UDP, ICMP, etc.) or application protocol (FTP, HTTP, SMTP). Because nearly every attack is linked to a particular protocol, it makes sense to classify attacks first by protocol. So you can develop a list of the protocols that are of interest in your environment, and use this list as the first step in choosing which attacks to run.

The next step is to determine which attacks should be used for each class. The first point to be considered is, are you interested in only those attacks that match your systems, or are you interested in all attacks? For example, if you only use Apache Web servers, do you care about Microsoft IIS attacks being detected? If not, then your task is much simpler - for your HTTP attacks, only consider those that target Apache. Otherwise you will have to consider attacks against all types of Web servers.

Another important criteria to consider is the age of the attack. Old attacks pose little risk in most environments, while brand-new attacks may be of great interest. Consider both the age of the attack and the operating system or application version that it targets. An exploit could have been published two years ago for an FTP server that was itself released five years ago.

The Process of Choosing Attacks

Now that you've narrowed down the protocols that you're interested in and the age of attacks that you'll consider, it's time to start choosing which attacks you will use. You want to perform testing efficiently - that is, to test as much of the product as you can with as little effort as possible. Therefore, you want to select attacks that test different aspects of the NID signature set. If you choose five FTP attacks, but all of them utilize a buffer overflow in the `CWD` command, then you haven't done a good job of testing the breadth of the FTP signatures. It would be much better to choose five FTP attacks that each target a different type of FTP vulnerability. The attack suite you use for testing should also incorporate IDS evasion methods that are relevant for your environment and for the network and application protocols that you are interested in.

So how do you find out about the exploits, vulnerabilities and IDS evasion methods that pertain to the protocols you've selected? The [Common Vulnerabilities and Exposures \(CVE\) project](#) has compiled information on thousands of vulnerabilities into an organized and searchable form.

[SecurityFocus's Vulnerability Database](#) also contains detailed information on thousands of

vulnerabilities, including links to copies of exploit code. These resources have pointers to additional information, such as security announcements from vendors and Bugtraq postings involving vulnerabilities and/or exploits. Both CVE and the SecurityFocus Vulnerability Database can be used to assist you in determining which attacks are good candidates to be included in your attack suite.

Summary

Determining the quality of network intrusion detection signatures is a complex and technically challenging process, requiring the NID evaluator to possess a range of technical skills, as well as adequate time and resources to perform the signature testing. The foundation of a good NID signature test strategy is to choose a representative sample of attacks to test as many NID signature capabilities as possible in the allotted time. By thinking about the threats that your environment faces, and which threats are relevant for you in terms of intrusion detection, you can establish a list of network and application protocols (and specific applications, where relevant) that serve as a framework for the development of an attack suite. Sources such as CVE and the SecurityFocus Vulnerability Database hold vast amounts of information on exploits, vulnerabilities, attack techniques and other things relevant to NID testing. In the next article in this series, we will conclude the discussion on choosing attacks, and then provide recommendations for generating attacks and for creating a testing environment.

To read **Evaluating Network Intrusion Detection Signatures, Part Two**, click [here](#).

Karen Kent Frederick works for EDS, specializing in intrusion detection. She is a graduate of the University of Wisconsin-Parkside and is currently completing her master's in computer science, focusing in network security, through the University of Idaho's Engineering Outreach program. Karen has over 10 years of experience in technical support, system administration and information security. She holds several certifications, including SANS GIAC Certified Intrusion Analyst, GIAC Certified Unix Security Administrator, and GIAC Certified Incident Handler. She is one of the authors of "Intrusion Signatures and Analysis" and "Inside Network Perimeter Security: The Definitive Guide to Firewalls, Virtual Private Networks, Routers, and Intrusion Detection Systems".

[Privacy Statement](#)

Copyright 2006, SecurityFocus