

Implementing Networks Taps with Network Intrusion Detection Systems

Nathan Einwechter 2002-06-19

Implementing Networks Taps with Network Intrusion Detection Systems

by [Nathan Einwechter](#), Senior Research Scientist [Fate Research Labs](#)

last updated June 19, 2002

Editor's Note: *Michel Boucey has kindly translated this article into French, and it is available [here](#).*

Introduction

Over the past decade or so, the use of switches to replace hubs has increased substantially. This is largely due to the increased size of networks, and the requirement for increasingly faster and more efficient networks. On most networks, the data must now be dependable and timely. This transition from hubs to switches, however, has generated a conflict with already deployed and designed network intrusion detection systems.

To combat design conflicts between network intrusion detection systems (NIDS) and switches, network taps were created. Network taps essentially allow all traffic on a network device to be monitored. Network taps are also very useful for passive network troubleshooting and analysis. Further, the tap makes the related NIDS system more secure, preventing attackers from being able to directly attack the NIDS system. This article will offer an introductory overview of taps, including: what taps are, why they should be implemented, their role in improving network security, how they should be implemented, and the economic benefits of taps.

What is the Conflict?

To understand why and how to use network taps, the design conflicts between NIDS and switches must also be understood. Switches differ from hubs in one fundamental way. They differ in how they transmit data from port to port. To demonstrate this, imagine a 4-port hub in which each port has a distinct letter associated with it (A, B, C, D). A computer connected to port A wants to send information to a computer on port C. The packet is sent, the hub receives it, and sends the packet out to all the ports on the hub (A, B, C, D). This is illustrated in figure

1.




Figure 1

In this situation, the NIDS systems have no problem. Since all traffic is sent to every port, a NIDS can detect traffic no matter where it is being sent across the hub.

Switches, however, send data in a completely different way. Instead of sending data destined for port C to every port on the device, the switch sends this data only to port C. This increases efficiency by reducing packet collision, and optimizes bandwidth by reducing unnecessary transmissions. This can be seen in Figure 2.

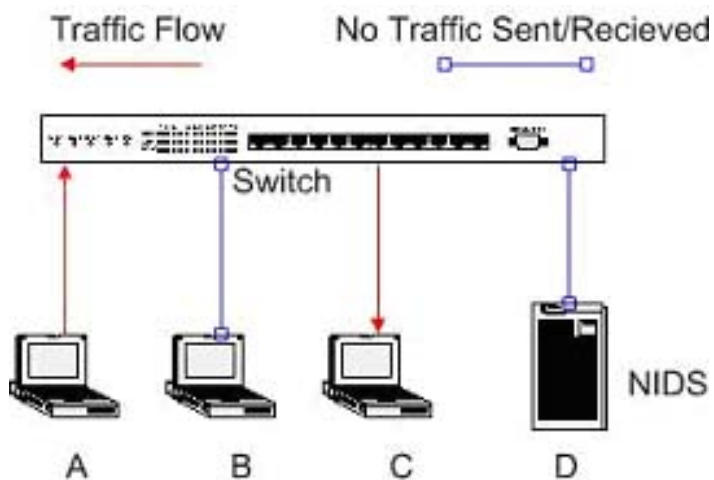


Figure 2

This diagram clearly demonstrates where the problem occurs. Absolutely no data is sent to the NIDS system (Port D), thus no event detection can take place. The only time the NIDS could detect an attack is if the attack were to be directed to the NIDS itself. This is, obviously,

unacceptable as it completely defeats the purpose of having a NIDS on the network in the first place.

What is a Network Tap?

Tap stands for "Test Access Port". Network taps allow all traffic on a network device (such as a switch) to be passively monitored. They are relatively inexpensive, reliable, and provide permanent access ports to monitor traffic through. Taps are usually separate devices, but can also be built into a switch itself. Two common tap solutions are offered by [NetOptics](#) and [Finisar](#). There are taps for just about any type of network in use today. This includes GigaBit SX, LX or ZX, ATM, DS3, T1, Fast Ethernet copper, and GigaBit TX to SX. This means that our NIDS can be deployed using a tap on basically any type of network setup imaginable. Further, taps are completely passive which means they will not disrupt your current network configuration, and are easily implemented within any existing network set-up.

Why Implement Network Taps?

Network taps are an ideal way to implement IDS into a switched and high-speed environment. To understand why taps should be used in these situations, it may be helpful to look at some other option for implementing an IDS into a switched environment. The most commonly used alternative is port spanning, also known as port mirroring. This option, although used often, has inherent flaws that create problems in implementing IDS systems with it. Most switches in use today come with this type of port. Port spanning or mirroring forces the switch to either send all packets from across the switch, or packets from a specified port, to a specific span/monitoring port in addition to delivering it to its intended recipient. This raises a few issues, the most obvious of which is that of packet loss to the mirror/span port. When utilizing such a port, it is much like a port on a hub. This means that there will be a higher rate of packet collisions, as the twenty other ports on the switch continually send packets to the mirror/span port. Furthermore, since this mirror/span port usually has the same amount of bandwidth as every other single port has, the packet loss is significantly increased again.

Port mirroring also presents problems in that it does not receive error packets or VLAN information, and only presents one side of a full-duplex connection. Thus, the IDS sitting on a mirror is severely limited by increased packet loss, and a complete blindness to half of the traffic on a link. What good is an IDS that can only see half the traffic at best, and less than that under most circumstances? The workings of port mirroring are demonstrated in figure 3.

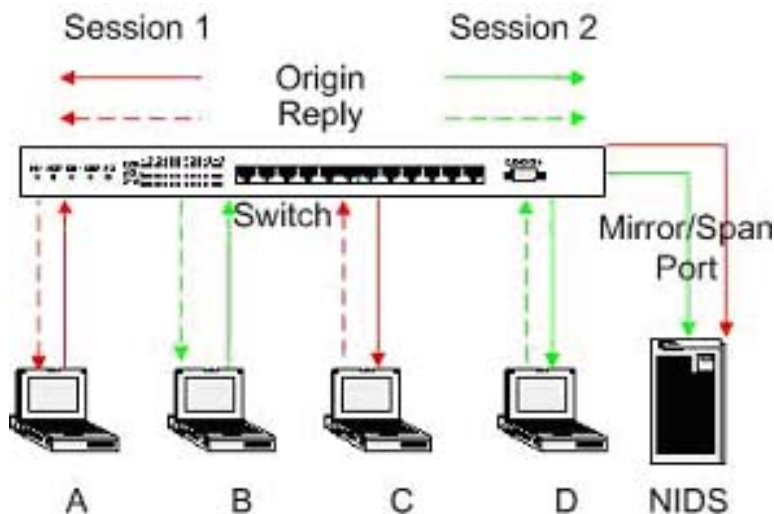


Figure 3

Note that half of the session is not being sent to the Mirror/Span Port; thus, as was described previously, the NIDS is not able to detect events within these messages.

The use of a tap fixes all of these problems. The tap is able to give IDSs the ability to view both sides of a full duplex conversation, reduce packet loss due to network hardware completely, and view all packets that are transmitted across the line. It is able to accomplish all of this in a passive mode that does not affect the network structure as a whole.

The Role of Taps in Increasing IDS Security

As mentioned in the introduction, taps can actually increase the security of an intrusion detection system installation. The reason for this is quite simple: IDSs behind taps do not require an address because the tap takes any and all data off the line and throws it directly to the IDS interface, thus eliminating the need for addressing. The IDS has no address; therefore, no traffic can be directed specifically towards the IDS. This prevents directed attacks against the IDS system, and can actually make attackers believe that no IDS is present to identify and track their attacks.

By preventing the detection of the IDS by attackers, the survivability of the system is significantly increased. After all, what good is an IDS if it can be attacked and disabled?

Explanation of Implementation

The installation and implementation of taps can be quite easy if you plan the set-up ahead of time, and if you have the right hardware. Most of the time, the tap will come off an expansion port that is commonly found on many newer switches, hubs, and routers. Often this is done through the utilization of data terminal equipment (DTE) and/or data communication equipment (DCE) interface(s). These interfaces will be discussed in greater depth later on. This allows the tap to be completely passive to the network, thus reducing the chances of the network disruption that many monitoring devices may cause. These lines also allow the high speeds required to permit tapping an entire device while keeping packet loss to an absolute minimum. The DTE/DCE interfaces are then fed directly to the fiber tapping panel that is stored on a server rack, which contains the tap monitoring port. Attached to the monitoring port is the NIDS we have decided to use. This is demonstrated in figure 4.

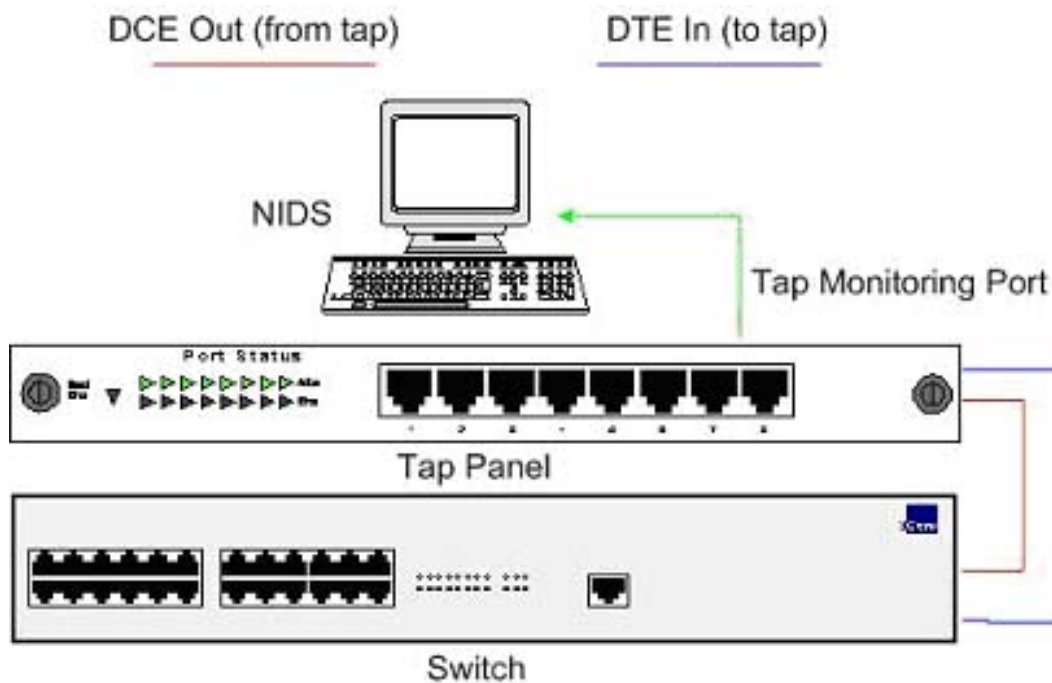


Figure 4

This is the most common type of set-up. Most network taps can be deployed on multiple devices across the network and then be interfaced together and with other systems at the server rack (as illustrated in figure 5).




Figure 5

The monitoring racks (tap panels) are usually designed for high speed (gigabit and better) systems to allow for all of this traffic to be collected at this central point with the minimum possible amount of packet loss. Due to the extremely high speeds these monitoring racks can stand, it is often necessary to use multiple IDS systems and load balance between them, as very few IDSs are currently able to handle these types of speeds. If multiple IDSs are used, then one may consider the idea of aggregating this data together on an analysis machine for further viewing and event detection at a later date. Jeff Nathan's [Gigabit Tap IDS illustration](#) gives a good demonstration of how this type of system would be set up and operated.

Data Terminal Equipment/Data Communication Equipment

Data Terminal Equipment (DTE) and Data Communication Equipment (DCE) interfaces, as mentioned earlier, are the most common means of transporting the data from a switch, router, or hub and into the tap panel. DTE is the actual device at the user end of a user-network interface that serves as a data source, destination or both. DTE connects to a data network (in our case to the tap) through a DCE device. DCE devices, on the other hand, provide a physical connection to the network, forward traffic, and provide a clocking signal for use with synchronous data transmission between DCE and DTE devices.

Now that the basic premise of DTE and DCE has been established, their role in providing data transfer for taps can be discussed. There are many different types and classes of DTE and DCE cables, interfaces, and connectors. There are, however, a few devices that are more commonly used with taps. The first of these is the High-Speed Serial Interface (HSSI). The HSSI is capable of handling T1 and OC-1 speeds, allowing high-speed connectivity between networks and devices utilizing the DTE/DCE interfaces. HSSI is able to achieve these speeds through utilization of differential emitter-coupled logic (ECL), which has been used in the Cray computer system interfaces for years. Other technologies such as frame relay, ADSL, ATM, and DS3 can be utilized through DTE/DCE interfaces for taps accordingly, as required by the implementing network.

The specifics of how to set up the DTE/DCE interface and the types of cabling used in the implementation will vary widely depending on the devices used throughout the network. Most tap systems and switches that support DTE/DCE devices will also have specific documentation on setting those devices up to work with each other. In most cases, the "analysis" port described in the documentation is the port you will be connecting your NIDS to.

Jeff Nathan has created a set of very good diagrams that demonstrate the generics of a network tap. These diagrams can be found at the [Snort](#) documentation page, and serve as good references.

Cost Benefits of Network Taps

Beyond the previously discussed technical advantages of utilizing network taps, there are also financial benefits. If a NIDS was in use before the tap was implemented, no change in IDS technology would be required to implement it into the tap. Since no new NIDS systems and deployments would be required, implementation costs would be significantly reduced. Further, since the NIDS would still be the same, the security and networking staff would not need to be retrained on a new system.

Utilizing network taps to allow the use of NIDS also creates cost benefits when one looks at some of the alternatives. The best example of this, is the transition away from NIDS in switched environments, over to network node (NNIDS) or host-based intrusion detection systems (HIDS). These systems would require that the IDS be installed on every system one wished to detect attacks against. To maintain the same amount of attack detection coverage across the

network as a NIDS can provide, the corporation installing the systems would have to pay for a copy of the NNIDS/HIDS for every system within a protected network segment. This, obviously, increases implementation costs significantly and requires staff be retrained on the various NNIDS/HIDS to be used. This cost is further magnified when multiple operating systems are in use, that each require different types of NNIDS/HIDS. In this situation, it is readily apparent that implementing network taps to allow NIDS to operate is significantly more cost effective than the alternative.

Conclusion

The current evolution from hubs to switches, coupled with increased security concerns and risks, makes the deployment of traditional NIDS systems problematic on most current networks. These problems can easily be overcome by utilizing network taps to allow the use of NIDS to maintain sufficient levels of security in just about any type of network currently in use. The use of network taps actually increases the security of the NIDS sensors while simultaneously allowing them to be more easily deployed on large high-speed networks that require switches and other high speed network devices. Taps also allow a reduction in IDS implementation costs

Taps are not currently in use as much as many security professionals believe they should be. There is, however, a general consensus that the use of taps will increase as the need for them, and the education about them, increases among corporations and other owners/operators of large networks.

[Nathan Einwechter](#) is currently a core member of BCN Group, developing a proposed national cyber defense system. He also worked previously as a System Developer/Incident Analyst with [myNetWatchman](#), and as a Senior Research Scientist/Head of Research and Development for [Fate Research Labs](#). Nathan is scheduled to speak at the Canadian Security and Intelligence Conference this summer..

[Privacy Statement](#)

Copyright 2006, SecurityFocus