INTERNET
SECURITY
SYSTEMS

HOW TO GUIDE

Internet Security Systems

Intrusion Detection Systems

INTERNET SECURITY SYSTEMS

How To Guide-Implementing a Network Based Intrusion Detection System

Written by Brian Laing, (brian@laing.org)
Contributors
Jimmy Alderson (blue0ne@igloo.org)

Sovereign House
57/59 Vaster Road
Reading RG1 8BT

Table of Contents

Chapter 1

Introduction

"Threats against corporate data are on the rise, and more companies are suffering financial losses because of attacks on computer systems. A greater amount of information is vulnerable to theft…" (Information Week, The Security Facade, Bob Violino, Oct 21, 1996, page 36)

What is an Intrusion Detection System (IDS). Intrusion Detection Systems look for attack signatures, which are specific patterns that usually indicate malicious or suspicious intent. When the IDS looks for these patterns in network traffic via a promiscuous interface it is considered a Network Based IDS. There are three forms of a Host based IDS. Of the two main ones, the first examines the logs of the host looking for attack patterns; the second examines patterns in the network traffic (this is not done in promiscuous mode like the Network IDS). The third one is a solution that executes both Log based and Stack-Based IDS.

Network-Based IDS

Network-Based Intrusion Detection Systems (IDS) use raw network packets as the data source. The IDS typically uses a network adapter in promiscuous mode that listens and analyses all traffic in real-time as it travels across the network. A first level filter is usually applied to determine which traffic will be discarded or passed on to an attack recognition module. This first level filter helps performance and accuracy by allowing known un-malicious traffic to be filtered out. An example of this would be if an event for suspicious SNMP get was detected, and a known SNMP management station generated this event. Using Filters SNMP traffic from this machine could be filtered out of the examined traffic. Caution must be taken when using filters as traffic can be spoofed, and mis-configurations can cause more traffic to be filtered than desired. At the attack recognition module, typically on of three methodologies are used for attack signatures; pattern, frequency, or anomaly based detection. Once an attack is detected a response module provides a variety of options to notify, alert, and take action in regards to the attack at hand.

Host Based IDS

Host Based Intrusion Detection actually started in the early 1980's before Networks were as prevalent, complex and inter-connected as they are today. In the 1980's it was common practice to review audit logs for suspicious and security relevant activity. Today's host-based IDS still use various audit logs but they are much more automated, sophisticated, and real-time with their detection and responses. Host-based systems use software that continuously monitor system specific logs. On Windows NT these include system, event, and security logs, while on most flavours Unix they include Syslog and OS specific log files. As soon as there is a change to any of these files the host-based IDS compares the info with what is configured in the current security policy and then responds to the change accordingly. One method of host-based IDS is to monitor log activity in real-time, while other solutions run processes that check the logs periodically for new in formation and changes. Being that the IDS is monitoring these logs continuously or frequently the detections and responses are considered to be in near real-time. Some host-based IDS can also listen to port activity and alert when specific ports are accessed, this allows for some network type attack detection.

Stack-Based IDS

This is the newest IDS technology and varies dramatically from vendor to vendor. Stack-Based IDS works by integrating closely with the TCP/IP stack, allowing packets to be watch as they traverse their way up the OSI Layers. Watching the packets in this way allows the IDS to pull the packets from the stack before the OS or the Application have a chance to process the packets. To be complete Stack-Based ID should watch both incoming and outgoing network traffic on a system. By monitoring network packets destined only for a simple host, the principle is to make the IDS have sufficiently low overhead so that every system on the network can run Stack-Based IDS.

Strengths of Networked Based

The network-based IDS has many strengths due to its real-time packet capture and analysis functionality that cannot easily performed with a host-based IDS alone. Below are some its strengths that make network-based IDS clearly a needed component to any security systems and policy.

1. **Cost of Ownership**. The network-based IDS allows strategic deployment at critical access points to view network traffic destined to numerous systems that need to be protected. It does not require software to be loaded and managed on a variety of hosts as does the host-based IDS. Being fewer detection points are required, it makes the cost of management more effective for an enterprise environment.
2. **Packet Analysis**. The network-based IDS examines all packet headers for signs of malicious and suspicious activity. Many of today's IP based denial of service (DOS) attacks are detected by looking at the packet headers as they travel across a network. For example; a LAND attack is a forged packet that has both source and destination IP addresses and source and destination ports being the same as the target host machine. The forged packet which tells the target to open a connection with itself and can cause target to operate slowly or crash. This type of attack can quickly be identified and responded to by a network-based IDS being it is looking packet stream in real-time. Also, attacks that use fragmented packets like TearDrop can also be detected with packet level analysis. A host-based IDS cannot detect these types of attacks. In addition to

looking at the packet headers, a network-based IDS can also investigate the content of the payload looking for specific commands or syntax used with a variety of attacks. Many of these commands are indicative of an attack, whether successful or not. Example: An attacker probing for the new Back Orifice exploit on systems that are not infected with the Back Orifice software can be detected by examining the packet payload. A host-based IDS would not be able to detect these types of payload embedded attacks.

3. **Evidence Removal**. The network-based IDS uses live network traffic for its attack detection in real-time and a hacker cannot remove this evidence once captured. This captured data not only has the attack in it but information that may help lead to his/her identification. This captured network traffic may also be needed as evidence leading to prosecution. One problem sometimes experienced with host-based IDS is that hackers understand most of the audit log files and that is usually the first place they go to cover their tracks and remove or damage this information.

4. **Real-Time Detection and Response**. The network-based IDS detects malicious and suspicious attacks as they are occurring in true real-time and provides faster response and notification to the attack at hand. Example: A hacker initiating a network based denial of service (DOS) based on TCP can be instantly stopped by having the IDS send a TCP reset to terminate the attack before it crashes or damages a targeted host. Many times with host-based IDS, the notification from a particular log or event may be detected after the damage is already done or not at all if the systems has crashed. Having real-time notification allows you to quickly react in a desired fashion. Some may want to allow further penetration so they can gather more information in a surveillance mode while other may opt for immediate termination of the attack.

5. **Malicious Intent Detection.** A network-based IDS can also be very valuable in determining malicious intent. If a network-based IDS is placed outside of detection Firewall it can detect attacks intended for resources behind the Firewall, although the firewall may be rejecting these attack attempts. A host-based IDS could not show these rejected attacks because they never hit the Host but are important to know the frequency and types of attacks being thrown at your network.

6. **Complement and Verification**. The network-based IDS can also complement existing components of your implemented security policy. In the case of encryption, the network-based IDS although it may not be able to read all encrypted traffic, it can be used to detect any not encrypted traffic that may be present of on your network. In the case of a Firewall, the network-based IDS can help verify if it is truly keeping out certain types of traffic and addresses that it should be rejecting.

7. **Operating System Independence**. The network-based IDS is not dependent on the host operating system for its detection source, as is the host-based IDS solution. All of the log information used with a host-

based solution requires the operating system functioning properly and not compromised in any way.

The Strengths Host-Based

Being the Host Based IDS resides on specific Hosts and uses the information provided by the operating system, it adds capabilities not found in the Network Based IDS. Some of the major strengths include:

1. **Attack Verification** – Being the Host Based IDS uses logs containing events that have actually occurred, it has the advantage of knowing if the actual attack or exploit was successful. This type of detection has been deemed as more accurate and less prone to false positives. Many Network Based attacks can trigger numerous false positives because of normal traffic looking very close to malicious traffic. In addition it is hard for a Network Based IDS to know whether or not an attack was successful or not.
2. **System Specific Activity** – The Host Based IDS can quickly monitor user and file access activity. Anytime a Login or Logoff procedure is executed it is logged and the host-based IDS can monitor this based on its current policy. In addition it can also monitor various file access and also be notified when specific files are open or closed. This type of system activity cannot be monitored or detected by Network Based IDS being it may not necessarily propagate traffic on the network. Example: Someone walking up to a Keyboard and open a non-shared file. The host-based IDS can also monitor activities that should and can only be executed from an administrator. Anytime user accounts are added, deleted, or modified this information is logged and can be detected as soon as the change is executed. Also if any audit policy changes are made affecting what the systems logs and does not log the Host Based Ids will immediately pick up this activity. A network-based IDS also could not detect these types of administrator changes.
3. **Encrypted and Switched Environments.** Being the host-based IDS software resides on various hosts throughout an enterprise it can overcome some of the challenges faced by network-based IDS. In a purely switched environment it can be challenging to deploy a network-based IDS due to the nature of switched networks having numerous separate collision domains or segments. It is sometimes hard to get the required coverage being the network-based IDS can only reside on segment at a time and numerous ones need to be protected. Traffic mirroring and Span ports on switches have helped but still present challenges getting enough coverage and have limitations. The host-based IDS can provide greater visibility into a purely switched environment by residing on as many critical hosts as needed. As for Encryption, there are certain types of encryption that can also present a challenge to the network-based IDS depending where the encryption resides within the protocol stack it may leave the network IDS blind to certain attacks. The host-based IDS do not have this problem. If the host in question has

log-based analysis the encryption will have no impact on what goes in to the log files. Stack-Based IDS allows us to monitor the packets as they traverse the TCP/IP stack, this allows the IDS to examine the packets before the OS or the Application see the packets, but after the TCP/IP stack has decrypted the packets.

4. **Monitoring Key Components.** The host-based IDS gives you the ability to monitor important system components such as key executables, specific DLL's and the NT Registry. All of these files could be used to breach security and resulting in systems and network damage. The host-based IDS can alert you when these files are executed or modified. Also items like disk space usage can be monitored and alerted on at certain levels. This can be very helpful in detecting if a hacker is using your server hard drive as a storage facility. These types of internal files cannot be detected with a network-based IDS.

5. **Near Real-Time Detection and Response.** Although a host-based IDS relying on Log Analysis is not true real-time, if implemented correctly it can be extremely close. "Near Real-Time". Instead of having a process check the status and content of log files are defined intervals, some of today's host-based IDS can detect and respond as soon as the log is written to and compared to the active attack signatures. This near real-time detection can help catch the hacker before he/she has time to do extensive damage and remove evidence of being on the system.

6. **Real-Time Detection and Response**. Stack-Based IDS can examine inbound and outbound packets and determine in real-time if an attack is being executed. If the Stack-Based IDS detections an attack it can then respond in real-time as well.

7. **No Additional Hardware** The host-based IDS does not require additional hardware to do intrusion detection. It easily resides on your existing network resources (File Servers, Web Servers, and other shared or critical resources). This can make the host-based IDS more cost effective in some cases. In addition: "its not another box on the network" that requires addressing, maintenance, and management.

8. **Firecell ©.** Firecell is a technology that is part of ISS Stack-Based IDS Server Sensor. Firecell allows either by pre-configuration, or as a response to an attack, to have specific traffic refused. This is done similar to a firewall in that the Stack-Based IDS examines the packets and if they match a Firecell rule the packets are dropped. For example if your company policy is for no HTTP servers to be installed on workstations, a Firecell rule could be implemented that dropped all inbound packets to port 80.

The need for both types

As you can clearly see both network and host-based IDS solutions have unique strengths and benefits over one another and that is why the next generation IDS must evolve to include a tightly integrated host and network component. There are no "Silver Bullets" when it comes to network security but adding these two required components will greatly enhance your resistance to attack. Below is a quick graphical representation that helps represent the independent network and host-based IDS scenarios and the benefit of implementing both a Network and a Host Based Solution

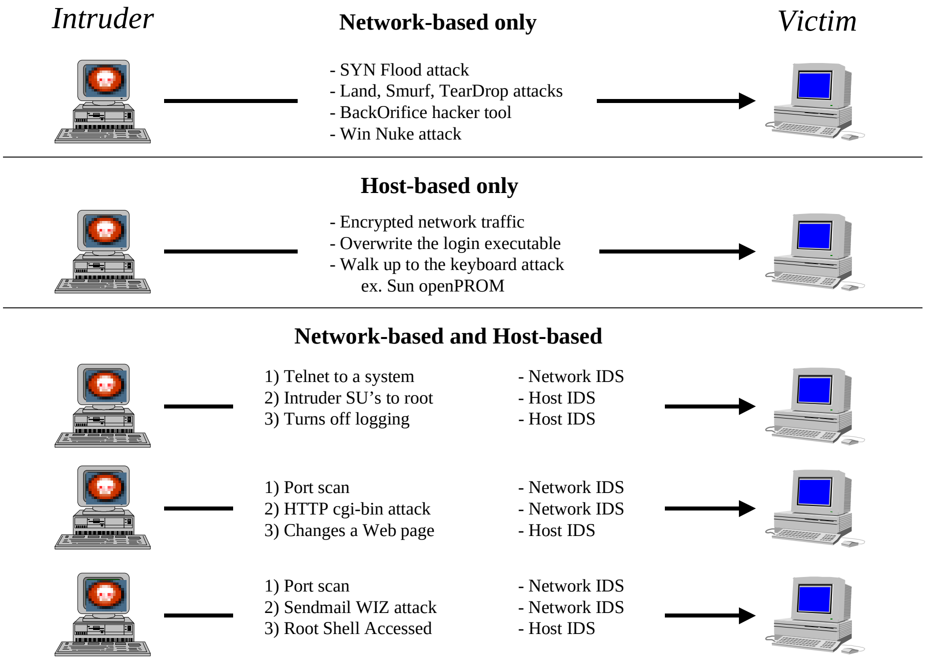

Implementing an IDS

An implementation of an IDS has all the normal hurdles for implementing a piece of Security Software. These include testing impact to various systems, managing the installed software and the install itself. Implementing a Host Based IDS is fairly standard having similar problems to implementing any other Host Based piece of software. Implementing a Network Based solution has these problems as well as the problem of tapping into the communication flow.

Many networks are now being built or upgraded to a switched Infrastructure instead of a Shared Media Infrastructure. This makes tapping into the communication flow an exercise in engineering that must be fully planned out. This document will outline the problems and pros & cons of the various solutions, as well as demonstrating them being used in an E-Commerce and a Perimeter solution.

Switched VS Shared Media

Many of today's Ethernet Networks are shared media, i.e. they share the total available bandwidth among many users. A 12-port 10mpbs hub will deliver 10mbps/12 users = .83mbps average throughput per users (this assumes each user generates the same amount of traffic constantly). Today's traffic is neither constant nor consistent, increasing the already present number of *Collisions* degrading the performance below the .83mbps, which is nowhere near the 10mbps speed the wire is capable of.

Collision: When packets are transmitted by one port on the Hub, the hub echoes these packets to all the other ports on the Hub. If two machines broadcast packets at the same time then their packets will collide, causing both machines to retransmit their packets.

A switched environment has a dedicated amount of bandwidth for each user. So a 12-port 10mbps switch will supply the full 10mbps throughput regardless of the network traffic. Since the media is no longer shared Collisions do not occur, and traffic can flow near wire speed.

If you think of the network as a small meeting room, and the network used Shared Media then everyone hears everyone else's conversations i.e. if someone were to say the room is on fire then everyone in the room would hear this. While a meeting room using switched media is like all the participants in the meeting using a cellular phone to communicate. This makes it very difficult to hear the other conversations going on in the room without the implementation of further technology.

Monitoring a Switched Connection

There are three main ways to tap into a switched connection, each one has its advantages and disadvantages . The three main methods are TAPS, Hubs and Spanning ports. The paragraphs below will outline how to monitor the traffic between the router and the switch show in Figure-1, as well as issues in managing the IDS once it is in place.

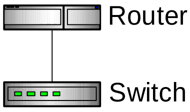

Figure-1

Span Port

The Switch Port Analyzer (SPAN) port is typically put to use by a Network Sniffer to monitor network traffic. The port works by configuring the switch to copy TX/RX/Both from one port or VLAN to another port. For instance in Figure-2 the switch is set to span both TX and RX from the port the router is attached to, to the port the IDS is installed on. This allows the IDS to monitor any traffic that passes between these two devices. Other than the added traffic passed to the SPAN port, the port is a standard port, which means managing the IDS can be done by any machine that can route IP packets to the IDS.

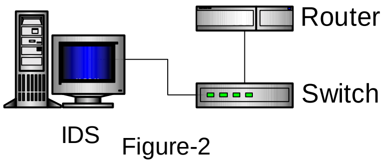

Figure-2

The advantages of using this solution are:

- Ease of installation: the IDS can be placed on the switch without modifying the core infrastructure.

- Managing the IDS requires no additional hardware or special configuration changes.

- Terminate Sessions and Firewall reconfigurations are not impacted.

The disadvantages of using this solution are.

- You can only have one span port per switch. To monitor more than one port you must span a range of ports, i.e. span ports 1-5 to port 6, or span an entire VLAN.

- Spanning more than one port is not feasible as it can very quickly overload the span port. This is especially true when monitoring a full-duplex environment.

- Without additional changes to the IDS the IDS is vulnerable to attack. This can be avoided by implementing the IDS in Stealth Mode.

- Without using a Stealth Configuration the Alerts generated by the IDS can cause additional problems with overloading the port.

Stealth Config; This is defined as a process of using two network interface cards in the IDS. The monitoring card has all Network Protocols unbound from the card. This prevents the exposed interface from being attacked. The second interface is then routed to a management LAN via a secure network. See Appendix A for a diagram.

- Switch performance degradation.

- Some span ports are uni-direction so terminating sessions is not supported.

- Inability to mirror errors such as undersize and oversize packets, and packets with a bad CRC.

Hub

This configuration is not recommended. The chance of causing a major network problem is too high. Using Hubs or TAPS is a very similar solution; the hub or tap is placed between the connections to be monitored. This is usually between two switches, a router and switch, or a server and switch, etc. In Figure-3 a hub has been placed between the router and the switch. This allows traffic to still flow between the router and the switch, while the properties of the hub cause a copy of the traffic to be copied off to the IDS.

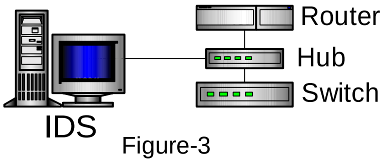

Figure-3

The Advantages of using the this solution:

- Easy to configure: doesn’t require any intrinsic knowledge to implement.
- Managing the IDS requires no additional hardware or special configuration changes.
- Terminate Sessions and Firewall reconfigurations are not impacted.
- 4 port hubs are very economical

The Cons of using this solution

- Due to limitation of shared media, this cannot be used if the connection between the switch and router is a full-duplex connection, as collisions will degrade the throughput.
- Due to the limitation of shared media, if the management of the IDS is done through the hub it will increase the number of collisions impacting the flow of traffic between the Router and Switch.
- Low cost hubs tend to be prone to failure

Taps

As stated before the Tap solution is very similar to the hub solution. Taps are by design fault tolerant having the main connection (i.e. the connection between the resource and the switch), hardwired into the device, preventing failure. For more detailed information on the Tap Please see Appendix 1. Once the Tap is in place there are several ways to route the traffic it collects.

Taps come in several configurations with the only difference between the configurations being how many connections the tap can monitor. Currently on the market are taps for a single port, 4 port, 8 port and 12 port. All of these but the single port tap can be purchased as a rack mountable unit.

Generic Pros of using a Tap

- The Tap is fault tolerant if power fails the connection between the router and the Switch is hardwired in and requires no power to function.
- Tap does not impact the traffic flow
- Once tap is in place changes to the IDS infrastructure do not impact the overall network
- The Tap prohibits people from establishing a direct connection to the IDS protecting it from some attacks.
- Taps eliminate the need for network connections to be broken and re-cabled each time a network segment needs to be analysed. Allowing a tapped/ TopLayer solution to scale quite large over time with no impact to the network as changes are made.
- Century Taps plus IDS are much more cost-effective in networks with multiple full-duplex links.
- No degradation to the network.
- Allows the IDS to monitor errors such as undersize and oversize packets, and packets with a bad CRC.

Generic Cons of using the Tap

- Taps can be costly
- Without extra modifications Terminate Sessions is not supported
- Without extra modifications the solution cannot monitor traffic in both directions.
- Requires Stealth Configuration.

There are various methods for using these three technologies to monitor network traffic. The first method demonstrates a method that can easily be scaled to growing environments. The Taps are consolidated into a TopLayer Application Switch (See Appendix C for more information). There are several items, which makes this switch different from other switches.

- The ability to have more than one span port. For example you can span ports 1-5 to port 6 and ports 7-11 to port 12.

- The output from the taps can be load balanced across multiple span ports. For example ports 1-6 could be spanned and load balanced to ports 7-9.

Using these abilities the output from taps can be consolidated into one TopLayer Switch without the worry that the port will be overloaded prohibiting the IDS from gaining access to the packets it needs to monitor. Executing this solution is best done using a rack mountable multi port tap. On busy connections a 4 port rack mountable tap in conjunction with a single TopLayer switch, and 4 IDS sensors work very well. The number of taps can be increased or decreased based on the utilization of the network.

As shown in Figure-8 the out put from the Taps has been routed into the TopLayer Switch. Two IDS are then placed on the Switch to handle any traffic that may overload a single port.

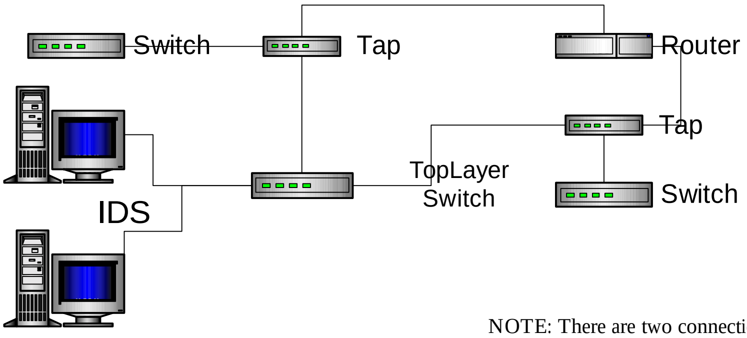

Figure-8

NOTE: There are two connections between each tap and the TopLayer Switch. One connection has been removed from the diagram to make it more readable.

Additional Pros of this configuration.

- IDS is not taken out when the traffic reaches volumes that would normally overload a single port.

- Tapped connections can be consolidated into a fewer number of IDS's lowering the Total Cost of Ownership of the Security Infrastructure.

- IDS is protected from connection based attack by the Taps, reducing the need for a stealth configuration. Management of the IDS can be done via the TopLayer Switch.

- Very Scailable; additional IDS can be added to the TopLayer switch without downing the network.

Additional Cons of this configuration.

- Requires minimal expertise with a new piece of equipment.

The TopLayer switch also has a configuration that supports gigabit Ethernet. This can be used to break the traffic apart into manageable 100mbps chunks. This is shown in the diagram below. The TopLayer Switch functions as the tap in this solution. Depending on the Utilization of the Gigabyte link this solution would need between 1-10 Network Sensors.

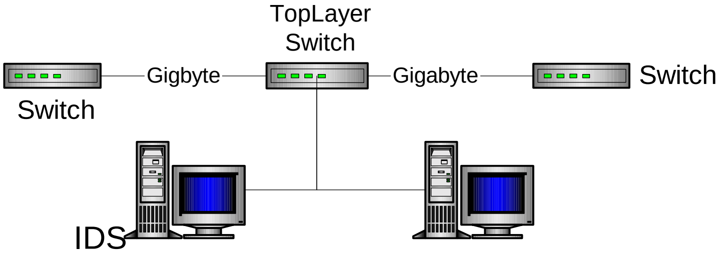

The second solutions is very similar to the previous one in that it uses switched media to monitor the packets collected by the taps. The switch used in this configuration however, is a standard switch, such as a Catalyst 5000. The switch is configured so that all of the incoming ports are part of a single VLAN. This VLAN is than spanned or mirrored to the IDS.

Additional Pros of this configuration.

- In a full-duplex environment the Switch will do its best to buffer any packets that would normally overload the port.

Additional Cons of this configuration.

- The switch has a limited buffer. So you should take into consideration the aggregate bandwidth of the networks being tapped to decide how many taps to place on the same switch. If the aggregate bandwidth of eight networks to be monitored is less than 100MB of traffic, then you will need 16 ports to accommodate the output from the taps. However, if the aggregate bandwidth of eight networks to be monitored is greater than 100MB than you will require more than one switch and require fewer ports per switch for the output from the taps.

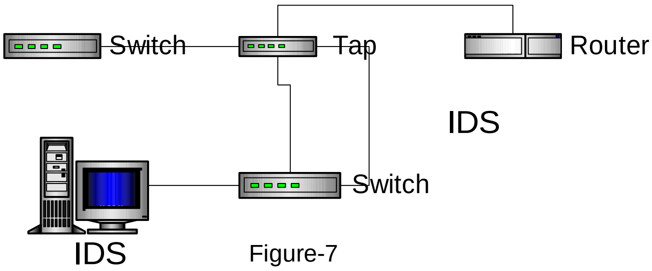

Figure-7

While the next three configurations are possible they are not recommended they are included purely for completeness. Implementing any of these solutions can cause attacks to be missed, thus reducing the strength of your IDS.

Tap the connection and only monitor the traffic in one direction. The IDS will also not detect attacks going in the opposite direction.

There are no additional Pros to this configuration, the only additional Con is.

- Since the IDS is seeing packets from a single direction the IDS cannot detect the following three attacks unanswered arp's, Ipduplicate, and Synflood, because they require the packets going in both directions.

- Since the IDS is seeing packets in single direction only attacks inbound or outbound are monitored.

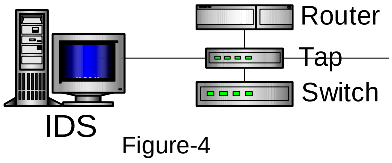

Figure-4

Figure-5 demonstrates one way of monitoring both ends of the connection; in this solution an IDS has been placed on both outputs of the Tap. This allows attacks in both an inbound and an outbound direction. This solution still will not detect unanswered arp's, Ipduplicate, and Synflood.

There are no additional Pros to this configuration, the only additional Con is.

- Extra hardware and software required.

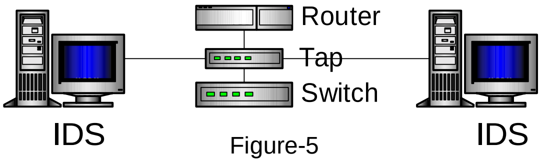

Figure-5

Figure-6 demonstrates a configuration that takes both outputs of the Tap and routes them into shared media for the IDS engine to monitor.

Additional Pros for this configuration are.

- Can monitor both directions from single IDS, allowing the detection of unanswered arp's, Ipduplicate, and Synflood.

Additional Cons for this configuration are.

- Due to limitation of shared media, this cannot be used if the connection between the switch and router is a full-duplex connection. This is caused from the Tap passing the packets coming in from both directions to the hub simultaneously, causing a collision. Normally packets colliding like this will get retransmitted, but since the Tap is a uni-direction flow the transmitting devices never see the collision and do not retransmit the packets.

- Using a hub to monitor more than one tap will also have an increased problem with collisions.

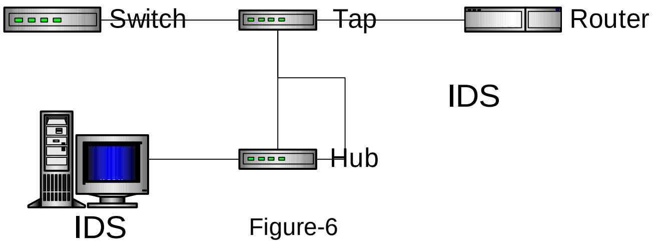

Figure-6

Terminate Sessions

Since Taps only allow traffic to flow in one direction the IDS cannot terminate sessions. To allow sessions to be terminated we must introduce an extra tap into the configuration.

Figure-9 shows an extra tap being routed back on to the primary network, this allows the session to be terminated while the tap prohibits a connection from being established.

The IDS monitors traffic by way of Tap-1. This traffic is consolidated into the Toplayer switch. If the IDS needs to generate a kill the kill would flow back along this path, however when it reaches Tap-1 the packets will be dropped.

To allow Kills Tap-2 is inserted with the TX from the IDS to the TopLayer Switch connected into the primary network (typically this is done by way of one of the switches being monitored). Now when a kill is issued it will flow the same path as stated before with the packets also being copied to the switch which will then pass the packets on to the appropriate location.

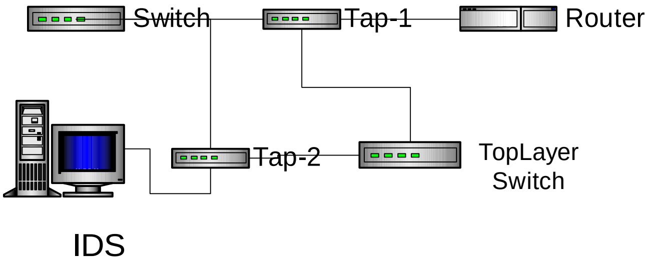

Figure-9

Note!! Careful attention needs to be taken to assure the correct cable from the Tap is feed back into the network. If the wrong cable is feed back into the network then network traffic could be duplicated and a packet storm created.

Example Implementations

The configurations shown above can be put to use in a variety of situations, however most implementations are done on either an Internet gateway or an E-Commerce site. While both of these configurations are very similar they each will done done slightly differently. This section will outline an implementation for each of these two sites as well as any additional configurations needed to make this solution work.

Internet Gateway

This configuration will outline a typical implementation on a Corporate Internet Gateway. All outbound traffic (HTTP, FTP, etc.) by internal employees will use this Gateway. Additionally the Gateway supports the companies Internet presence. Figure-10 outlines the Gateway Topology. This gateway does not require high availability so it is limited to a single ISP connection and a single Firewall.

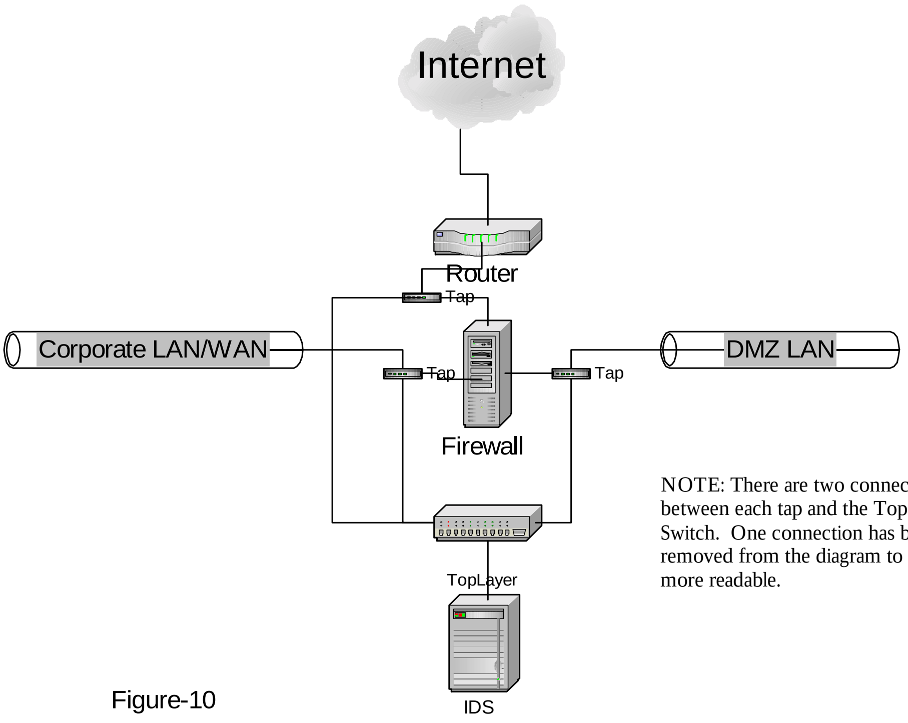

NOTE: There are two connections between each tap and the TopLayer Switch. One connection has been removed from the diagram to make it more readable.

Figure-10

In this configuration the taps do not generate enough traffic to flood a single IDS, so only the consolidation features of the TopLayer are being used. If traffic increases to a level where the port the IDS is connected to becomes overloaded a second third etc IDS could be added as needed without taking down the network infrastructure to put it in place, this is shown in Figure-11.

Once the additional IDS have been put in place, one IDS can be used specifically for the traffic on the outside of the firewall, while the second is used to monitor the DMZ and the Corporate LAN. Once this is done the IDS monitoring the DMZ and Corporate LAN can mimic the firewall rule-set to determine when it has been compromised or mis-configured.

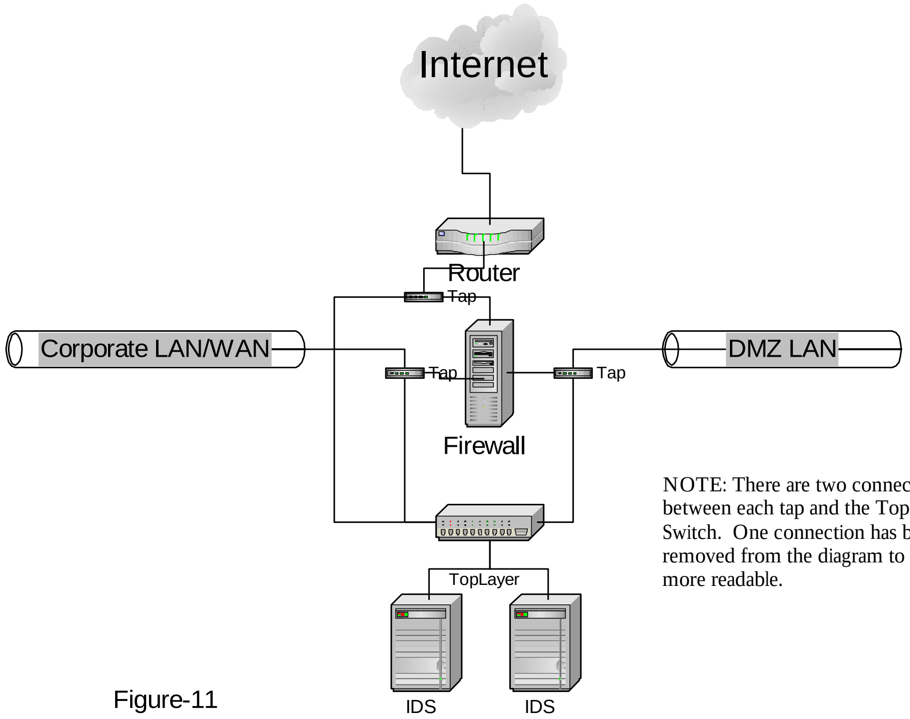

NOTE: There are two connections between each tap and the TopLayer Switch. One connection has been removed from the diagram to make it more readable.

Figure-11

Pros of this configuration

- All pros of using Taps and the TopLayer Switch
- Multiple engines are not needed saving costs
- Very Scalable; if network traffic increases additional IDS can be placed without impact to the network.

Cons of this configuration

- Kills are not supported unless the changes shown in Figure-9 are executed.
- Typically an IDS placed behind the firewall has connection based rules configured to validate the firewall rule-set. In Figure-10 the IDS is listening to traffic from both in front of and behind the firewall so this cannot be done.

Connection Events; These are rules that look at the network layer for Source and Destination information. For example alert when Machine A telnets to Machine B. These rules are used most commonly on an IDS behind a firewall to validate that the firewall rule-set has not been compromised.

Figure-12 outlines a very robust fault tolerant E-Commerce infrastructure, which is being used by many E-Commerce sites today. The site has full redundancy throughout the infrastructure. The solution however is very similar to Figure-10 and Figure-11.

The first set of taps monitors traffics against the outside of the firewall, with the second set monitoring attacks that come through the firewall, as well as mimicking the Firewall rule-set. Due to the utilization of the site two IDS are needed to start with, leaving 2 additional ports for adding one or more IDS in the future.

Management of the whole site is done through the Management LAN, for this segment taps have been placed between the Management LAN and the firewall. From this location attacks into and out of the DMZ from the Management LAN will be detected. With the Taps on the inside of the firewall it cuts down the number of Taps that need to be put in place.

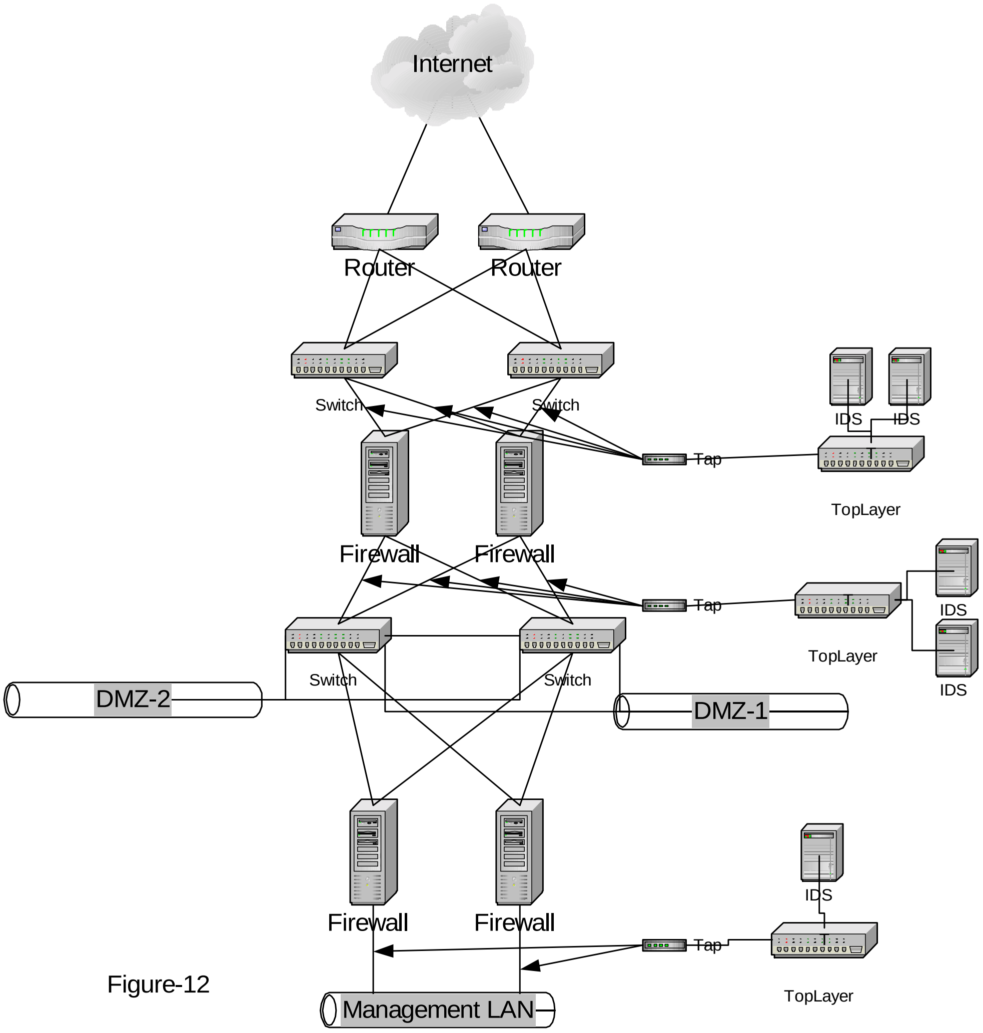

Figure-12

Additional Pros of this configuration,

- Since IDS are spread through the entire infrastructure the policies being used can be specifically tailored to the traffic on that segment
- While the overall number of IDS has increase the complexity of the solution has not increase significantly

Additional Cons of this configuration,

- None

Appendix A

The stealth configuration consists of an IDS with two interfaces. The first interface has all network bindings removed. This prevents the IDS from being attacked. The second interface is then routed to a Management LAN. This allows full management of the IDS without risking the IDS being attacked.

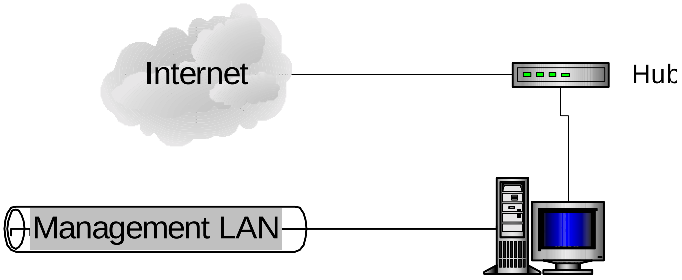

Figure-13

Appendix B

Tap Vendors

There are two main vendors of taps Rioco Direct Ltd (www.rioco.co.uk) who is the only manufacture with a rack mountable multi port tap. Shomiti (www.shomiti.com) also makes a single port tap as well as fiber taps (which can be used to tap Gigabit Ethernet). Rioco is also a reseller for the Toplayer and ISS software.

Shomiti sell mainly through resellers so you will have to try your local Value added resellers for information.

Rioco Direct Ltd. does sell direct and you can contact them at

Darren Murphy
Rioco Direct Ltd
Court 1
Bedfont Lakes North
Challenge Road
Ashford
Middlesex
TW15 1AX
Tel: 08000 360360
Fax: 01784 264364

Generic Tap information

Taps are a four-port device that provides a means of tapping into network traffic viewing traffic on a full-duplex or half-duplex 10/100 Ethernet segment.

Typically, the Tap is deployed on a critical link in a network where network monitoring and analysis capabilities are important. The Century Tap is also fault-tolerant; any loss of power to the device will not impact network connectivity or performance.

Tap Components

The diagram below represents the rear view of a Rioco Tap on the left and a Century Tap on the right. The only real difference between these two is that all four ports are on the front of the Rioco Datatap, making it much easier to use. The Rioco multiport tap is simply four or more of the single port taps mounted in a chasses, with fail over power supplies.

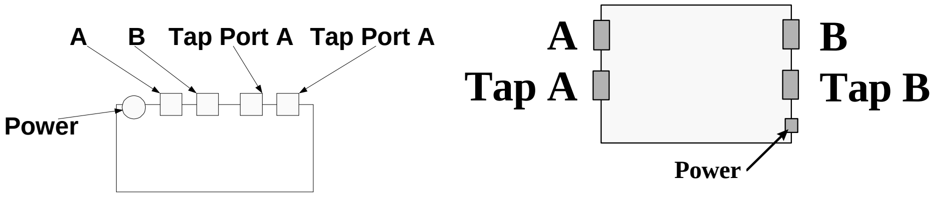

- The A/B ports are directly connected to the network segments you wish to monitor.

- The A/B "Tap Ports" are connected directly to an IDS or into a consolidation device..

- The Power LED depicts whether the Century Tap is receiving power.

Tap Cabling and Operation

The diagram and the following three boxes depict the information necessary to successfully install the Century Tap:

- The diagram shows the basic circuitry between the A/B ports and the A/B Tap Ports.

- The Operation box shows how the A/B Tap Ports are able to monitor the A/B ports. It is imperative the user realizes that Tap Port A (B) mirrors the data received into Port A (B) from the device(s) attached to Port A (B).

- The Cabling Guidelines box depicts the cabling required to achieve the necessary connectivity between the network devices a consolidation device, which could be a hub, regular switch or a TopLayer Switch. You will require two straight-through cable, and two crossover cables. If the traffic from the tap is not consolidated than the two crossover cables are replaced with standard straight-through cables.

- The Cabling Distances box show the maximum distance that will work between any two end points connected through the Tap. The Tap, in essence, is equivalent to a 10-meter length of cable. The diagram for the Rioco tap would be identical to the one below.

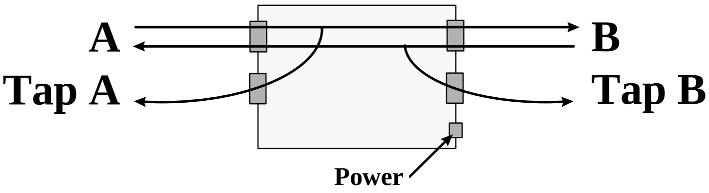

Figure 13. Century Tap Components (logical)

Operation

Tap Port A mirrors traffic received into Port A from the device(s) attached to Port A
Tap Port B mirrors traffic received into Port B from the device(s) attached to Port B

Cabling Guidelines

Port A to Network Link:	Existing Cable
Port B to Network Link:	Straight-Through Cable
Tap Port A to Cisco Switch:	Cross-Over Cable
Tap Port B to Cisco Switch:	Cross-Over Cable

Note: If no link light appears on network devices, swap cables between Port A and B.

Cabling Distances

90 Meters maximum distance between:

Network device connected to Port A and Network device connected to Port B
Network device connected to Port A and Cisco Switch connected to Tap Port A
Network device connected to Port B and Cisco Switch connected to Tap Port B

Appendix D

The Top Layer AppSwitch is not a traditional packet switch – it is a flow switch. What does this mean?

Packet switching involves looking at each packet in isolation to determine where to forward it based on the Layer 2 MAC address, Layer 3 IP address or Layer 4 TCP or UDP port. All three of these addressing methods may be used to make forwarding decisions. But the point is – any of these methods can be deployed by looking at a single packet.

Flow Switching involves looking at traffic as bi-directional flows between end systems on the network, and using information from previous packets to make decisions on packet forwarding. This is critical to recognizing dynamic Layer 4 TCP or UDP port assignments that are part of FTP, H.323 and many other applications. Only be looking at information present in bi-directional flows can a switch learn the dynamic port assignments in FTP flows, or H.323 streaming flows. Once these dynamic port assignments are learned, a flow switch can apply traffic policies to the flows.

Why is this important for Intrusion Detection? It is critical for an IDS to see both directions of a connection to be able to detect Syn Floods and many other types of attacks. The AppSwitch can assure that both directions of a connection are sent to a single IDS. When TAPs are deployed, the AppSwitch allows both TX and RX streams from the TAP to be sent to IDS engines.

Futures for the AppSwitch

- While today the AppSwitch mirrors traffic to the IDS engines, in the future we could allow the IDS to send kills to the AppSwitch directly – stopping suspicious traffic quickly.
- Currently there are a few intrusion filters implemented directly in the AppSwitch. These are single packet intrusions that are easy for a switch to detect. We could add more of these in the future. As more of these are implemented in a switch, the IDS can be configured to look only at the more complex intrusions, thus freeing up CPU cycles in the IDS.
- The IDS could detect a suspicious flow, and tell the AppSwitch to re-direct that flow to a “honey pot” such as the Resource device. This allows the intruder to “think” they are successfully in the network, when in reality, forensic data is being collected on them.

Competition

There are other switches that do Layer 4-7 switching. Arrowpoint, Alteon and Foundry make switches that are certainly Layer 4 capable, but their Layer 7 functionality is limited to HTTP URL filtering. None of these switches today support dynamic port recognition. And more importantly, none of them support the FlowMirroring function so critical to distributing traffic to Intrusion Detection devices. This is not to say they never will, but none of these vendors talks about this today.

Product Overview -- AppSwitch 2000 and 2500

The AppSwitch™ family from Top Layer Networks is the industry's first 7 Layer switching family providing Application Control™ to intelligently route traffic at full wire speed. These easy-to-use networking devices identify and control specific business-critical applications and provide prioritization and Admission Control based on stateful analysis of information contained in the higher layers of each information packet. This level of high-speed intelligent processing is enabled by custom Application-Specific Integrated Circuits (ASICs) that allow the AppSwitches to provide simpler and more effective switching than traditional Layer 2, Layer 3, or even Layer 4 routing and switching devices.

The award-winning AppSwitch family opens up new opportunities for Business Driven Networks™. You can tune your infrastructure to best support your most critical business applications. You can guarantee bandwidth for high-priority or real-time applications and implement policy-based networking that allows you to resolve network congestion, minimize WAN costs, and gain increased control over network traffic. By controlling application traffic using information up to the "Top Layer" you can best allocate your network resources to support your business priorities.

Application Control of Network Resources

The AppSwitch family provides the Application Control you need to manage the behavior of your network and optimize the value of your infrastructure investments. AppSwitches enable Business Driven Networks that implement policy-based networking to make sure that network resources are allocated according to business priorities.

The computing and communications requirements of the enterprise network are dynamic, but network infrastructure until now has been relatively static compared to the application environment. The vast majority of desktops are networked, and the enterprise continues to add distributed applications that consume network resources. Both the volume and the types of traffic are constantly changing. The AppSwitch family intelligently monitors this behavior and automatically responds by ensuring that you receive the prioritization, Admission Control, and security you require for your business-critical applications according to the policies that you establish.

AppSwitches allow you to monitor network utilization by application, user, and workgroup so that you can evaluate behavior and tune the network to reflect business priorities. Configuration of the network is finally easy, because AppSwitches identify and index users, host machines, and both IP and MAC addresses. The AppSwitches then enforce policies according to your business rules, allowing you to monitor and define network behaviors.

For example, if you find that a low-priority application, workgroup, or user is consuming excessive network resources while a mission-critical e-commerce application is starved for bandwidth, you can easily fix this problem by using the web-based interface to change network behavior. Similarly, if your company has made the investment in enterprise software applications such as SAP, Oracle, or PeopleSoft you want to make sure payback is realized and that these mission-critical applications are given the high-priorities they deserve. The "touch-every-packet" switching engine of the AppSwitch family reshapes the intelligence, performance, and future-readiness of today's networks by delivering unprecedented control over the availability and performance of business-critical applications and transactions.

Business Driven Networks with the AppSwitch 2000

There are several elements necessary to build a Business Driven Network. The Infrastructure Layer is the installed base of cabling, switches, and routers in the network up to Layer 3. The Relationship Layer incorporates the host and user names of all devices on the network along with their Layer 3 and MAC addresses. The Service Layer is where the real benefits of 7 Layer Application Control come into play. At congestion points in the network, applications are prioritized according to business policies.

Admission Control is equally important, and AppSwitches allow you to limit bandwidth by application. For example, you can limit the amount of bandwidth available for lower-priority traffic flows, such as file transfers or newsfeeds, or constrain the network resources available to lower-priority workgroups.

Included in Admission Control is a built-in firewall capability allowing you to manage relationships by building secure extranets with your customers, prospects, partners, and suppliers. You can create intelligent firewalls based on applications, users, or workgroups to safeguard access to enterprise resources while guaranteeing availability for high-priority applications and limiting network resource consumption.

Services accounting is supported via the TopFlow™ protocol to enable integration into your existing billing applications for departmental chargeback. Application Layer traffic monitoring allows you to review flows with increased granularity, without the need for costly RMON probes. Data is presented in a standard form for review in leading data analysis software packages.

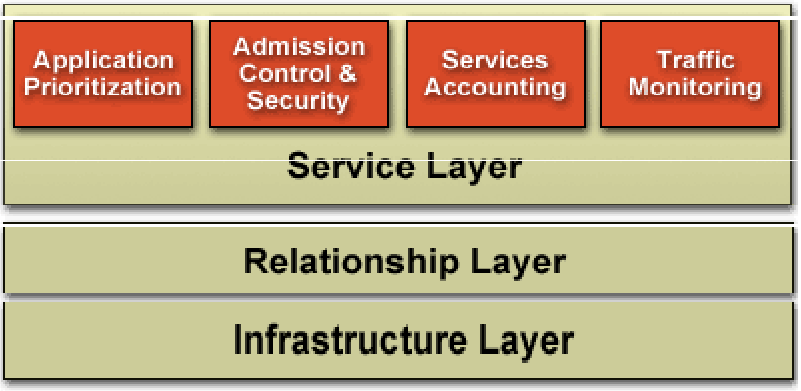

Introducing 7 Layer Application Control

The prices of network switches have declined drastically in the past few years to the point where it is common to see 10/100 Layer 2 desktop switches deployed in many wiring closets, with Gigabit Ethernet links to the network center. The LAN core is increasingly built today with Layer 3 switches instead of classical routers, and some of these switches also perform some level of Layer 4 prioritization. Managers of enterprise networks looking to the future may wonder what the next step might entail.

Top Layer Networks believes it is the deployment of Business Driven Networks built with networking devices that perform 7 Layer Application Control. Even with the low cost of bandwidth available today in the LAN, there are congestion points in the network. These include the LAN/WAN border, where WAN bandwidth is not "free" at all, the server farm where the ability of servers to deal with multiple requests is limited, and desktop switch aggregation points.

Here are just some of the reasons for deploying 7 Layer Application Control in your enterprise network:

- Businesses invest a lot of money in mission-critical business applications such as SAP, Oracle, and PeopleSoft. Top Layer Networks enables **tuning of the network to take maximum advantage of mission-critical business applications**.

- The types and **amounts of traffic flowing in networks constantly changes**. Fixed configurations for traffic prioritization in Layer 3 or 4 network devices are not sufficient to deal with this. Only 7 Layer Application Control technology from Top Layer Networks can know about this behavior and enforce the desired priorities for business applications.

- **It's not just about traffic prioritization — Rate Limiting** is just as important. One of the biggest advantages of 7 Layer control is the ability to limit bandwidth by application.

- **Voice over IP** is making its way into data networks. Top Layer Networks assures crisp clear voice signals while keeping business applications flowing smoothly.

- Support of some type of **video** (such as conferencing or distance learning) traffic will be required. Top Layer Networks can guarantee the video bandwidth needed while keeping business applications performing as expected.

- **Security** — 7 Layer technology provides the highest level of firewall security available, and stateful analysis of all flows enables powerful firewall capabilities.

- **Configuration** of a network is finally easy. Automated configuration and the Web-based interface tools make building the network as easy as possible.

- **Save on WAN costs** — The AppSwitch family makes the most efficient use of WAN links, allowing the use of lower-speed and lower-cost links than would otherwise be possible by prioritizing and/or limiting the use of the WAN by application.

- **Policy** based networking is realizable today with 7 Layer switching. AppSwitches are ready to work with established industry standards, such as LDAP.

- **Monitoring** of network behavior just took a giant step forward. Top Layer Networks technology provides more information about network traffic than has been possible up till now, without the need for costly RMON probes.

- All of these benefits are realizable today at a **cost** similar to today's Layer 3 and 4 switches.

Directory Enabled, User-Based Policy Management

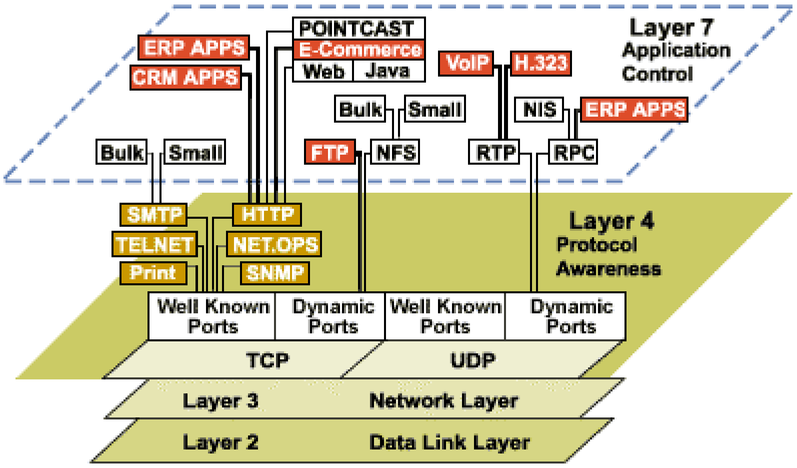

Benefiting from 7 Layer Application Control

Layer 4 switching provides only a subset of the functionality of 7 Layer Application Control. Many applications use a fixed control port but dynamic data ports. Stateful traffic inspection allows you to determine traffic types even if traffic is associated with dynamic ports. In this example, the HTTP traffic uses a fixed port but carries multiple transaction types. The AppSwitch family allows you to distinguish these flows so you can further segment the HTTP flows. For example, you could "turn up" ERP and e-commerce applications at the expense of Pointcast and Web browsing.

7 Layer Application Control allows you to automatically analyze a stateful session to segment traffic flows. For example, although FTP file transfers have a fixed control port, the actual file transfer ports are dynamically assigned. Since FTP traffic is often large files with the potential to clog the network during peak periods, you need to be able to automatically provide a stateful inspection of the whole FTP sequence so large file transfers can be prioritized. AppSwitches can analyze the session setup, map the port

assignments, and prioritize file transfers. Similarly, you can use AppSwitches to limit the amount of bandwidth FTP can consume.

A Scaleable Switching Network Architecture

Built on a scaleable and extensible architecture, AppSwitches eliminate bottlenecks at the LAN/WAN edge, server connections, and LAN traffic aggregation points. Application Control requires high-speed processing and analysis of traffic by an intelligent switching engine. Top Layer Networks has achieved this with its TopFire™ chip sets of patent-pending ASICs.

The first chip is an MII Octal MAC (MOM) providing 8 ports of MII interface circuits used for Fast Ethernet switching including MAC layer support. The packets then travel to the next chip—the Queue Manager (QM™)—which shuffles traffic in and out of the high-speed packet memory. The QM passes the packet to the Relay Engine (RE™) chip, which performs the stateful flow analysis and determines the proper forwarding and QoS parameters for the packet. The RE chip next instructs the QM to transmit the packet. In the AppSwitch, a second MOM chip is connected to the internal high-speed non-blocking bus, and another RE chip is used as a switch control processor.

TopPath™ software works directly with the TopFire chip set. By using the wire-speed analysis of protocol data within data frames, TopPath automatically monitors, analyzes, and maintains a table of up to 16,000 traffic flows. A second table maintains a priority list for business applications. This table automatically assigns priority to the traffic flows based on the application, thus eliminating the burden of manual setup typical of traditional routers.

The basic three-chip set is extended to include another MOM chip and another RE chip used as a Background Engine. You can now benefit from Application Control and the ability to cope with growth and change, and automatically prioritize business applications and perform Admission Control for limiting bandwidth or security.

Custom ASIC and TopLayer Software Enable Application Control

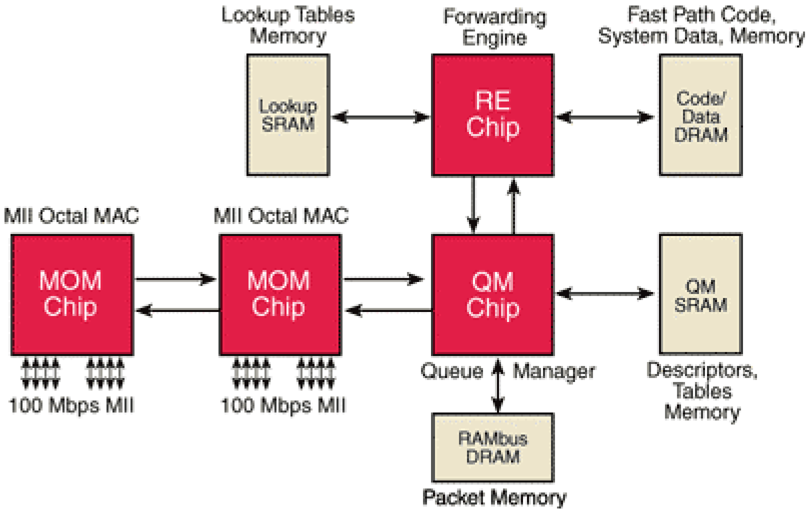

AppSwitch 2000 and 2500 Technical Specifications

General Features

- Dual Power Supply Inputs
- 16 MB Background Engine Memory
- 4 MB Forwarding Engine Memory
- 15 MB Compact Flash Storage

- Console on Any Port
- Command Line Interface
- Installation CD
- Fan Speed/Motion Detection
- Basic Logging

Layer 2 Features:

- Protocol Independent Data Link Bridge
- 802.1d Compliant
- Wire-Speed Forwarding
- 8,000 MAC Addresses
- Port Failover (Spanning Tree)
- Fast Port Failover (Non-Spanning Tree)
- Cisco-Compatible Link Aggregation
- 802.1q VLAN Tag-Friendly
- SNMP Bridge MIB
- Fast Ethernet Link Aggregation
- Fast Ethernet Link Failover Resiliency

Layer 3 Features:

- L3 IP Forwarding
- Wire-Speed Forwarding
- IP-Static Routes
- IP-RIPv1
- IP-RIPv2

Layer 4 Features:

- L2/L3 Forwarding with L4 QoS
- Wire-Speed Forwarding
- 16,000 Simultaneous Flows
- Flexible TCP/UDP Control

Layer 7 Features:

- Service Classes
- 802.1p/TOS/DiffServ Bit Setting

- TopFlow Application Monitoring and Data Collector
- FTP Dynamic Ports
- HTTP URI Applications
- Sun RPC-Based Applications
- Layer 7 Arbitrary Offset Applications
- Application Graphing

TopView™ Web Management:

- Embedded Welcome Page
- Application Definition Library (ADL)
- Top Layer ADL Help
- TopFlow Data Collector Utility
- Enhanced Graphing and Reporting
- Basic Authentication
- Distributed IP Address Setting
- Designated AppSwitch Election
- Configuration Files Editor
- State Browser Accessibility
- Graphical Model of Logical Network
- Multiple Zones
- Default Policy Set Templates
    - Firewall
    - DMZ Firewall
    - Standard
    - No Access
    - Baseline
- User-Defined Policy Sets
- Telnet Support
- Graphs for Monitoring
- MIB 2 SNMP Agent
- Mini-RMON — Groups 1,2,3,9

Physical Interfaces:

- 10 BASE-T/100 BASE-TX (RJ45 Connectors)
- 100 BASE-FX (SC Connectors)

Safety Certifications:

- Agency Certifications: UL1950
- CSA22.2 No. 950
- EN 60950, IEC 60950
- AC Protection: 4 A Fuse
- Over Temperature Protection: Automatic Warning at 60º C (140º F)

Regulatory Markings:

- CE Mark
- C-TICK Mark
- TUV/GS Mark

Electromagnetic Emissions:

- FCC Part 15B Class A
- ICES003 Class A
- CISPR22 Class A
- EN55022 Class A
- AS3548 Class A
- VCCI Level A

Immunity:

- EN 50082-1

LAN Interface Specifications:

802.3, 802.1d, 802.3u, 802.1p

AppSwitch Models:

AS2002
12 10 BASE-T/100 BASE-TX Ports and
2 100 BASE-FX Ports

AS2006
8 10 BASE-T/100 BASE-TX Ports and
6 100 BASE-FX Ports

AS2014
14 100 BASE-FX Ports

AS2050
Redundant 48 VDC Supply

AS2502
12 10 BASE-T/100 BASE-TX Ports and
2 100 BASE-FX Ports
Console Port
Internal Power Supply
Standard Rackmount Design

Product Dimensions:

	AppSwitch 2000	AppSwitch 2500
Height	7.5 cm (3 inches)	6.5 cm (2.55 inches)
Width	20 cm (8 inches)	43.8 cm (17.25 inches)
Depth	30 cm (12 inches)	33 cm (13 inches)
Weight	2.7 kg (6 pounds)	5 kg (11 pounds)

Environmental:

Operating Temperature	0º C to 40º C (32º F to 104º F)
Non-Operating Temperature	-25º C to 70º C (-13º F to 158º F)
Relative Humidity	5 - 95% Non-Condensing
AC Input **(External AC/DC Converter)**	100 to 250 VAC Auto-Ranging
Frequency	47 to 63 Hz
AC Input Current	0.85 A @ 110 VAC 0.42 A @ 240 VAC
Power Consumption	100 Watts (Maximum)
External AC to DC Converter	48 VDC
DC Input	42 VDC to 56 VDC

Glossary