securityMETRICS

APPLIANCE

INTEGRATED VULNERABILITY ASSESSMENT, INTRUSION DETECTION AND PREVENTION

A TECHNICAL WHITE PAPER

INTRODUCTION, IMPLEMENTATION AND TECHNOLOGY

Table of Contents

About SecurityMetrics

SecurityMetrics, Inc. is committed to keeping data secure with comprehensive solutions that maintain network security and block Internet attacks in real-time. SecurityMetric develops security appliances and Internet vulnerability testing services. SecurityMetrics is a privately held corporation headquartered in Orem, Utah.

SecurityMetrics, Inc. was founded in February of 2000 and received initial funding from Software Development Corporation, the company that developed WordPerfect for UNIX/Linux for Novell, Inc. and Corel Corporation.

The SecurityMetrics Appliance provides a dynamic security solution, coupled with intrusion detection, intrusion prevention, vulnerability assessment and firewall protection. The Correlated Intrusion Assessment feature makes our Security Appliance the most advanced Intrusion Detection System on the market today.

SecurityMetrics offers a **risk-free evaluation** of the Security Appliance. To arrange an evaluation, please call 801.724.9600 or email sales@securitymetrics.com.

IDS Introduction

Firewall technology has not been able to successfully safeguard private data. Firewalls do a great job to filter or mask ports on a host system. However, most companies must open ports on their firewalls to provide web, email, FTP, domain name (DNS) and other services.

As soon as an administrator opens ports on their firewall then those ports are no longer protected. A good analogy would be to lock every window in your house with the back door, the basement door, the sliding family-room door, even the pet door locked, but leave the front door open. Everything is very secure except the front door. An attacker would have little problem compromising your home.

Intrusion Detection Systems (IDSs) have more intelligence and are built to fill in the gaps left open by firewalls. An Intrusion Detection System is a device that monitors all network traffic. It analyzes the traffic in real-time to determine if someone is sending attacks or malicious traffic on your network. The analysis normally incorporates pattern matching and other techniques that are fast enough to analyze all packets on busy networks.

On basic Intrusion Detection Systems an IT administrator is notified in real-time when an attack occurs. Also, good Intrusion Detection Systems maintain a database of all attempted attacks so that system administrators can look at a central IDS source for all attack information. This circumvents the need to look at each individual computer log file to obtain attack information.

On advanced Intrusion Prevention Systems the attacker may be automatically disconnected from your network in real-time as soon as an attack is identified. Also, Upstream Providers may be notified of attacks originating from their networks. And finally, the duration of the attacker disconnection can be determined by the IT administrator, so the IP address may be used by potential customers minutes after an attack has been blocked.

Network-based Intrusion Detection Systems

There are two common types of Intrusion Detection Systems. The most common type of IDS is network-based. This is due to the fact that only one network-based IDS may be needed on a simple network.

Network Intrusion Detection Systems require little maintenance because no agents or software need to be installed on any of the computers on the network. The Network Intrusion Detection System analyzes all packets on the network whether they originate from outside or inside your firewall.

Network-based Intrusion Detection Systems use a number of methods to analyze all network packets:

Pattern Matching – The IDS checks each packet for an attack signature pattern.

Frequency – The IDS watches for the frequency of certain types of packets.

Anomaly – The IDS searches for packets with known suspicious anomalies.

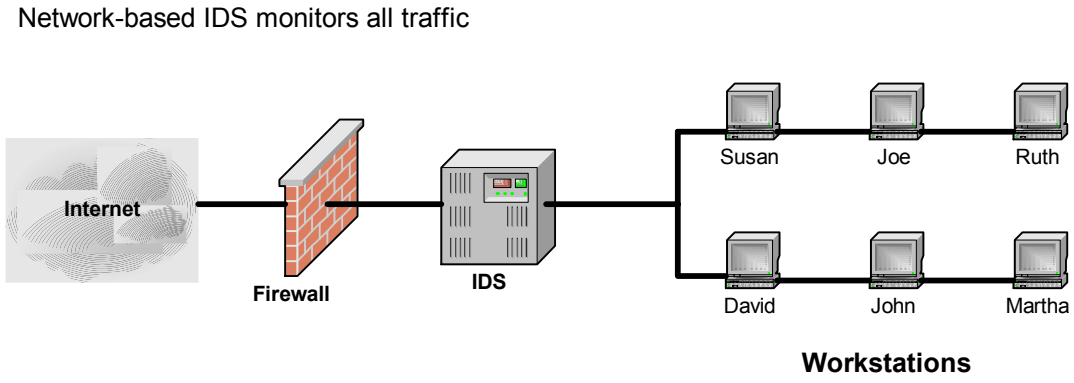

Host-based Intrusion Detection Systems

As the name infers, a host-based IDS resides on the computers or hosts to be monitored. The software is installed and maintained on each host to be monitored. The installation and maintenance of software on all computers to be monitored can be a burden to an IT administrator.

There are some advantages to a host-based IDS, such as the use of audit and system logs to corroborate an attack. However, a well-tuned Network IDS is able to find the attacks on the network.

Even the use of logs does not guarantee against the occurrence of false positives. Also, host-based Intrusion Detection Systems tend to be costly. Many IT administrators are not convinced the use of host-based Intrusion Detection Systems justify the extra labor, maintenance and cost.

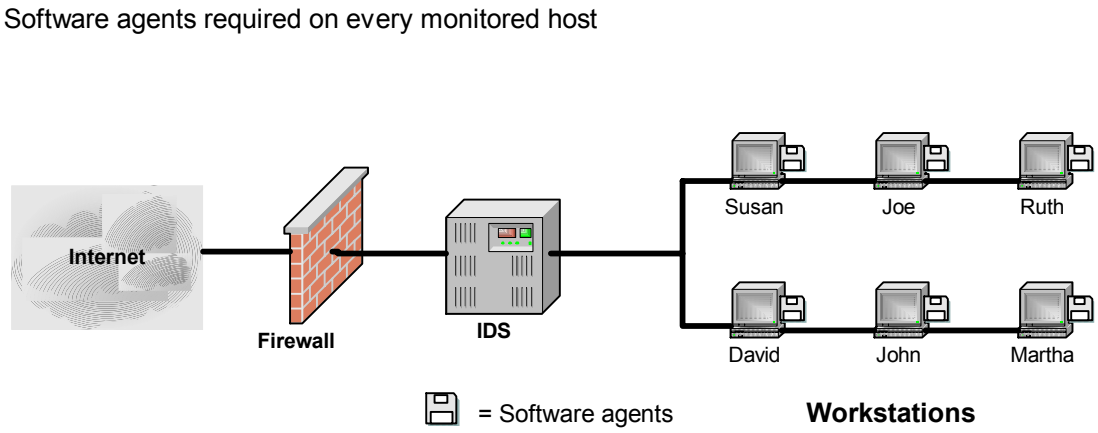

Intrusion Prevention Systems

An Intrusion Prevention System is a module added to a base Intrusion Detection System. This module provides the ability to perform specific tasks automatically. An IT administrator can define the actions to be taken by the IPS when the attack severity reaches a pre-determined threshold.

This allows an IT administrator to specify that any attack event at the denial of service (DoS) level or greater will result in the source IP address being filtered. The filter duration can be set from 15 minutes to permanently.

The advantages to Intrusion Prevention Systems are numerous:

- An attacker's ability to attack the target network can be automatically blocked any time 24x7.
- The filter duration can be specified so the attacker's IP address is not permanently blocked.
- Real-time email notification can be sent to the IT administrator.
- The attacker's Upstream Network Provider can be notified immediately when an attack occurs.

See the illustration below.

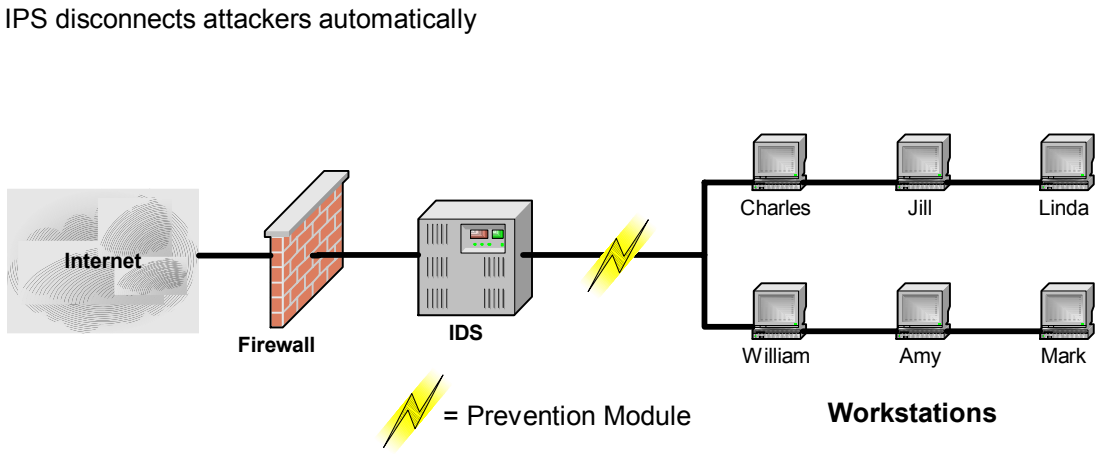

Network Hubs

Network hubs are shared media devices. This means that all computers connected to the hub can view all data passing through the hub. This is a convenient feature for an IDS which needs to analyze all data on the network.

One of the drawbacks to network hubs is total network bandwidth capacity. The network speed of a hub is shared among all network ports or connections. For example, a 10-port hub with 10 Mbps (megabits per second) throughput is limited to 10 Mbps. This hub can only support 10 Mbps even if 10 computers are connected at the same time. So if all 10 users are copying files at the same time the network hub limits each users network speed to 1 Mbps (10 users / 10Mbps = 1 Mbps).

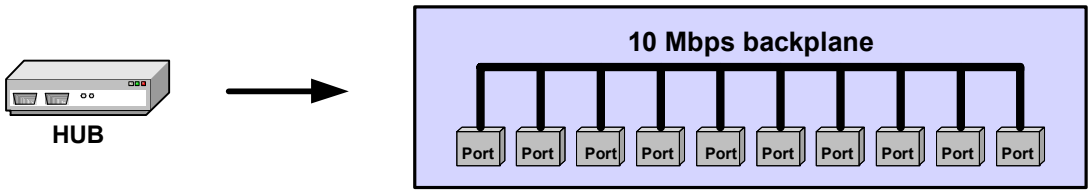

All data visible to every connection

Network Switches

Network switches are different from hubs in many respects. A switch filters and forwards packets between LAN segments. It does not send network traffic on the same "bus" for each port or connection like a network hub does.

A switch has much more total network capacity than a network hub. For example, a 10-port 100 Mbps switch has enough network capacity to support 1 Gbps total bandwidth (10 * 100Mbps = 1,000 Mbps or 1 Gbps).

A switch does not share information between all computers connected to it. So unlike a network hub, all computers connected to the switch cannot view all data that passes through the switch. This is a problem for Intrusion Detection Systems since their primary function is to analyze all data and a switch is keeping all data private.

One solution to allow an IDS to see more data on a switch is the use of port spanning. Port spanning is typically a feature found on managed switches. Port spanning allows data from one or more switch ports to be mirrored by a specified port. This allows a spanned port with an IDS connected to it to analyze data from one or more switched ports.

While switch port spanning is a useful tool to copy data to an IDS port there can be traffic saturation problems. For example, if traffic on three ports is running at 100 Mbps and all three ports are spanned to a single 100Mbps port then traffic loss will occur since the destination port can only support 100Mbps of traffic.

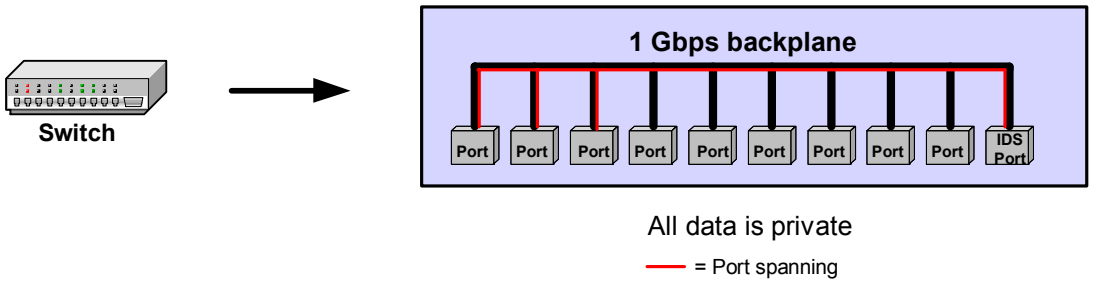

TAPS

A network TAP is a shared media device similar to a hub. A TAP is designed to allow an IDS to analyze data on networks where switches or other obstacles exist.

TAPS are designed specifically for the purpose of using Intrusion Detection Systems in a complex network.

TAPS have a number of main benefits and weaknesses which should be understood:

Benefits

- In the event of power failure the TAP does not interrupt network data flow
- The TAP does not restrict data flow on a network
- A TAP can be used for devices on dissimilar network segments

Weaknesses

- TAPS are expensive
- TAPS have network distance limitations
- Intrusion Prevention System functionality is difficult to provide when using a TAP

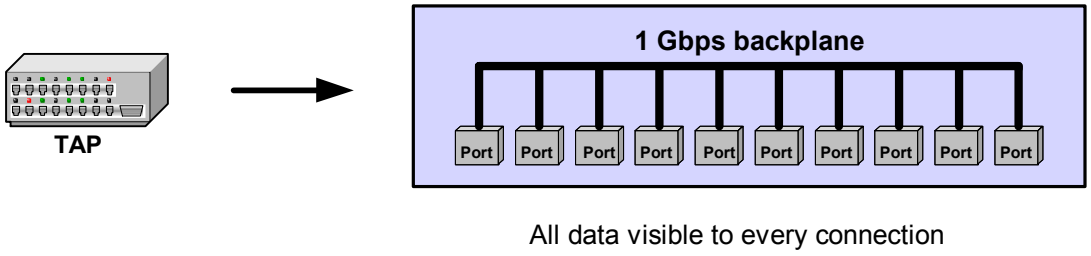

All data visible to every connection

Intrusion Detection/Prevention System Challenges

Intrusion Detection and Prevention Systems are necessary to understand and prevent network attacks that originate from the Internet or from your internal network. The need for IDS/IPS is increasing as network attacks become more sophisticated and frequent.

There are a number of challenges associated with the implementation of IDS/IPS systems in a production network. This section discusses Intrusion Detection/Prevention System implementation challenges to help IT administrators formulate sound implementation strategies.

Network Detection Zones

Intrusion Detection/Prevention Systems are placed in different types of network environments. For simplicity sake, we have identified three types of network detection zones as shown below. Each network detection zone has unique characteristics and the IDS must be able to adapt to each zone.

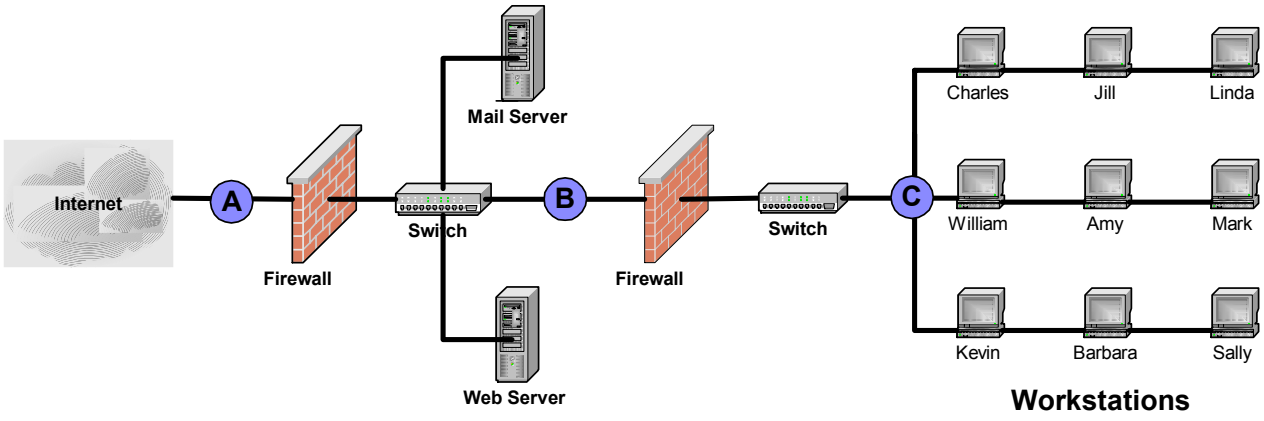

Zone A

This zone is in front of the main firewall. The main characteristic of this zone is the number of attacks logged. Frequent port scanning attempts, worm attacks and other network attacks are found in this network detection zone. The IDS must have the following characteristics to operate in this zone:

- Employ firewall protection on the external interface
- Allow logging of all attacks while offering user selectable alert notification for critical attacks
- Trigger alerts originating from both internal and external networks

Zone B

This zone is behind the main firewall so the number of attacks is dramatically lower than those experienced in Zone A. When the IDS triggers in this zone the threat is more serious in nature. IPS threshold settings may be tightened to lower or more sensitive levels in this zone.

Zone C

In this network detection zone a properly configured IDS will see fewer alerts than Zone B. The IDS and IPS threshold settings may be tightened to the lowest levels in this zone.

Network Traffic Saturation

Care must be taken when designing the implementation or placement of an Intrusion Detection/Prevention System on a network. There are three main issues to consider regarding traffic saturation and network media:

1. **Network Hub** – If a network hub is used to allow an IDS to analyze network data then the hub may restrict network data. Also, the use of a hub may render network connections too slow.
2. **Switched Media** – Since switches do not allow sharing of data between connections, port spanning may be used to allow the IDS to analyze network data. If network traffic is moderate to high then the spanned ports may overwhelm the destination IDS port and traffic loss may occur. An IT administrator should analyze network traffic before a decision is made regarding switch port spanning.
3. **TAPS** – TAPS allow traffic to flow unrestricted and they allow all data to be analyzed. The main problems are the expense of TAPS and the ability to use IPS functionality.

Frequent Updates

A significant strength of Intrusion Detection Systems is the number of data attack signatures they "watch" for. Since new attacks are discovered almost daily it is important to have the latest attack signatures downloaded daily to your IDS.

A good IDS will offer frequent updates. However, the delivery mechanism should require authentication and encryption to ensure your data has not been tampered with. Also, the IDS should perform the update function automatically so no additional burden is placed on IT staff.

False Positives

There is much discussion regarding false positives generated by Intrusion Detection Systems. False positives are annoying since they may cause an IT administrator to waste time looking for attack information, when an attack never occurred. There are a number of issues, which can lessen or remove the impact of false positives in a well-designed IDS.

False positives typically occur for low severity attack signatures such as port scans and ICMP traffic. High severity attacks such as denial of service (DoS) attacks or privilege escalation attacks are less likely to be false. This is because these attacks are typically more sophisticated and the attack pattern is typically unique.

An IDS should allow tuning so that it can be placed in any type of detection zone on a network. Using tuning parameters, the IDS sensitivity can be adjusted to provide accurate data.

Your IDS should allow an IT administrator to quickly enable or disable rule groups as well as individual rules. This allows the flexibility to keep your IDS recording the attack data you require while removing data that is unnecessary.

Advanced Intrusion Detection Systems will log all data into a database. The database notification module should contain intelligence to determine if the IT staff should be notified of an attack.

Network Integration Strategies

Integrating an IDS/IPS into your network requires careful thought and planning. Your main objectives must be considered before selecting a strategy and purchasing equipment.

You should consider your requirements and objectives before determining an IDS or IPS solution. Also, you will need to understand what "zones" you wish to place your IDS.

Here are a number of questions you should consider to help determine your needs:

1. Do you want to monitor all external attack attempts?
    a. If so, place the IDS/IPS outside or before the existing "edge" firewall.
2. Do you wish to monitor internal traffic as well as external traffic?
    a. You may wish to install the IDS/IPS on your main Internet connection.
    b. You may require multiple interfaces on your IDS/IPS or a network TAP may be required.
3. Are you required to monitor traffic for all computers on the network or only critical servers?
    a. You may wish to install the IDS/IPS on your main Internet line.
    b. You may need to monitor several network segments.
4. Do you have multiple subnets?
    a. Multiple interfaces on your IDS/IPS or a network TAP may be required.
5. Do you have multiple sites?
    a. You may require multiple network-based IDS to act as "sensors" at multiple locations.
    b. You may require a centralized management console to collect data from the sensors.
6. Is IPS functionality a requirement?
    a. You will need to determine the best location to place the IDS/IPS on your network.
    b. Use of a network TAP may be required.

In-line Intrusion Prevention

The SecurityMetrics Appliance uses inline Intrusion Prevention System technology. An Inline IPS functions like a network switch. It can be plugged into a network without any special configuration.

The Appliance analyzes all data packets flowing on the network. If a packet contains an attack signature then the Appliance will turn off the packet flow only for the attacker's network connection. The IDS will turn off any attack regardless if the attack is launched from an external or internal IP address.

The following two sections detail how the SecurityMetrics Appliance using Intrusion Prevention could benefit an organization.

External IPS Example

An external hacker is attempting to launch a Microsoft IIS Extended Unicode Directory Traversal Attack (CVE-2000-0884) at your web server. If the SecurityMetrics Appliance is placed between the firewall and your web server in your demilitarized zone (DMZ) then the Appliance will detect the attack and automatically turn off the attacker's connection into your network.

You (the IT Administrator) receive an email informing you of the action taken, how long the attacker is filtered from your site, and when the attack occurred. If you had turned on the Upstream Provider Notification IPS module, the attacker's upstream provider would also receive an email with this information so he can investigate the attack.

Internal IPS Example

An employee is doing some web research on a customer database system. As a result of the research the employee is surfing to sites he is unfamiliar with. The employee is visiting a rough looking website and notices the website is suddenly unavailable to him.

The employee contacts the IT administrator who looks in the IDS log or his email to discover the website contained malicious cross-site scripting code designed to obtain private information from the employee.

Normally an employee would have no way to know that their personal data had been "skimmed" by a cross-site scripting attack. However, this attack was automatically stopped before the employee's web browser obtained the malicious cross-site scripting request. The IPS intercepted the traffic and stopped the attack before the employee could receive the data to his browser.

IDS versus IPS Network Implementation

Historically it has been easier to implement an Intrusion Detection System into a network than an Intrusion Prevention System. This was mainly due to the fact that Intrusion Prevention Systems require insertion into the flow of traffic (two network interface connections) where an IDS simply sniffs existing traffic and doesn't need to block the data flow (one network interface connection).

SecurityMetrics Appliance uses advanced Layer 2 Network Bridging technology that acts as a switch on your network. Configuration is very simple. The Appliance only needs two valid network IP addresses and all packets will begin routing. No network configuration changes are required.

The following two sections show a sample network implementation of an IDS and IPS in simple network environments. While the SecurityMetrics Appliance is designed for use both large and small networks, a simple network is used for illustration purposes.

Sample Intrusion Detection System Network Placement

In this example we'll use a simple small office network. The IP addresses shown below are for example purposes. Your network IP address schemes may vary from the example shown below.

Network Before IDS Implementation

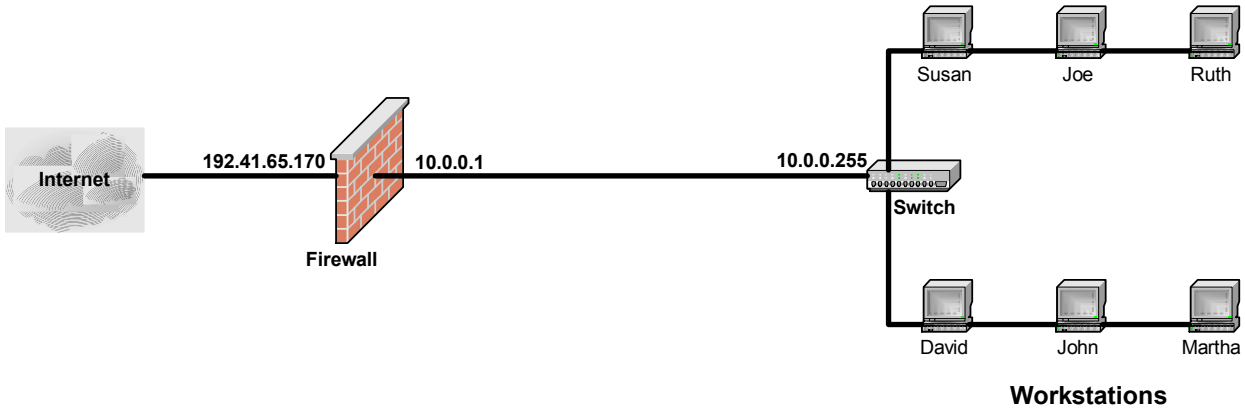

In this example our objective is to monitor internal and external Internet traffic using Intrusion Detection. This can easily be accomplished using an inexpensive hub since the Internet connection bandwidth is only 1.5Mbps. The IDS is connected to the network hub and all Internet traffic is monitored. The IP Address of the IDS can be any IP Address in the 10.0.0 subnet or it can use DHCP.

If the switch shown below was a managed switch then port mirroring could be used. In this case, the IDS would be connected directly into the switch and no hub would be required.

Network After IDS Implementation

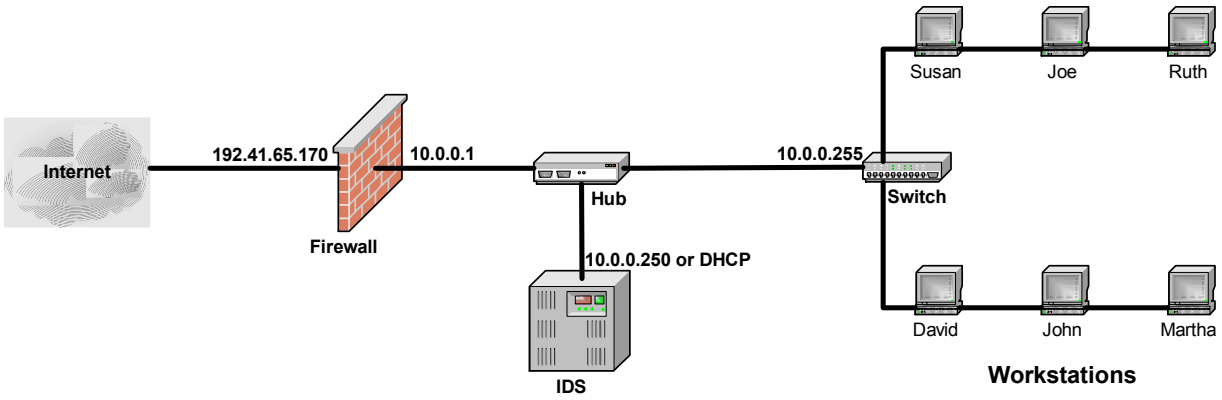

Sample Intrusion Prevention System Network Placement

In this example we'll use the same small office network. The IP addresses shown below are for example purposes. Your network IP address schemes will vary from the example shown below.

Network Before IPS Implementation

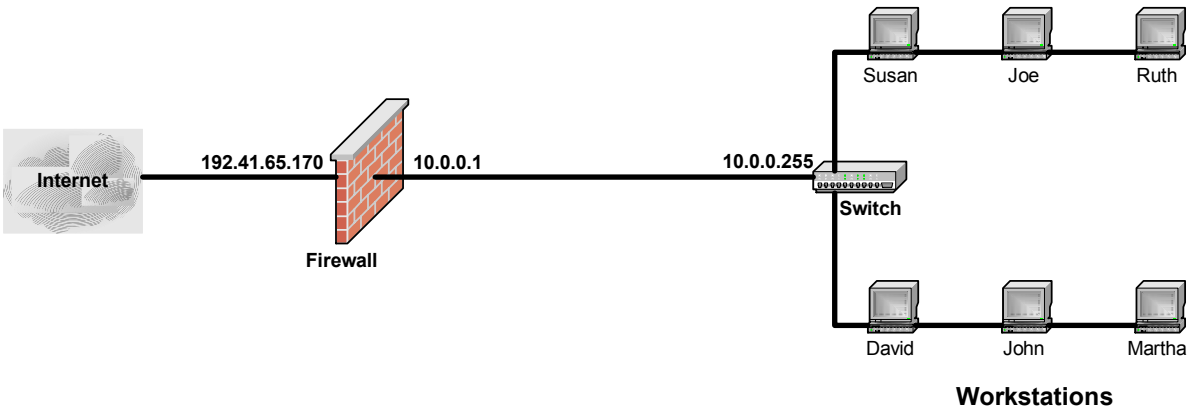

In this example our objective is to monitor internal and external Internet traffic using automatic Intrusion Prevention. This can be accomplished by placing the IPS "inline".

Since the SecurityMetrics Appliance may be configured as a Layer 2 Bridge, no configuration changes to the internal network are needed. Simply assign the Appliance to use two available network IP addresses as illustrated in the figure below.

The SecurityMetrics Appliance can be shipped to you with the IP addresses pre-configured or with DHCP enabled. You may change the IP address configuration at any time.

Network After IPS Implementation

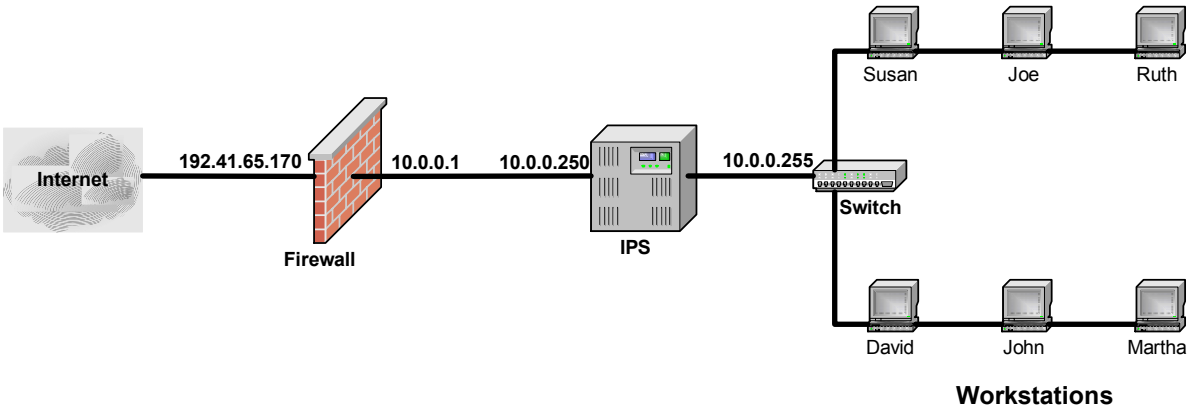

SecurityMetrics Vulnerability Assessment Technology

Now that we have a background for Intrusion Detection Systems lets briefly discuss vulnerability assessment.

Vulnerability assessment is perhaps the most ignored security technology today. It is inexpensive and deadly to IT administrators who have never used the technology for their systems. One of our favorites quotes on the use of vulnerability assessment is found on page 125 in "Linux Exposed":

> "There is one simple countermeasure that will protect you, should a hacker scan your machines with a [vulnerability assessment scanner] – scan your own systems first. Make sure to address any problems reported by the scanner, and then a scan by a hacker will give him no edge."

It is difficult to beat a good vulnerability assessment system. It is hard to recover from the use of a poor VA system. False positives can become extremely time consuming, frustrating and a waste of time.

A good vulnerability assessment system will point out holes you could never have found yourself and tell you of password problems, programming errors and basic architecture issues without the high price tag of a security consultant.

The SecurityMetrics Appliance uses five components in each vulnerability assessment scan:

1. Port scan of up to 65,535 TCP ports and the most common UDP ports
2. Vulnerability assessment of over 1,500 industry standard vulnerabilities
3. Brute force testing of 698 of the most common default username/password combinations on FTP and Telnet ports
4. Mail open-relay testing which determines if your system is being used as a mail relay
5. Website spidering two levels deep to find HTML errors or XSS code

The SecurityMetrics vulnerability assessment is the best out there. Our sales guys are happy to do a head to head comparison for you with any competitor.

Integrated IDS and Vulnerability Assessment

The importance of integrating Vulnerability Assessment with Intrusion Detection Systems is clear when you consider real world networks and tools.

A good Intrusion Detection System is sensitive. It will have a large rule database (as opposed to anomaly detection only) that will spot any potential attack. The price for a sensitive and accurate IDS is a large amount of data.

Now data is important and vital to intrusion detection. You don't want to turn off potential attack information but you also don't want to be alerted at 3:00am to find out an inexperienced hacker is attempting MSSQL attacks on your DNS server.

It is important to receive all alerts but the need for intelligence is becoming critical. SecurityMetrics Appliance integrates vulnerability assessment with its IDS. When an attack is launched the system automatically looks at the last vulnerability assessment database for the attack target. An analysis is initiated to discover if the target is vulnerable to attack. If the target is not vulnerable then no alert is sent to the administrator. If the alert is real then an alert is sent to the administrator.

The graphic below illustrates the IDS Alert Process using vulnerability assessment correlation:

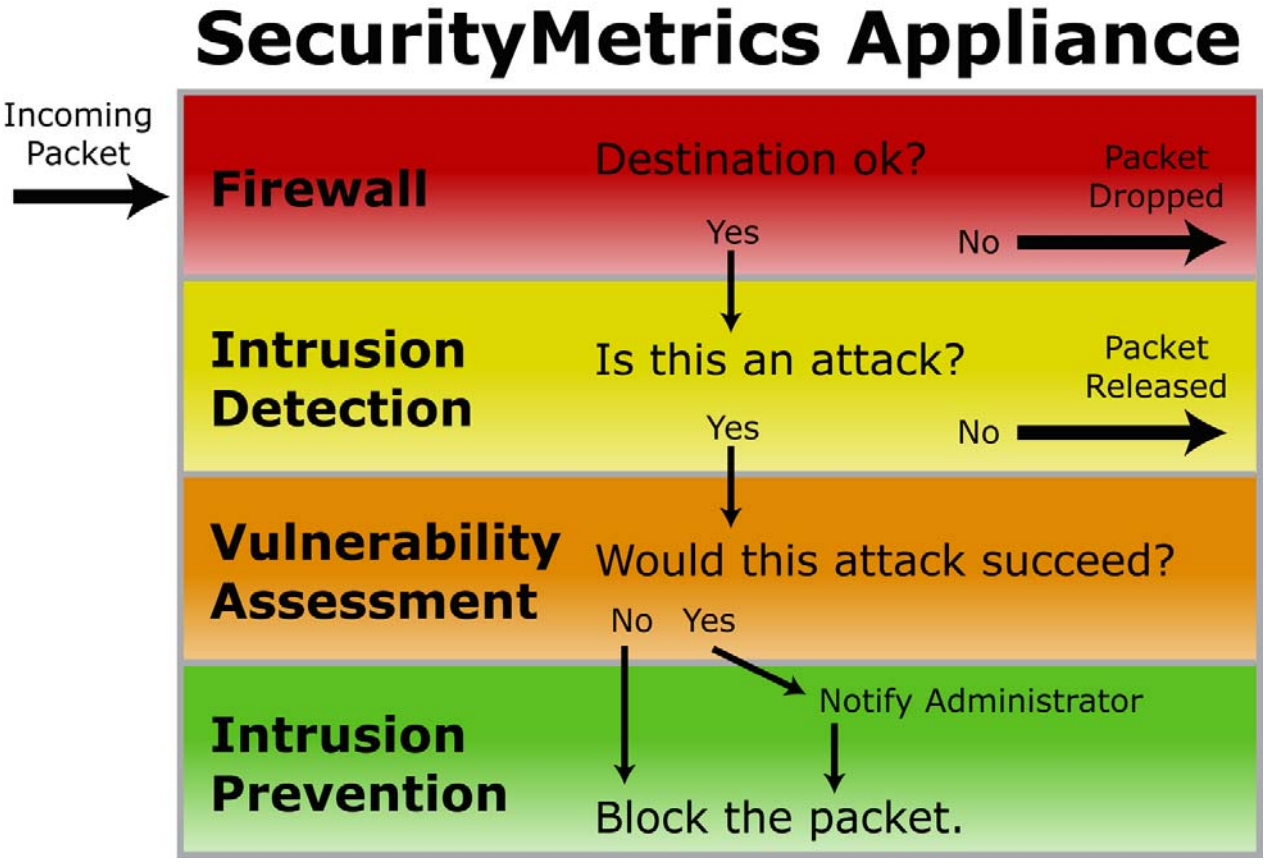

Vulnerability Assessment Reliability

Using vulnerability assessment to correlate intrusion detection reduces workload for IT administrators. However, the importance of valid vulnerability assessment becomes a necessity. If an IDS determines a target system is not vulnerable to a specific attack when it really is then the system has failed.

SecurityMetrics offers managed vulnerability assessment services (Site Certification) to its customers including online merchants, banks, credit unions, medical companies and more. Since we verify all of our vulnerabilities daily we have the cleanest vulnerability assessment database in the industry.

This is the level of reliability needed to ensure you do not receive false negatives.

Sample Vulnerability Assessment Report

The SecurityMetrics Appliance provides complete vulnerability assessment capabilities. All computers in a network can be grouped and may be scheduled to run at low usage times.

All computer test results are automatically compared to the last time they ran. If a computer experiences a risk increase then the system administrator is automatically notified by email. This reduces the need to review mounds of data and keeps vulnerability assessment and repair efficient.

One of SecurityMetrics original design specifications was concise reports. Anyone in a large organization who has to manage hundreds or thousands of computers can appreciate that a 20-page report per computer system is unmanageable.

Below is a sample report illustrating SecurityMetrics Vulnerability Assessment reports.

Executive Summary		
Test Result: **Fail**	Date: **2003-05-22**	Target IP: **10.0.0.31**
Test ID: **33**	Test Length: **1.20 Minutes**	DNS Entry: **No DNS entry**
Total Risk: **14**	Start Time: **13:55:35**	Finish Time: **13:56:48**
Full OS Description: **Windows NT 3.51 SP5, NT4 or 95/98/98SE**		

The computer **fails** because a risk of 4 or more was found. Look in the Security Vulnerabilities section below for instructions to reduce your security risk.

Attackers typically use foot printing, port scanning and security vulnerability testing to find security weaknesses on computers. This report provides information on all these categories.

Footprinting
Find public information regarding this IP, which an attacker could use to gain access: IP Information

Port Scan
Attackers use a port scan to find out what programs are running on your computer. Most programs have known security weaknesses. Disable any unnecessary programs listed below.

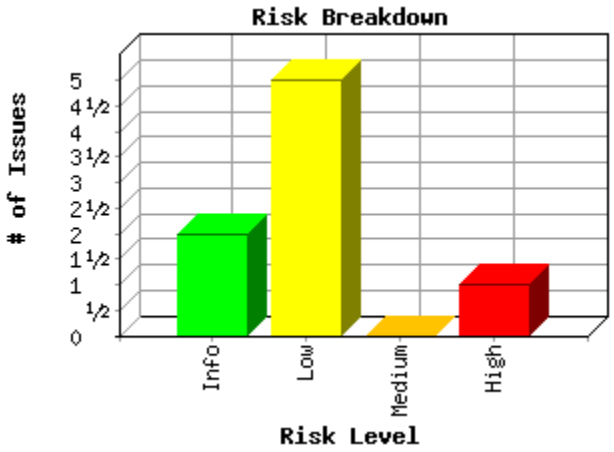

Port Scan					
Protocol	**Port**	**Program**	**Status**	**Summary**	**Turn Off**
ALL		Firewall	Absent	Your computer does not appear to be behind a firewall. We recommend installing and using a properly configured firewall.	
ICMP	Ping		Accepting	Your computer is answering ping requests. Hackers use Ping to scan the Internet to see if computers will answer. If your computer answers then a hacker will know your computer exists and your computer could become a hacker target. You should install a firewall or turn off Ping requests.	HowTo
UDP	137	netbios-ns	Open	The NetBIOS Name Service (NBNS) provides a means for hostname and address mapping on a NetBIOS-aware network. NBNS does not specify a method for authenticating communications, and as such, machines running NetBIOS	HowTo

				services are vulnerable to attacks.	
UDP	138	netbios-dgm	Open	The Netbios Datagram Service exposes characteristics of the system. A hacker can use the information exposed to break into the system. If you are a dial up user then turn off NetBeui for your dial up connection. Turn off file sharing if possible.	HowTo
TCP	139	netbios-ssn	Open	NetBIOS is a networking protocol used by Microsoft Windows to provide easy networking. If this port is open, any computer with Microsoft Windows can connect to yours and potentially use shared resources on your computer. This makes it possible for an attacker to copy, delete, or modify your data or install malicious programs on your computer.	HowTo

Security Vulnerabilities
An attacker probes your computer for weaknesses using vulnerability detection tools. The following section lists all security vulnerabilities detected on your computer.

Each vulnerability is ranked by risk on a scale of 0 to 9, with 9 being critical. The computer will fail if any vulnerability has a risk of 4 or more.

Security Vulnerabilities				
Protocol	**Port**	**Program**	**Risk**	**Summary**
tcp	0	general/tcp	**7**	The remote host has predictable TCP sequence numbers. An attacker may use this flaw to establish spoofed TCP connections to this host. **Solution** : If the remote host is running Windows, see http://www.microsoft.com/technet/security/bulletin/ms99-046.asp **Risk Factor** : High CVE : CVE-1999-0077
udp	137	netbios-ns	**3**	The remote host has the following MAC address on its adapter : 0x4a 0x42 0x20 0x20 0x20 0x20 If you do not want to allow everyone to find the NetBios name of your computer, you should filter incoming traffic to this port. **Risk Factor** : Serious Low CVE : CAN-1999-0621
icmp	0	general/icmp	**1**	The remote host answered to an ICMP_MASKREQ query and sent us its netmask (255.255.255.0) An attacker can use this information to understand how your network is set up and how the routing is done. This may help him to bypass your filters. **Solution** : reconfigure the remote host so that it does not answer to those requests. Set up filters that deny ICMP packets of type 17. **Risk Factor** : Low CVE : CAN-1999-0524
icmp	0	general/icmp	**1**	The remote host answers to an ICMP timestamp request. This allows an attacker to know the date which is set on your machine. This may help him to defeat all your time based authentication protocols. **Solution** : filter out the ICMP timestamp requests (13), and the outgoing ICMP timestamp replies (14). **Risk Factor** : Low CVE : CAN-1999-0524
tcp	0	general/tcp	**1**	The remote host uses non-random IP IDs, that is, it is possible to predict the next value of the ip_id field of the ip packets sent by this host. An attacker may use this feature to determine if the remote host sent a packet in reply to another request. This may be used for portscanning and other things. **Solution** : Contact your vendor for a patch **Risk Factor** : Low
tcp	0	general/tcp	**1**	The remote host uses non-random IP IDs, that is, it is possible to predict the next value of the ip_id field of the ip packets sent by this host. An attacker may use this feature to determine if the remote host sent a packet in reply to another request. This may be used for portscanning and other things. **Solution** : Contact your vendor for a patch **Risk Factor** : Low
tcp	0	general/tcp	**0**	Remote OS guess : Windows NT4 or 95/98/98SE CVE : CAN-1999-0454

SecurityMetrics Intrusion Detection Technology

The story is told of a retail store that was using an SCO Point of Sales system. For a number of years people had been tripping over an old computer in the back closet. Finally they decided to throw that old computer away. When they did, they found all their cash registers stopped working. The SCO Point of Sales System was the old computer in the back closet. No one knew what it was because it worked so reliably.

This low maintenance design was one of SecurityMetrics Intrusion Detection System design goals. The SecurityMetrics IDS is designed to save time for your IT staff. Here is a short list of design goals for SecurityMetrics Intrusion Detection System:

- Design all features to run automatically if chosen by the IT administrator
- Ease the installation and configuration of the Appliance
- Offer the best Intrusion Detection engine technology
- Focus on reliability by using only proven hardware and software components
- Layout the features in a well-organized interface
- Provide adequate information on each screen to remove the need for costly training
- Support the appliance at no charge so IT staff can ensure they are fully utilizing the IDS

SecurityMetrics Intrusion Detection System is comprised of a number of subsystems or modules. Each of these components performs specific features. The following illustration shows the main components:

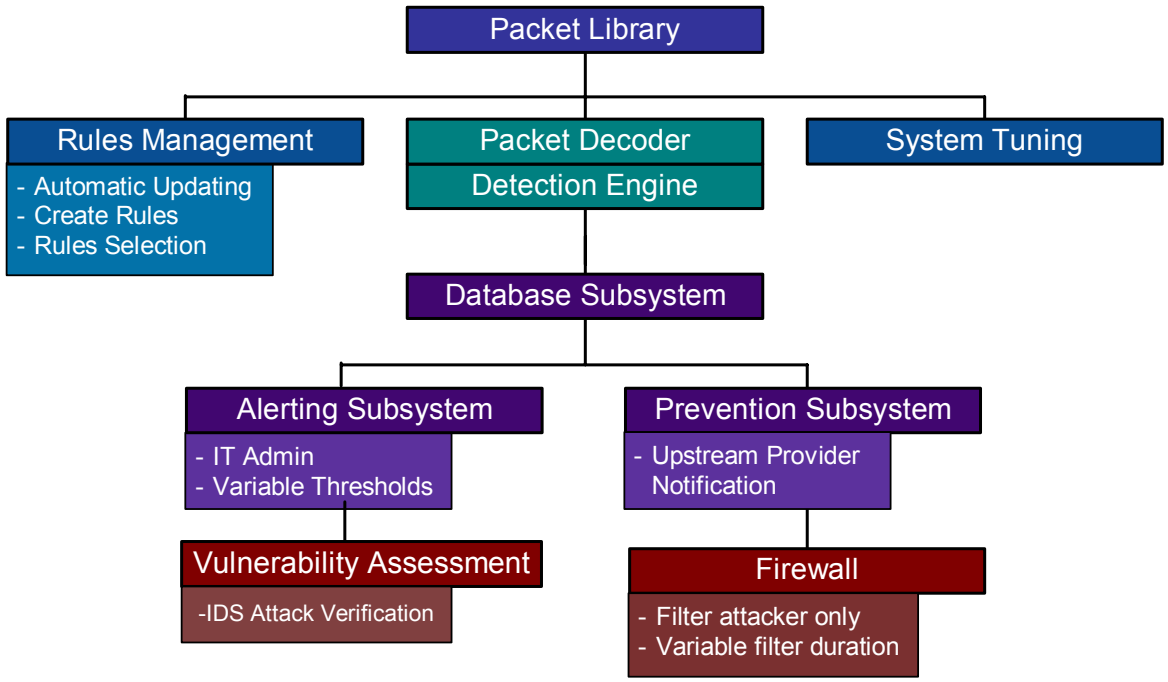

Rules Management

Rules Management is the primary tool an IT administrator uses to tune the system. SecurityMetrics provides an easy interface that shows all rule groups. Any rule group can be disabled or enabled by selecting a check box.

Inside each rule group an IT administrator can enable or disable any specific rule inside the group. In some instances a specific rule may be providing little value but may be generating false positives. For example, some port scan rules may be generating noise in the database. An IT administrator can simply turn off those rules.

The rule configuration directly affects all packet analysis for the Intrusion Detection System. Packets are flagged and logged by the IDS only if they match enabled rules.

System Tuning

There are fine tuning options that allow the removal of alert "noise" from the IDS. Certain computer relationships need to be entered into the system so they are less sensitive (at the IT Administrator's option). An example is domain name servers (DNS), which are often being accessed by the internal network and can be noisy. Also, some computers in the internal network may have activity that is potentially noisy and those computers can be flagged.

Packet Library

This library (LibPCap) allows the SecurityMetrics appliance to grab packets in their entirety. This is needed since operating systems will strip a network packet as it is processed. In an effort to ensure all packet data is analyzed a low-level packet library such as LibPCap is necessary.

Also, the default behavior for an operating system is to not look at data that isn't destined for its device. The packet library allows the appliance to analyze all packets on the network independent of the destination or port.

Packet Decoder

The packet decoder overlays data structures on the raw network packets. This sets up packet data for later analysis by the detection engine.

Detection Engine

The detection engine is optimized for performance. A busy network has a large number of packets traveling on it and each packet must be analyzed. If the detection engine is not able to handle all packets sent along the network then some data can be missed and a false negative will occur.

The engine condenses the data from the IDS rules into two sets of data: a common data group and a unique data group. The packet information is searched by the data groups recursively. The packet information is then searched again for verification. If one of the rule data groups finds a match then the detection engine triggers the actions contained in the specific rule.

Database Subsystem

Once the detection engine finds a matching rule then the data is recorded in a high performance database system. The data is stored before any further actions are taken.

The database logs all data but no actions need be taken. This allows an IT administrator to stop false positives from generating false alerts. It should be noted that through rule configuration most false positives would already be removed. But through the use of intelligent database triggers, an IT administrator can control alerting and Intrusion Prevention thresholds.

Working With SecurityMetrics

SecurityMetrics understands that the implementation of an Intrusion Detection System in your network requires configuration and setup. SecurityMetrics works with each customer to understand your network configuration, and each IDS is pre-configured to suit your needs and objectives.

SecurityMetrics will consult with you determine the best location within your network for the SecurityMetrics Appliance. This will be based on your objectives and your responses to the questions in the Network Integration Strategies of this White Paper.

Once the location within your network is determined, SecurityMetrics will pre-configure the Appliance for easy installation. SecurityMetrics support personnel are available for any last minute questions or subsequent configuration assistance.

SecurityMetrics is providing a powerful tool for maintaining security and making the process easy for installation and upkeep. Call today to get started on the path to improving your security and reducing your security workload.

SecurityMetrics offers a free 30-day evaluation for the Appliance to approved organizations. This allows you to see the benefits of the SecurityMetrics Appliance in your network environment. If you are approved, SecurityMetrics will require a conditional Purchase Order prior to shipment.

If you wish to apply for a free 30-day evaluation of the appliance call (801)724-9600 and ask for the SecurityMetrics Appliance Evaluation department.