



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials Bootcamp Style (Security 401)"
at <http://www.giac.org/registration/gsec>

Controlling Hostile Network Traffic in Excessively Heterogeneous Environments

Mike Denka

May, 2002

GSEC Practical Assignment version 1.3

Abstract

Certain IT environments must tolerate tremendous diversity in the traffic patterns on their networks. Internet Service Providers (ISPs), collocation facilities, large educational facilities and other organizations providing unlimited internet access to customers or users fall into this category. Guaranteeing unlimited or nearly unlimited internet access implies excessive heterogeneity in terms of the categories, volumes and endpoints of user traffic on these networks. Excessive heterogeneity can make hostile traffic quite difficult to detect. Providers of internet service, however, are not excused from the obligation to protect their networks as much as possible from the threat of hostile traffic. Indeed, since a large percentage of hostile traffic emanates *from within* these networks, service providers have a special responsibility to protect the rest of the internet from hostile traffic originating within their borders.

This discussion will examine various tools and procedures that may be used in these excessively heterogeneous environments to evaluate traffic for its threat potential. These tools and procedures are not fundamentally different from those used in more homogeneous environments. The focus will be the approach taken with tools and procedures that may be more conducive to threat detection and prevention in excessively heterogeneous environments where the line between expected and unexpected events is hazy at best, and where false positives rule the day.

Characteristics of an Excessively Heterogeneous Environment

Most traditional corporate, private or organizational networks are largely homogeneous in terms of internal traffic patterns and utilization. In typical homogeneous networks, much of the network is expected to be private. Most company documents, data and other resources should be exempted from public view.

Inevitably more and more private organizations are linking up to the public internet for access to outside services, to reduce communication expenses and to promote their own products and services. Nevertheless, their public exposure is normally restricted to a limited set of services and access policies. Depending upon an organization's mission, size and numerous other defining factors, it may operate one or more public services such as web, dns or email. Some organizations may use virtual private networking to communicate between remote offices and staff, between suppliers and customers. And many of these organizations permit more or less limited internet access to internal staff. But the security policies that these organizations set up should strictly define the types of

inbound and outbound traffic and connections permitted between the public internet and internal network components. While the diversity of this public traffic can be complex for many organizations, it is finite, containable and ultimately definable.

Service providers live at the other end of the predictable traffic spectrum. For the purposes of this discussion, we will define service providers as an organizational class providing unlimited or nearly unlimited internet service to a large and diverse group of users. These offerings can be in the form of dialup or network connectivity via WAN (Frame Relay, ISDN, fractional T1, T1 or above) or LAN (xDSL, Ethernet) connections. Service providers can be traditional Internet Service Providers (ISPs) offering connectivity for sale; collocation services, offering high speed connectivity and value added services for sale; colleges and universities offering full or nearly full access to students and faculty, or any other organization providing full or nearly full access to a group of fundamentally unknown and largely uncontrolled users.

The environment of the service provider is likely to be “excessively heterogeneous” in terms of the type and quantity of traffic expected on its user access network segments. For dialup and connected customers or users, nearly any kind of protocol in any direction can be normal and acceptable. It is impractical to use access lists to filter most classes of traffic to or from these customers since the customer is purchasing (or at least feels entitled to) “internet access” as a product and is expecting “full” access without restrictions. Similarly, it is not practical to use firewall protection on these user access network segments.

Distinguishing hostile traffic that imposes high risk from less risky and acceptable traffic on these networks can be a challenging proposition. Network Intrusion Detection Systems (NIDS) running on these networks have a far more difficult time separating true and significant hostile events from normal or mildly hostile but largely irrelevant events. Signature based NIDS that base their decisions about hostile traffic upon recognizable patterns in network packets will generate unmanageable numbers of false positives. Anomaly based NIDS are practically useless since it is nearly impossible to formulate accurate behavior models on such heterogeneous networks.

Still, the service provider has private data and services that must be protected. And the service provider has the responsibility for keeping his and his customers’ network free from internal and external threats and attacks. The service provider has the additional responsibility to protect the rest of the public internet from potential misbehavior, intentional or otherwise, from his customers. So a service provider must be especially wary of threats traveling in all directions across his network.

Network Topology Considerations

Before going further let’s generalize some sample service provider network topologies. The specifics of any service provider’s network will vary widely, but there are a few common properties that should exist. For example, at least one internet connection is required. If only a single connection is available, all customer connectivity and both

public and private services obtain access through this single connection. Larger ISPs may have multiple internet connections through one or more higher tiered service provider or peering partner. Some may use one or more of these connections for customer access and one or more different connections for public and private services. As we will see below, the problem of isolating hostile traffic becomes less cumbersome as providers approach the multi-homed, multi-network model.

Next, we will postulate that all service providers will offer some public services. Nearly all will provide domain name services (DNS). Most will also offer mail services and most will offer web services. Some will offer additional services such as news, ftp, time, and more. Finally, most service providers will support private services for internal use that may require limited internet access.

Figures 1 and 2 represent sample topologies for two hypothetical networks. Figure 1 depicts a sample topology for a small service provider where a single connection exists to the internet over which all connected traffic must travel. Figure 2 exemplifies a larger service provider with multiple connections to the internet. For the purposes of this discussion I will refer to providers whose networks can be roughly approximated by figure 1 as 'type A' providers. Those networks that more closely resemble the multi-homed, multi-network structure in figure 2 will be termed 'type B' providers.

Generic Type A Network

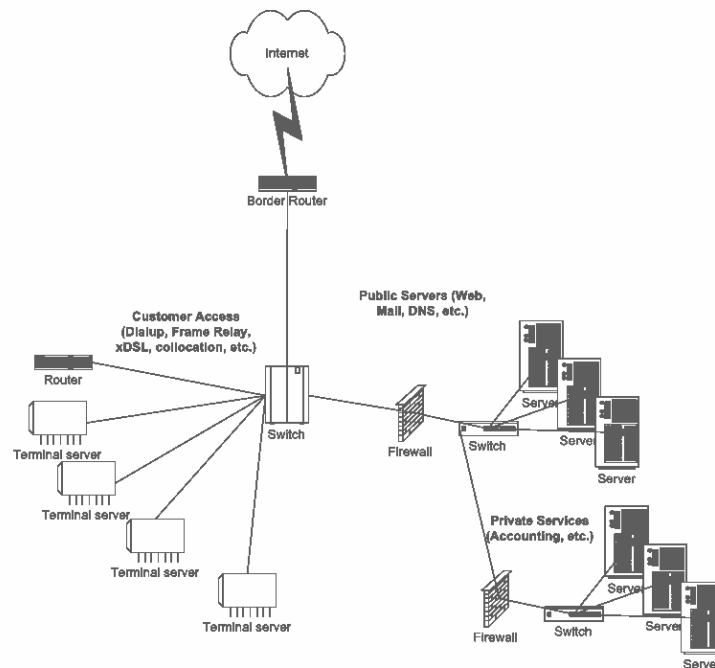


Figure 1

Generic Type B Network

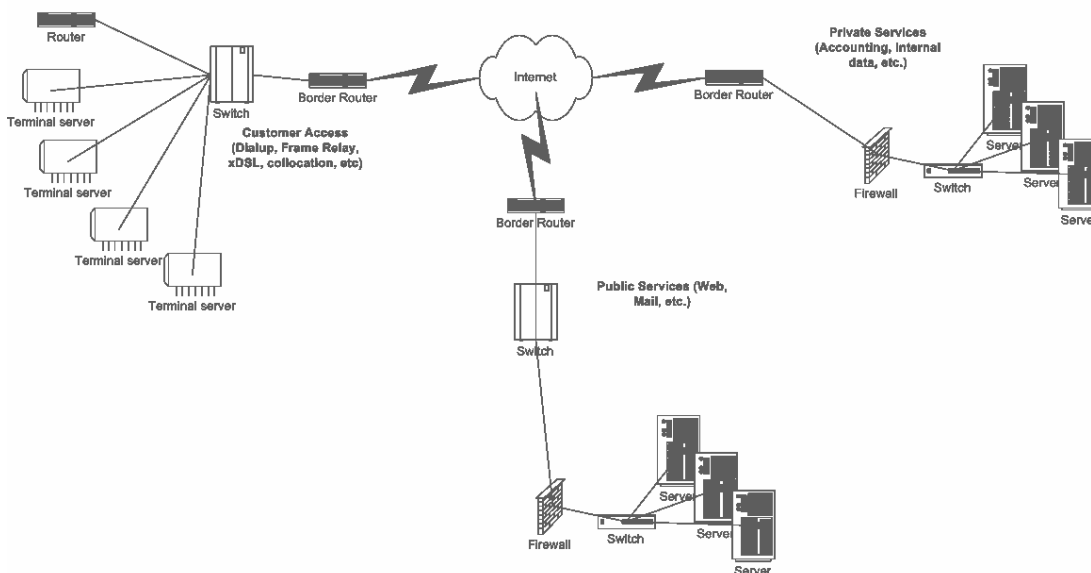


Figure 2

Since all traffic must ingress and egress the type A provider's network over the same connection in figure 1, this scenario represents the more challenging case of excessively heterogeneous traffic. The type A topology has private and public service data entangled inside the border with the excessively heterogeneous data from the customer access segment of the network. Protecting and monitoring the confidentiality and integrity of that data destined for internal network segments becomes proportionately more difficult (depending, to some extent, upon NIDS sensor placement as discussed in the "Network Based Intrusion Detection" section below).

As a service provider grows and adds multiple connections to the public internet, that provider's network topology has the potential to approach the type B representation in figure 2. The topology represented in figure 2 shows confinement of the excessively heterogeneous traffic to a single border router. So the private intranet and public service networks in the type B topology are basically homogenous. This homogenous space is more easily monitored and hostile traffic is more easily detected and inhibited using standard containment and intrusion detection mechanisms such as access lists on the border router, network and host intrusion detection systems, and firewalls to tighten up access to internal facilities. The type B topology has only one excessively heterogeneous network segment to worry about: the customer access segment.

The first goal for any type A or near type A topology should be to first become more type B. Once that goal is achieved, protecting private intranets and public services becomes a

typical homogeneous network security issue. Then, only the excessively heterogeneous public access network segments need to be treated specially. We will discuss conversion from type A to pseudo type B below under “Filtering at the Border”.

The Role of Policy

The security policy of a service provider should play a major role in helping to isolate hostile traffic on his or her network. This security policy should be developed in much the same way as it would be for any organization with connections to the public network. A fundamental difference in the way a service provider’s policy is developed has to do with differences between the customer side of access (public access policy) and the access and policy decisions affecting the provider’s staff and internal network segments (private and public service access policy).

The public access policy must allow maximum flexibility and, to the extent possible, full access rights to the internet. At the same time, public access policy must restrict that behavior which would threaten other connected systems. Public access policy should be included in the Terms and Conditions agreed to by all customers or users. Public access policy should be as explicit as possible. For example, it should be expressly forbidden to perform port scans of any network (without express written permission of the owner). Sending Unsolicited Commercial Email (UCE or spam) is usually prohibited. IP address spoofing and other forms of deception should be prohibited explicitly. And, of course, direct attacks in the form of Denial of Service (DOS), Distributed Denial of Service (DDOS) and any form of system compromise should certainly be prohibited. Terms and Conditions should prohibit all forms of harassment or illegal behavior to cover the general case, but the more specific the security policy descriptions in the Terms and Conditions, the easier it will be to match your public access policy to your rules and filtering profiles as described below.

Private and semi-private access policies are not fundamentally different from those defined by traditional organizations with more homogeneous network traffic. These policies relate to permissible behaviors and traffic patterns on the protected portion of the service provider’s network and will be used to transcribe rules and filters pertaining to those network segments.

After these policies are formulated and included in the Terms and Conditions and the organization’s security policy, the policies can be used to help define rules and filters on network components. Filters or access lists will be applied to routers and rules to firewalls and intrusion detection systems that we will discuss in more detail below.

The important point here is that there is a strong link between policy, both public and private, and the procedures we use to detect and limit hostile traffic on our heterogeneous networks. It is important to maintain that relationship. If a provider makes substantial modifications to his or her security infrastructure, the policies known to staff and customers alike may need to reflect these changes. Likewise, if changes are made to Terms and Conditions or security policies are radically modified, the question should be

asked whether similar changes need to occur in security monitoring processes. This link between policy and process keeps network users aware of expectations and it keeps security administrators more attuned to the types of network traffic that may be hostile. Since, as we noted before, almost anything can be legitimate on excessively heterogeneous networks, it is especially important in these environments to be able to distinguish as accurately as possible any clues to traffic that is illegitimate. The policy decisions we make often illuminate many of those clues. Examples of this relationship between policy and process will be given shortly.

Filtering at the Border

The best place to nab a foreign terrorist, smuggler or other miscreant out to do you harm is at your national border. It's much more difficult to apprehend the n'er-do-well once she has infiltrated the general population. This same philosophy applies to any network. If you can keep most hostile traffic from getting past your border, the job of identifying hostile traffic within your network becomes proportionately simpler. Unfortunately, it is at the network border that the differences between excessively heterogeneous networks and homogeneous networks are most predominant.

Small homogeneous networks can do a pretty good job protecting themselves with a simple firewall on the border (emphasis on "pretty good" – acknowledging that a firewall is not to be considered a complete security solution). In such a network, the security administrator/owner often knows precisely the types of traffic that are permitted and prohibits the rest by configuring the rules in the firewall at the perimeter.

Border routers can also be used to exclude many common types of traffic that simply have no business entering the simple homogeneous network. No reason for telnet, ftp, RPC, NFS or Napster traffic to be inside your network? Cut it off at the border. No need to allow anyone to ping your internal hosts? Stop those messy echo requests at the border. Cut down on a lot of needless traffic and make the job of monitoring your internal network traffic simpler by an order of magnitude or so.

Not so easy, though, for the service provider managing an excessively heterogeneous network. Especially as the service provider's network approaches the type A topology described above. Since the type A service provider has all customer traffic entering and exiting at the border, and since we've already acknowledged that the customers or users of this service provider's network are expecting 'full access', what can our provider do at the border?

One partial solution for the extreme type A provider with meager resources might be to split his border to try and appear more type B. This could be done, for example, by adding an additional interface to the border router and running all private segments and public services through the added interface (Figure 3). Adding interfaces at the border (or, if resources permit, adding internal routers separating internal network segments) will allow more granular filtering capabilities for each interface or routed segment. Now the private and public services segments have become homogenous networks and

managing hostile traffic on them has become proportionately less complex. There is still at least one excessively heterogeneous customer access segment that must be managed, but that segment no longer includes private data and equipment. The down side is that now more physical equipment may be required to detect and limit hostile traffic across all the new segments. More routers or high speed interfaces, for example, and perhaps more NIDS sensors (see section on sensor placement under “Network Based Intrusion Detection”, below).

Modified Type A Network to Pseudo Type B Network

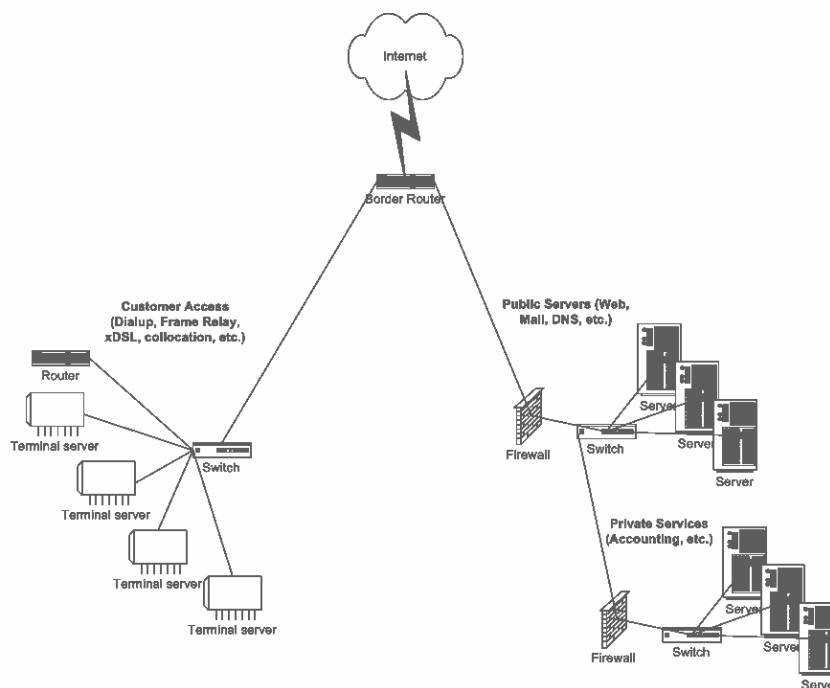


Figure 3

Nonetheless, some potential and mandatory border filters should be considered for all excessively heterogeneous networks and network segments. One type of traffic that can and should be filtered even in the most unrestricted environment is defined by RFC 2267 [1]. This RFC suggests a mechanism for preventing spoofing from RFC 1918, reserved address space based attacks. More importantly, it suggests a mechanism for preventing most types of spoofing attacks originating from the service provider’s network. It is imperative that all service providers implement these spoofing prevention techniques either on their border or access routers. Applying these filters to both border and access routers will reduce somewhat the amount of traffic transiting their networks. But if all service providers religiously block spoofing from within their networks, a large number of DOS attacks can be prevented or easily blocked.

Another once common attack that can be easily blocked at the border without affecting legitimate traffic into a heterogeneous network is the “smurf” type attack. Denying

directed broadcasts at the network border is a simple mechanism to prevent internal systems from becoming innocent intermediaries in this type of DOS attack [2].

TCP small services, chargen, echo and discard, for example, have been used in simple hostile attacks. These can usually be safely blocked at the border since they were developed for testing and should not be accessed from the public internet for legitimate purposes.

Blocking most other types of data at the type A provider's border, however, can be tricky. The nature of excessively homogeneous networks, as we have said, is that they are expected to provide *full internet access*. But if a provider would like to argue the definition of full access on some finer points, there are a few services that I believe may qualify as exceptions to the full internet access policy and might be considered for filtering at the border.

The services that I believe could qualify as exceptions to the full access policy have a few common characteristics. First, due to the nature of the information they communicate, these services should not be used across the public internet unless they are encrypted in a encrypted tunnel or Virtual Private Network (VPN). Second, these groups of services are very often installed on systems by default with vulnerable default configurations. Third, these services are very often used as the starting point or attack point of hostile traffic.

One such group of services that might be considered for blocking at the border of excessively heterogeneous networks might be NetBIOS services – udp and tcp ports 135 through 139 and 445. These ports are used in numerous Windows enumeration attacks and are often utilized in successful compromises of Windows systems [3]. Since the vast majority of systems inside dialup service providers' networks are windows systems, and since NetBIOS services contribute in large part to successful attacks of Windows systems, and since NetBIOS services are configured insecurely by default on most Windows systems [3] and it is not practical to reconfigure all end-user systems, and since the exchange of data across the internet can be accomplished in other more traditional ways or more securely over VPNs, it might be argued that NetBIOS services should be blocked at the customer access border.

Similar arguments can be made for other services that are infrequently used for productive purposes across the public network (outside VPNs) but frequently used for hostile purposes. RPC, SNMP, NFS and NIS are examples of such services.

If a provider makes a decision to block one or more of these or other services at her border, she should make that decision very clear in her Terms and Conditions as part of her public access policy. That way, customers are aware of the limitations of their internet access (provided, of course, they actually *read* their Terms and Conditions).

Filtering traffic at the border can reduce the load on network intrusion detection systems, firewalls and internal routers by reducing total traffic, blocking potentially hostile traffic and reducing heterogeneity. Border filters can also be modified easily in the event that

conditions change. Those filters could be easily changed in response to customer demands – if, for example, everyone demands to open their shares to the public.

Tools of the Trade

Once inside the border, we concentrate on procedures and tools to monitor the internal traffic. We need to consider mechanisms to protect assets and to detect likely attack patterns and signatures in a sea of heterogeneous traffic. Then we need to devise the means to mitigate attacks and their affects as soon as possible with minimal impact on the network.

We will discuss specific examples from a few categories of tools used to perform these activities. To monitor the traffic on the wire, we will discuss a NIDS and two traffic monitoring programs. Each of the traffic monitoring programs focuses on different aspects of network traffic. Firewall protection is useful for encapsulating the protected intranet where the private and public service assets are housed. But that portion of the network has now been homogenized. Firewall protection is not practical on our excessively heterogeneous segments. These segments, as we have said, demand full or nearly full access. There is nothing for a firewall to block on an excessively heterogeneous public access network segment.

Intrusion Detection

An important set of tools that should be central to any security administrator's bag of tricks is a robust intrusion detection package. The package should include both network and host based intrusion detection systems.

Host Based Intrusion Detection

Host based intrusion detection is run on individual hosts in the form of agents that monitor properties of files and directories and, in some cases, the occurrence of important events on their host platforms. Many of the most recent attacks on all fronts have been in the form of “. . . hybrid threat[s] that combine virus payloads with multiple, automated attack scripts against common computer vulnerabilities . . .” [11]. These hybrid attacks can launch from anywhere in a network. Host based intrusion detection packages are a necessary element of any security strategy to combat intrusions into systems that originate from unexpected network segments and that may not be blocked by firewalls or detected by NIDS. The service provider should seriously consider using host based intrusion detection systems¹ on public access servers and private servers on his protected network segments.

However, due to a large base of customer connections, the service provider cannot deploy host based intrusion detection agents on all the hosts directly connected to his customer access network segment. He may have thousands of host platforms to which he has very

¹ An example of a couple host based IDSs and instructions for using them are referenced below [4]

limited access. This is a huge blind spot on the provider's network that must be corrected to the maximum extent possible. To compensate for the 'blind host' problem, the service provider must be vigilant on other fronts. The most promising other front is the network based intrusion detection front.

Network Based Intrusion Detection

Network Intrusion Detection is the last bastion of protection on the wild edge of the excessively heterogeneous network. Unfortunately, it is the very diversity (and often the associated volume) of traffic on this segment that limits the effectiveness of network based intrusion detection systems.

The astute security administrator must monitor network traffic to compensate for the "blind spot" comprising a service provider's customer connections – the war zone where almost anything goes and where host based intrusion detection and access list filtering or firewall protection is not practical. The biggest problem in the excessively heterogeneous network segment, though, is that mostly unfiltered² network traffic hammering on this service provider's NIDS can result in a tsunami of false positive events and a mountain of alerts. How does the conscientious security administrator deal with all this data?

Before attempting to answer this question, an important consideration to ponder is NIDS sensor (that NIDS component that actually "listens to" raw network traffic) placement. Many factors must be considered here - topology, budget, traffic flow, etc. Stephen Northcutt [5] suggests that, ideally, sensors should be placed inside and outside the firewall(s). According to Northcutt, placing sensors in front and behind the firewalls often allows the analyst to determine whether attacks are actually getting through the firewall or originating from inside the firewall. It also allows the analyst to see the types of attacks that the firewall is exposed to. That's fine for the public and private segments of our network that we have affectively homogenized. But what about the excessively heterogeneous network that remains? And what if we don't have the resources to homogenize our private and public services segments?

If a service provider is unable for whatever reason to escape a type A topography (Figure 1) he would ideally require three sensors - one at the connection between the border router and the core switch to catch all traffic, including excessively heterogeneous public access traffic, passing between the service provider's network and the internet, one between the public services firewall and the internal switch, and a third between the private services firewall and its internal switch.

If a service provider is able to "homogenize" segments as suggested above, then NIDS sensors are required both in front and behind the public service firewall and another behind the private intranet firewall. Another NIDS sensor would, of course, be required at the departure point of our excessively heterogeneous customer access segment. So the

² Refer to the section titled 'Filtering at the Border' for a discussion of filtering raw customer network traffic

penalty for homogenizing the type A network is, in addition to the additional router interface, an additional NIDS sensor.

The NIDS sensors behind the firewalls will watch for intrusion attempts or anomalous traffic patterns into and out of the protected zones. This traffic should be relatively homogenous – anything passing through the firewalls should be predictable and traceable back to the internal security policy for the private and protected public segments of the network. Anomaly detection is a possibility at these intersections, especially behind the private server firewall, since patterns here should be most predictable.

In network topologies approaching type B in Figure 2, sensor placement would ideally be in front and behind each firewall on the public and private services networks and just inside the border on public access connection(s). The protected networks are not excessively heterogeneous and can be monitored like any other traditionally contained network.

Now let's return to the question of handling excessive false positives on our customer access network. The argument will be made that the customer access portion of these networks, especially the type B or modified type A models, need not be monitored at all. This is a war zone – a free for all where nearly anything goes and monitoring this mass of traffic is an exercise in futility.

But remember our service provider's responsibility. She has the responsibility to protect, to whatever extent she can, the security of her customers. And she very definitely has the responsibility to protect the rest of the internet from hostile traffic that could potentially originate from her network. So the issue becomes one of operating a NIDS on this network segment that is not made irrelevant by a barrage of false positive events and alerts.

We have found a couple of mechanisms that allow us to run a NIDS on excessively heterogeneous networks without being swamped with false positives. The first mechanism we use is to carefully trim the rule sets configured into the NIDS. The second mechanism has to do with meticulously tailoring the NIDS configuration to fit site specific characteristics of each service provider and its unique heterogeneous environment. Each mechanism has potential to reduce false positives so that serious events of interest can be emphasized.

Trimming rule sets can be done in two ways. The first way involves a labor intensive, iterative process of examining NIDS alert output on the heterogeneous network segment, determining the rules or signatures most responsible for the false positives, and removing or re-writing those signatures.

The second way to trim rule sets is to examine blocks or categories of rule sets (the Snort NIDS, for example, that will be used in further discussions below, breaks its rule sets out into compartmentalized blocks – e.g., web-cgi rules, web-iis rules, rpc rules, netbios rules, etc.) and make decisions about the desirability of using or disposing of certain

blocks. For example, some providers may decide that they will not take responsibility for web server security of their customers' systems. The argument could be made that there are thousands of false positives generated per day on some user access networks due to constant scanning for port 80 vulnerabilities [10]. The argument could further maintain that if a customer wants to offer a web service, that customer should be responsible for her own security. Of course it is also true that some, presumably far fewer, scans may find a working unprotected web server on the customer access network and compromise it. The question then is: given the compromise of a customer's web server, what are the ultimate consequences and risks to other customers and to the entire system? An ancillary question is: if you are receiving hundreds of alerts against hostile scans of port 80 and one of them results in a compromise, will you even notice it?

So the problem with trimming rule sets is that by removing rules that tend to generate false positives, the risk of false negatives increases. The false positive rate of any rule or block of rules that are candidates for removal must be weighed against the impact of ignoring real hostile traffic of the indicated type on the customer access network. Also, the likelihood of discriminating true positives from false positives in this mix must be considered. Are the occurrences of false positives so great that true positives would never be noticed anyway? Each service provider must make final rule set decisions based upon his assessment of the threats and risks involved in his or her unique environment.

The second mechanism we have used to reduce false positives from the NIDS running on our heterogeneous segment is by manipulating the NIDS configuration file to process different address blocks in different ways according to policies, expected threat types and directions, and other site specific criteria. A fine paper has been written by Roberto Nibali that illustrates this technique.

Roberto Nibali's "Introduction to Network-Based Intrusion Detection Systems Using Snort" [9] provides a detailed explanation of a strategy that can be used for configuring Snort³, a first rate open source network intrusion detection system [6], in a way that minimizes false positives and maximizes notification of serious hostile events. The techniques outlined in Nibali's paper outline specific steps to take to segregate network components, construct user defined rule types, associate rule types to policy statements, and order these rule types.

Dial up access customers, for example, can be identified as a group in your NIDS configuration and new rule types can be made to apply only to this group of users. Alerts generated against the dialup access pool can then be output to a special file or report. XDSL customers, WAN clients and directly connected customers could each receive special consideration, have special rule types governing the traffic destined to or sourced from their address blocks, and logs for each of these groups of customers or users can be individualized. Each of these sub-groupings contributes to the pseudo-homogenization of the war zone by the NIDS

³ Obtaining and installing Snort are beyond the scope of this discussion. But excellent papers have been written on this subject and are referenced below [7, 8, and 9]

While using Nibali's methodology offers an *approach* for reducing false positives and accentuating true hostile events, the best practices for a particular heterogeneous network must be developed using iterative experimentation with Nibali's methodology as a guide. Outlining a step by step procedure for developing a Snort configuration file and a set of rules appropriate to minimizing false positives and highlighting serious hostile events in an excessively heterogeneous network could be a source for a separate discussion but is slightly beyond the scope of this paper.

Nibali's paper also describes a method for partially automating the creation of configuration files for multiple versions of Snort running on multiple interfaces of a single platform. This methodology could be used to monitor both the excessively heterogeneous customer access network and the homogeneous private and public service network from one sensor platform. I'm not sure I would rely on a single platform to monitor the potentially large volume of heterogeneous traffic on most service providers' networks. But the option is there for small providers with very tight budgets.

Another way to reduce false positives on heterogeneous segments, at least with Snort, is to disable the portscan preprocessor on incoming traffic. In large heterogeneous environments with hundreds or thousands of users and hundreds or thousands of ip addresses, daily portscans from all around the world are a matter of course. They will occur with startling regularity and there is little that can be done to prevent them. Knowing that they are occurring may reinforce your security administrator's paranoia, which is a good thing, but know that they are there and forget logging them.

Running the portscan preprocessor on outbound packets to the internet will probably generate large quantities of false positives unless connections to frequently hit ports like 53 and 80 are exempted. On the other hand, detecting port scans emanating from inside providers' networks is a useful and noble objective. It may be difficult to accomplish, but every effort should be made to stop hostile traffic from inside out and certified portscans are a sure sign that someone on the internal network is looking for some trouble to get into.

Network Monitoring Tools

Finally, at the upper layer of our defense in depth strategy, there are some pretty standard network monitoring tools that are useful to help detect hostile traffic on excessively heterogeneous networks. These tools are most helpful to quickly isolate denial of service or distributed denial of service attacks. Two tools useful for identifying these attacks are MRTG and the Netflow technology available primarily for cisco routers⁴ and switches. These tools are designed for tracking network utilization, but they are also useful in a security context for quickly displaying evidence of sudden changes in network traffic volume.

⁴ Certain other router and switch manufacturers have begun to include the netflow technology in their products.

MRTG

MRTG is the Multi-Router Traffic Grapher [12]. MRTG is a collection of perl scripts and C programs that communicate with routers, switches and other SNMP capable platforms to monitor traffic volume across network interfaces. It can be run at any interval – the default interval is 5 minutes - to graph traffic volumes across interfaces on a service provider's network. A wily network or security administrator can easily put together a background on his favorite desktop machine using the MRTG graph of his most heavily utilized heterogeneous network interface as seen from the border router. Within 5 minutes of any major traffic change, the administrator should see evidence of that change on his screen. This information can then be used to drill down on that interface and look for trouble. A good tool for drilling down on the trouble is Netflow.

Netflow

Netflow was originally developed by cisco as a set of high performance switching features that captures traffic statistics on routers and switches and exports them to Netflow collectors - host platforms running collection applications [13]. In order to take advantage of Netflow technology, either cisco or other Netflow capable routers or switches must be used. If these products are not available, similar types of information can be captured by other network monitoring tools. A couple of open source tools that will provide similar information are Ntop (<http://www.ntop.org/ntop.html>) and IPTraff (<http://cebu.mozcom.com/riker/iptraf>). I prefer Netflow if you can use it for a couple of reasons: 1) you don't need an extra platform tapping into your heterogeneous network segment; the data comes right off your switch or border router 2) the tools available to do analysis on Netflow data are extremely flexible and the traffic statistics can be visualized and optimized for many purposes.

However, if you do use Netflow, you will need a platform somewhere on your network to collect, analyze and display Netflow results. And you will need software running on that platform to perform the collection, analysis and display tasks. Cisco provides collection and analysis tools at a pretty hefty price. But they are built specifically for the task and come ready to use. Alternatively, there are some good open source tools that do the job nicely. Two open source packages that work together to perform Netflow collection and analysis are cflowd [14] and FlowScan [15]. Cflowd is billed as a complete traffic flow analysis tool that includes the collections, storage, and basic analysis modules for cflowd. But I have had good success using FlowScan to produce “graph images that provide a continuous, near real-time view of the network border traffic” [15].

These images provide a useful breakdown of network flows through the border that will isolate dramatic changes in traffic patterns to a specific protocol, layer 7 service or traffic flow. Further, FlowScan, in its default setup, creates a list of the 10 top traffic users on your network. This can be used to quickly isolate the problem user(s) in the provider's heterogeneous segment in cases of sudden changes in traffic patterns. This data can also be viewed historically to determine, for example, if past alerts generated by the NIDS on

the customer access segment may have been accompanied by associated changes in traffic flows. If they did, this may be evidence of the success of a particular attack.

Summary

Internet Service Providers, large educational facilities, collocation centers and other organizations offering internet service to large groups of customers or users form a special class of service providers. The independent groups of customers or users within these service providers' networks are free to employ any internet protocols in nearly any combination and in any manner. The service providers' networks through which these customers access the internet thus become rich with excessively heterogeneous traffic – traffic whose data volumes and patterns are extremely varied and unpredictable.

Excessively heterogeneous traffic on network segments can complicate the problems already inherent in isolating hostile traffic on those segments. The diversity and volume of excessively heterogeneous traffic make it proportionately more difficult to analyze the contents of that traffic. Specifically, signature based network intrusion detection tools lose much of their effectiveness due to the large number of false positives generated in such environments. Anomaly based NIDS are virtually useless in these circumstances since “normal” patterns only exist on a very large scale that is seldom granular enough to isolate hostile traffic. Using firewalls or access lists on routers connecting these heterogeneous network segments is not really practical since any form of traffic filtering interferes with the concept of “full access”.

Nevertheless, service providers are not exempt from the responsibility of protecting, to the maximum extent possible, their users from the internet and the internet from their users. We have discussed a set of tools and procedures that can work together to encapsulate excessive heterogeneity and to deal with it where it must exist so that service providers can employ an in-depth offensive and defensive stance on these troublesome network segments.

[1] Ferguson, P., RFC 2267, “Defeating Denial of Service Attacks which employ IP Source Address Spoofing”, January, 1998

URL: <http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2267.html>

[2] Huegen, Craig A., “The Latest in Denial of Service Attacks: ‘Smurfing’ Description and Information to Minimize Effects”, Feb 8, 2000

URL: <http://www.pentics.net/denial-of-service/white-papers/smurf.cgi>

[3] McClure, Stuart; Scambray, Joel; Kurtz, George; Hacking Exposed, Third Edition, Osborne/McGraw-Hill, ISBN 0-07-219381-6

[4] Hrivnak, Allison, “Host Based Intrusion Detection: An Overview of Tripwire and Intruder Alert”, January 29, 2002 URL: <http://rr.sans.org/intrusion/HIDS.php>

[5] Northcutt, Stephen, Network Intrusion Detection, New Riders Publishing, 1999

[6] The NSS Group, “Summary – Performance Testing”, “Intrusion Detection Systems”, URL: http://www.nss.co.uk/ids/ids_test_summary.htm

[7] Brennan, Michael P., “Using Snort for a Distributed Intrusion Detection System”, January 9, 2002 URL: http://rr.sans.org/intrusion/distributed_IDS.php

- [8] Boman, Michael, "Building and Maintaining a NIDS Cluster Using FreeBSD and Snort", August 30, 2001 URL: <http://rr.sans.org/intrusion/NIDS2.php> -
- [9] Nibali, Roberto, "Introduction to Network-Based Intrusion Detection Systems Using Snort", June, 2001
URL: <http://www.unixreview.com/documents/s=1234/urm0106j/0106j.htm>
- [10] Costello, Sam, "Server port 80 Plagues Internet Security", April 3, 2002
URL: <http://www.infoworld.com/articles/hn/xml/02/04/03/020403hniss.xml>
- [11] Internet Security Systems, "Internet Risk Impact Summary for December 22, 2001 through March 21, 2002" URL: <https://gtoc.iss.net/documents/summaryreport.pdf>.
- [12] "MRTG - The Multi Router Traffic Grapher", URL: <http://mrtg.hdl.com/mrtg.html> -
- [13] Cisco Systems, "Netflow switching Overview" URL:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/switch_c/xcprt3/xcdnfov.htm
- [14] CAIDA, the Cooperative Association for Internet Data Analysis, "cflowd: Traffic Flow Analysis Tool" URL: <http://www.caida.org/tools/measurement/cflowd>
- [15] CAIDA, the Cooperative Association for Internet Data Analysis, "FlowScan - Network Traffic Flow Visualization and Reporting Tool"
URL: <http://www.caida.org/tools/utilities/flowsan>

© SANS Institute 2000 - 2002, Author retains full rights.

Upcoming Training

Click Here to
{Get CERTIFIED!}



SANS Prague 2017	Prague, Czech Republic	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Boston 2017	Boston, MA	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS New York City 2017	New York City, NY	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Salt Lake City 2017	Salt Lake City, UT	Aug 14, 2017 - Aug 19, 2017	Live Event
Community SANS Omaha SEC401*	Omaha, NE	Aug 14, 2017 - Aug 19, 2017	Community SANS
Virginia Beach 2017 - SEC401: Security Essentials Bootcamp Style	Virginia Beach, VA	Aug 21, 2017 - Aug 26, 2017	vLive
SANS Chicago 2017	Chicago, IL	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Virginia Beach 2017	Virginia Beach, VA	Aug 21, 2017 - Sep 01, 2017	Live Event
SANS Adelaide 2017	Adelaide, Australia	Aug 21, 2017 - Aug 26, 2017	Live Event
Community SANS Pasadena SEC401 @ NASA	Pasadena, CA	Aug 23, 2017 - Aug 30, 2017	Community SANS
Mentor Session - SEC401	Minneapolis, MN	Aug 29, 2017 - Oct 10, 2017	Mentor
SANS San Francisco Fall 2017	San Francisco, CA	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Tampa - Clearwater 2017	Clearwater, FL	Sep 05, 2017 - Sep 10, 2017	Live Event
Mentor Session - SEC401	Edmonton, AB	Sep 06, 2017 - Oct 18, 2017	Mentor
SANS Network Security 2017	Las Vegas, NV	Sep 10, 2017 - Sep 17, 2017	Live Event
Mentor Session - SEC401	Ventura, CA	Sep 11, 2017 - Oct 12, 2017	Mentor
Community SANS Albany SEC401	Albany, NY	Sep 11, 2017 - Sep 16, 2017	Community SANS
Community SANS Columbia SEC401	Columbia, MD	Sep 18, 2017 - Sep 23, 2017	Community SANS
Community SANS Dallas SEC401	Dallas, TX	Sep 18, 2017 - Sep 23, 2017	Community SANS
SANS London September 2017	London, United Kingdom	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Baltimore Fall 2017	Baltimore, MD	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Copenhagen 2017	Copenhagen, Denmark	Sep 25, 2017 - Sep 30, 2017	Live Event
Community SANS Boise SEC401	Boise, ID	Sep 25, 2017 - Sep 30, 2017	Community SANS
Baltimore Fall 2017 - SEC401: Security Essentials Bootcamp Style	Baltimore, MD	Sep 25, 2017 - Sep 30, 2017	vLive
Community SANS New York SEC401	New York, NY	Sep 25, 2017 - Sep 30, 2017	Community SANS
Rocky Mountain Fall 2017	Denver, CO	Sep 25, 2017 - Sep 30, 2017	Live Event
Community SANS Sacramento SEC401	Sacramento, CA	Oct 02, 2017 - Oct 07, 2017	Community SANS
SANS DFIR Prague 2017	Prague, Czech Republic	Oct 02, 2017 - Oct 08, 2017	Live Event
Community SANS Charleston SEC401	Charleston, SC	Oct 02, 2017 - Oct 07, 2017	Community SANS
Mentor Session - SEC401	Arlington, VA	Oct 04, 2017 - Nov 15, 2017	Mentor
SANS Phoenix-Mesa 2017	Mesa, AZ	Oct 09, 2017 - Oct 14, 2017	Live Event