White Paper

Intrusion Detection

Deploying the Shomiti Century Tap

METASeS™

Shomiti Tap Deployment

Purpose of this Paper

The scalability of Intrusion Detection Systems (IDS) is often an issue when deploying an IDS solution on a large network or ambiguous perimeter. Augmenting the IDS with additional hardware is one way to overcome some of the scaling issues.

This white paper describes some possible solutions to traditional IDS problems through use of the Shomiti Century Tap and Cisco 2900 series switches.

Hardware components (an overview)

The hardware configuration has the following components:

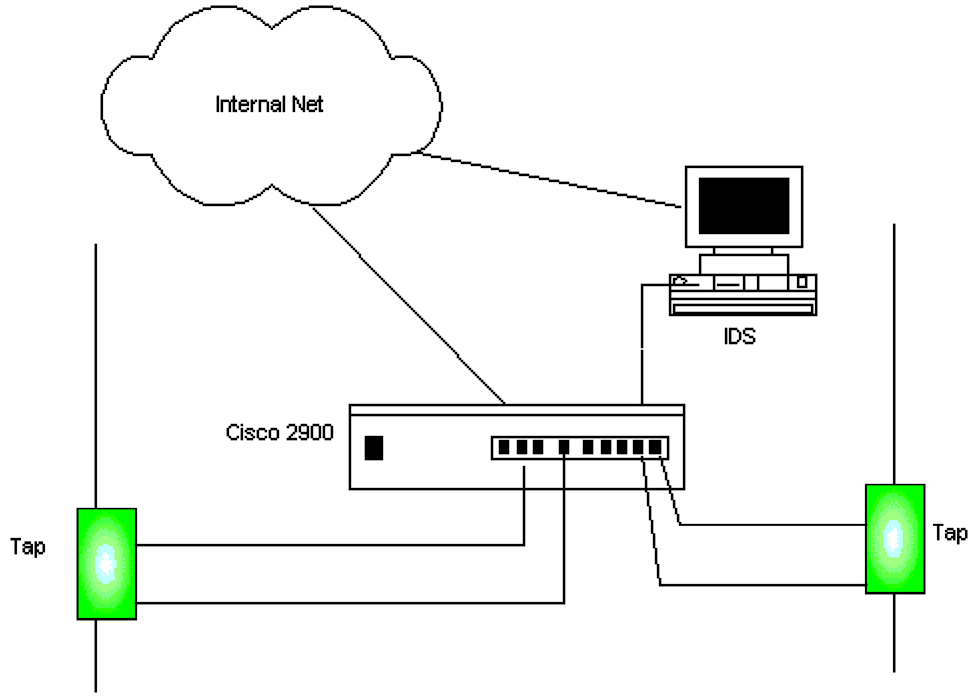

Figure 1. Tap/Switch/IDS deployment components

Taps

The Century Tap is a single port device that provides a method of directly viewing traffic on a full-duplex or half-duplex 10/100 Ethernet segment. Typically, the Century 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.

Switch

The low cost of the Cisco 2900XL Series switch coupled with the 3.2GB back plane (used for packet queuing) make it well suited for this configuration. The 2900XL switch comes in two models: the 24 port and the 12 port. Either of these models will work well, however several factors play a part in deciding which model to use. For example, if there is already a 24 port switch in the current corporate inventory that is not gainfully utilized then it will do just fine. Remember that each tap will have 2 output cables, so there will always be twice as many ports required as there are taps. Also 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.

IDS

The IDS can be any available IDS. This configuration has been tested with ISS RealSecure and NSW Dragon. The IDS will require two Network Interface Cards (NIC). One NIC will be used for monitoring the traffic, and the other will be used for remote management. This second NIC is often referred to as the ‘Out-of-Band’(OOB) NIC due to the fact that it is commonly configured to be both physically and logically separate from the network being monitored. A separate NIC for remote management and data feeds helps the Monitoring NIC to function at full capacity while also allowing a measure of security for the IDS itself.

Internal Network

Both the Switch and the IDS will need to be connected to the internal network. This allows you to remotely manage both as well as get data feeds from the IDS. When monitoring networks that are not saturated, the ‘Out-of-Band’ NIC of the IDS can be connected directly back into the switch to allow only one connection into the internal network while at the same time allowing secure management of the switch.

The Shomiti Century Tap Configuration

The Century Tap is a single port device that provides a method of directly viewing traffic on a full-duplex or half-duplex 10/100 Ethernet segment. Typically, the Century 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.

Century Tap Components

The diagram below represents the rear view of the Century Tap:

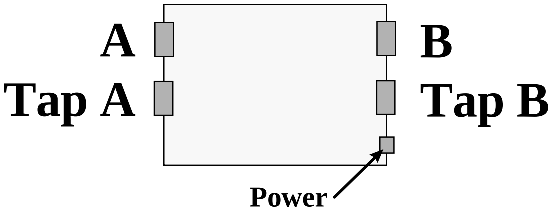

Figure 2. Century Tap Components (physical)

- The A/B ports are directly connected to the network segments you wish to monitor.
- The A/B “Tap Ports” are connected to the Cisco 2900XL Series Switch.
- The Power LED depicts whether the Century Tap is receiving power.

Century 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 realize 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 and the Cisco 2900XL Series Switch. You will require one straight-through cable, two crossover cables, and the existing network cable.
- The Cabling Distances box show the maximum distance that will work between any two end points connected through the Century Tap. The Century Tap, in essence, is equivalent to a 10 meter length of cable.

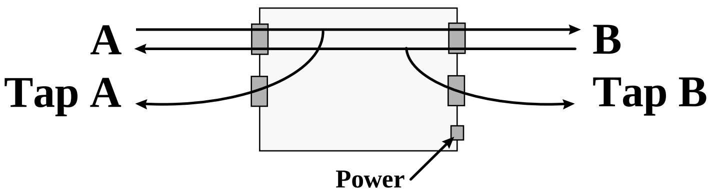

Figure 3. Century Tap Components (logical)

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

The Cisco 2900XL Configuration

The low cost of the Cisco 2900XL Series switch coupled with the 3.2GB back plane (used for packet queuing) make it well suited for this configuration. The 2900XL switch comes in two models: the 24 port and the 12 port. Either of these models will work well, however several factors play a part in deciding which model to use. For example, if there is already a 24 port switch in the current corporate inventory that is not gainfully utilized then it will do just fine. Remember that each tap will have 2 output cables, so there will always be twice as many ports required as there are taps. Also 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.

Out of the Box

When a new switch is first plugged in (there is no power switch), and a terminal is connected to it using the ‘Console’ cable, the user is prompted to enter network specific information such as the IP address to be assigned to the switch. The ‘enable’ password will also be set at this time. If, however, the switch is not new, the user must enter the ‘setup’ command to invoke this questionnaire manually.

It is a good practice to have an IP address already selected for the switch. This IP address should be a viable IP address on the internal network. All other network information (i.e. gateway) will typically be identical to all other devices located on the network.

Switch Related Concepts

It is important to review several concepts associated with a switch. They are the following:

- **Switch vs. Hub** – A traditional hub can be represented as an open hallway. As we plug twisted pair cables into the hub, it is as if we are opening up a door along that hallway for each cable we plug in. We need to assume that there is a person located in each doorway and when this person wishes to communicate with another he yells out the name of that person and then continues to communicate his message. In this way every one person in every doorway must hear this conversation even though they only respond when it is directed at them. In the real world of a hub, this causes network traffic to be duplicated for every network device connected to the hub and is not very efficient. This can turn a hub into bottleneck.

  A switch is similar to the hallway described above. The difference is seen when a person decides to communicate with another. Instead of yelling out his message for all to hear, he simply writes it down on a piece of paper, folds that paper into an airplane, and sends it directly to the intended party thus relieving the other people in the hallway from ever having to deal with this communication. Since the traffic being communicated is only duplicated once, this makes high bandwidth networks run faster and more efficient.

- **VLANs** – A basic switch can only support one logical sub-net. To make switches more flexible, the VLAN (Virtual Local Area Network) was invented. This allows the switch to be configured in a way so that logically it appears as two different switches. Imagine that instead of one large hallway as described above, we now place a brick wall in the middle of that hallway thereby dividing it into two separate hallways. Now each hallway could correspond to a different sub-net however, the sub-net is not required to be different. In summary, a VLAN is a superswitch, which allows dynamic reconfiguration of networks and creation of virtual sub-nets.

In the tap/IDS/switch configuration the reason for having two separate VLANs is due to the fact that the network being monitored and the internal network may be located on the same logical sub-net. This can have a couple serious implications. First, traffic that is being monitored could be duplicated by the switch and sent onto the internal network, causing erroneous packets to run amok. The second possibility could be that the monitored network is the perimeter on the outer side of the firewall and it has the same logical sub-net mask as the inside. In this scenario an attacker could send packets into the network, by-passing the firewall. Even though the output lines of the taps are only one way, once the packet hits the destination it is routed back out the network through the normal routes for the internal network.

- **SPAN** – The Switch Port Analyzer (SPAN) port is commonly used to view network usage statistics.   (This is not to be confused with the concept of a Spanning Tree). The SPAN port is also commonly referred to as the ‘spy’ port or ‘mirror’ port. On older models, there was a separate port hardwired as the SPAN port. On the 2900XL series however, it can be configured to be any port as long as it is located in the same VLAN as the ports it is configured to monitor. The switch will forward all traffic from the ports being mirrored to the SPAN port. This allows us to forward all incoming traffic from the taps to the one port that is connected to the monitoring NIC of our IDS.

Switch Management

The IP address of the switch will be available for remote management through any port located in VLAN 1. Originally all ports will belong to VLAN 1. Once an IP address has been assigned to the switch, the switch can be accessed using telnet. Once the switch is placed into production it is a good idea to first connect to the IDS system using secure shell (ssh) and then telnet into the switch so as not to allow the password to be transmitted in the clear where prying eyes could find it.

Switch Cabling Configuration

The figure below illustrates the common cabling configuration of a Cisco 2912XL. The switch is configured into two separate VLANs. Ports 1 & 2 belong to VLAN 1 and ports 4 – 12 belong to VLAN 2 (port 12 acting as the SPAN port). Port 3 is not used in this configuration.

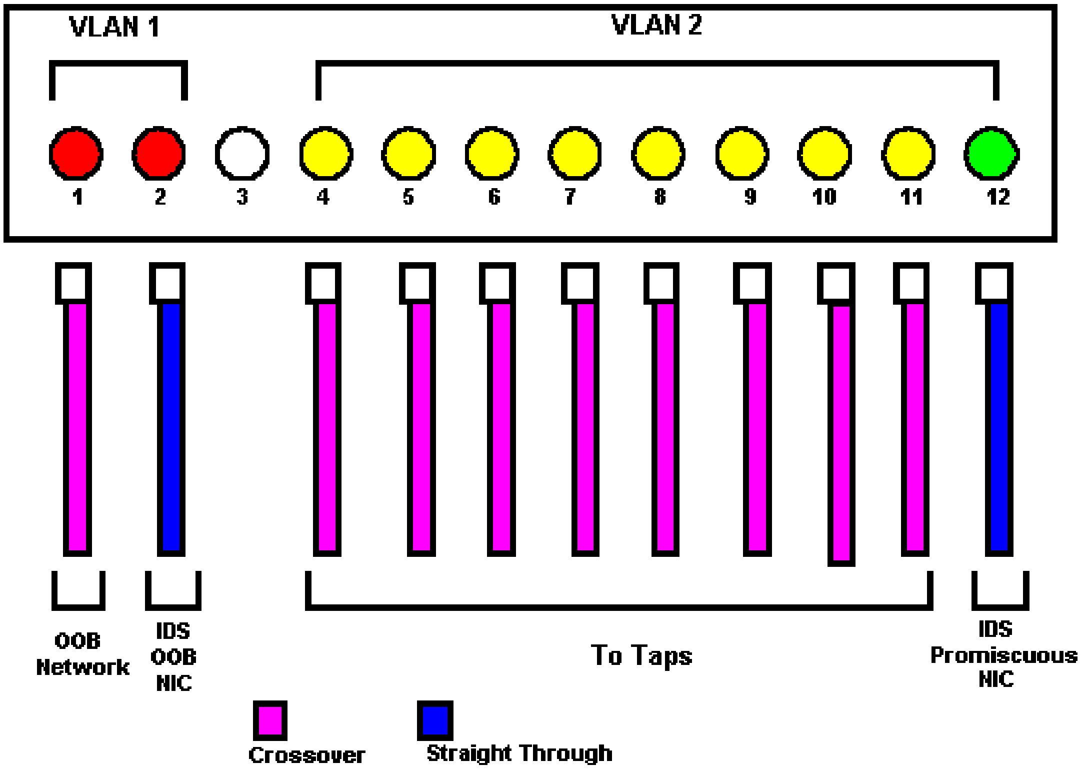

Figure 4. Cisco 2912XL cabling configuration (physical & logical)

Once the taps have been placed on the network segment to be monitored, two crossover cables per tap should be plugged into VLAN 2 of the switch. It is imperative to have separate VLANs, as packets from the tapped wire could run amok on the internal network if if the switch is configured with only one VLAN.

The last port in VLAN 2, (typically 12 or 24), should be designated the SPAN port. The main reason for doing this is so that you can identify the configuration by sight. The SPAN port will then be connected to the monitoring NIC of the IDS and will forward all traffic from the monitored wires to the IDS for examination.

The logical configuration of the switch can be accomplished through use of a script to make large deployments much easier. The script is included with this paper.

The IDS Configuration

The IDS can be any available IDS. The tap/switch configuration has been tested with ISS RealSecure and NSW Dragon. There should be two Network Interface Cards (NIC) on the computer that the IDS software is running on. This allows for communication to a console or central data repository.

Generic IDS Considerations

There are several general concepts that apply to IDS, yet they may differ in execution. They include the following:

- **Two Tiered Architecture** - The common IDS network architecture is composed of two separate devices. The first device, of which there may be many, is used for collecting data and is commonly referred to as the 'Engine' or 'Sensor'. The second device, of which there is commonly only one, is used for management of the ‘Engines’ and is referred to as the ‘Console’ or ‘Data Repository’. Both the 'Sensor' and the 'Console' are normally an Intel or SPARC based platform.

  The bulk of the work is done on the 'Engine', as it is where the traffic is actually examined and subsequently flagged or ignored. This determination is accomplished through use of filters, attack signatures, and other policy elements.

  It is on the 'Console' that the data is usually stored after being collected from the 'Engine'. Policies to determine what to monitor for are developed here and 'pushed' out to the 'Engines'. Reporting is also normally a console function. Optimally there will be many 'Engines' to one 'Console,' allowing for distributed management.

- **Dual NICs –** The platform that supports the IDS 'Engine' software should have two Network Interface Cards (NIC). One is used for monitoring the traffic and is connected to the SPAN port on the switch, and the other is used to for remote management and data feeds to the remote 'Console'. The latter NIC is commonly referred to as the Out-of-Band (OOB) NIC. In many cases there will be an onboard NIC (i.e. connected directly to the motherboard) and a PCI NIC. The IDS will yield the best performance if the onboard NIC is designated as the monitoring NIC.

- **Jurisdiction –** When deciding how to position an IDS 'Sensor' on a perimeter network, it is important to understand your goals in deploying an IDS. Depending on the amount of traffic on the monitored wire, it may be a good idea to simply monitor for inbound attacks. This can be accomplished with careful filter development as described below. Commonly, an IDS is designed to protect the corporation from malicious users both internal and external to the network. The question might be “Which do we wish to protect: our network from the Internet or the Internet from our network?” If down stream liability is a concern then it may be appropriate to add another IDS 'Sensor' to the perimeter due to performance concerns on the 'Sensor' monitoring inbound traffic. Since each 'Sensor' will have its own name and IP, this will make it easier to determine from which side of the network an attack is occurring when viewing the master Console.

- **Filters –** There are normally ways to designate traffic the IDS should not examine. Many times, use of these filters actually optimizes the performance of the IDS by allowing it to focus cycles on only the important traffic. These filters usually are defined by IP addresses, Netmask, Protocol, and Port. To be able to develop a good set of filters, it is important to become familiar with the topology of the network in question. An accurate physical/logical network diagram is invaluable in this regard. You will need some time after the initial install of the IDS and hardware configuration to get an accurate picture of the traffic that is actually travelling over the monitored segments. You should develop a log that includes the key servers and workstations discovered while monitoring network traffic. Many times the IDS will allow you to alias an IP with an identifying label. This makes it easier to determine what assets are involved in any particular incident or developing a policy that is relevant to the devices on the network.

- **Attack Signatures** – IDS software commonly is shipped with a database of known attack signatures. This is a collection of identifying characteristics of various known network attacks, which has been developed by the IDS vendor, against which collected traffic is compared. If the examined traffic exhibits the characteristics of a known attack in the database then it is flagged and handled according to the parameters allowed by the specific IDS. Many of these signatures may be irrelevant to the network being monitored. For instance, if it is known that there are no POP servers located on the monitored network then indicating that this check should not be performed helps to optimize the IDS.

- **Connection Based Event Collection** - This aspect of IDS allows the user to specify network events that are deemed important, yet which are not associated with known attacks. These rules are commonly defined by IP addresses, Netmask, Protocol, and Port. They are very similar to the rules designated as filters, and are defined in much the same way. An example would be to flag all traffic as high priority that originates from the Internet and is a TCP packet destined for port 23 on any devices associated with your network. This would be an example of someone using telnet to access your assets from an untrusted network.

ISS RealSecure Considerations

Although many concepts concerning IDS are universal, there are technical differences between IDS products. The key elements for the ISS RealSecure product are as follows:

- **Platform Support** – While the ISS RealSecure Console runs only on WinNT 4.0, the Network Engine has been designed to run on two platforms: Solaris SPARC and WinNT Intel. The primary difference between these two is the performance factor; RealSecure runs best on a SPARC architecture. It is commonly believed that a WinNT platform with extra processors will allow a performance boost, but this is not so, as the RealSecure software has not been engineered to take advantage of the multi-processor or multi-threading features of WinNT. Several factors affect the choice of platform. If remote management of the operating system is desired, then Solaris should be chosen. If the link being monitored is greater than a 10MB link, then Solaris should be chosen.

- **Solaris Installation** – Installing RealSecure onto Solaris is done merely with a pkgadd command. The readme file on the distribution CD will have all of the current filenames and path requirements. Once installed, the RealSecure Service can be started by using the command /etc/init.d/realsecure start or stop. By default RealSecure will open up two TCP ports. The first is 2998, referred to as the daemon channel and the second is 901, referred to as the event channel. It is important to note that the traffic flows bi-directional over these ports. This should be remembered if a firewall is placed between the console and the engine.

- **WinNT Installation** - Installing on WinNT is a bit more complex, but easily accomplished through the installation wizard. A key point to make, however, when installing either the Console or the Engine is that if the strong encryption is desired the Certicom Crypto that is located on the distribution CD should be installed first. The RealSecure service can be controlled through the ‘services’ icon in the control panel. The prior version of RealSecure should always be removed properly before installing the current version.

- **Authentication** – Keys are used to authenticate valid consoles to engines. When the RealSecure Console is installed, a key file made up of the host and user names will be created and placed in the \ISS\RealSecure\Keys tree. There will be two different folders in here with the extension of PubKey. The key will be placed according to the version: Domestic or Export. For the authentication to work correctly, this key must be placed in the same directory structure located on the Engine. The Engine service should be stopped and restarted for this to take affect.