TippingPoint
a division of 3Com

TippingPoint Best Practices Guide and Advanced Discussions

Table of Contents

Introduction

This guide is a compilation of best practices, questions and scenarios that our Systems Engineers and Account Services Specialists have encountered in the field. This guide is a living document. As features change and evolve, this guide will be updated.

Intrusion Prevention System (IPS)

Architecture and Background Information

IPS System Architecture

The main component of the IPS is the Threat Suppression Engine. The Threat Suppression Engine (TSE) reconstructs and inspects flow payloads by parsing the traffic at the application layer. As each new packet belonging to a flow arrives, the flow is re-evaluated for malicious content. The instant a flow is deemed malicious, the current packet and all subsequent packets pertaining to the flow are blocked. This ensures that the attack never reaches its destination. The TSE is a “flow” based network security engine. Each packet is identified as a member of a flow. A flow can have one or more packets. Each flow is tracked in the “connection table” on the IPS. A flow is uniquely identified by the port it was received on and its packet header information:

- IP protocol (ICMP, TCP, UDP, other)
- source IP address
- source ports (TCP or UDP)
- destination IP address
- destination ports (TCP or UDP)

Once classified, each packet is inspected by the appropriate set of protocol and application filters. The IPS filter engine combines pipelined and massively parallel processing hardware to perform simultaneous filter checks on each packet. The parallel filter processing ensures that the packet flow continues to move through the system with a bounded latency (on the order of microseconds) for the most part, independent of the number of filters that are applied. This hardware acceleration is critical in order to support massive amounts of filters without sacrificing performance. Out of the box, the IPS will identify flows in asymmetric mode – meaning the IPS only needs to see either the transmit or the receive side of a TCP connection (not both).

Core IPS Elements

The following variables, timers, and tables are core to the operation of the IPS:

- TSE Connection Table - Blocked Streams
- TSE Connection Table Timeout
- TSE Asymmetric/Symmetric Variable
- TSE Adaptive Filtering

- TSE Adaptive Aggregation

TSE Connection Table – Blocked Streams

All packets received by the IPS are identified as a member of a flow (packet stream). A flow can consist of one or more packets. All packets received that are classified as a member of a "blocked stream" are discarded. Packets will only be blocked if they match a filter that has an action set of block.

TSE Connection Table Timeout

This global timer applies to all "blocked streams" in the TSE connection table, and designates the amount of time that must elapse after a flow is marked as "blocked" before it will be "unblocked." While blocked, any incoming packets for that stream are discarded. After a flow is unblocked, the next packet for that flow is allowed but may be dropped and the flow blocked again based on the IPS filters.

For normal operations in production environments the TSE Connection Table Timer should be left at its default value (1800 seconds). However, for lab testing, this timer can be set to its minimum value (30 seconds) in order to make filter changes become more immediately apparent via seeing repetitive log updates from the same source IP address. Another way to immediately see the effects of filter changes is to "flush" the blocked stream in question from the Connection Table.

*Note – Changing filters to "unblock" a flow must be done in combination with "flushing" the blocked flow from the TSE Connection Table. Otherwise, the filter changes will not take effect for the "blocked" flow until the TSE Connection Table timer expires for that flow.

TSE Asymmetric/Symmetric Variable

Out of the box, an IPS identifies flows in asymmetric mode. Asymmetric network traffic is defined as traffic that takes a different outgoing path than incoming. It is very common for traffic to be asymmetrical in both Service Provider and larger Enterprise networks due to the nature of routing within large complex environment that have multiple ingress/egress points.

Since the bulk of the IPS filters are flow based (meaning state kept per flow versus per session), attacks are detected in either send or receive directions. The only time you must use Symmetric mode is when you are monitoring DDoS conditions regarding connection based attacks.

TSE Adaptive Filtering

On rare occurrences, the IPS system may experience extreme load conditions that may cause the device to enter "High Availability" (layer-2 forwarding) or blocking due to traffic congestion caused by filter failure or page faults. To prevent the IPS from entering HA mode, "Adaptive Filtering" disables the filter causing the possible congestion of traffic. If the Adaptive Filter Configuration (AFC) feature fires, an entry will be written in the system log.

TSE Adaptive Aggregation

Because a single packet can trigger an alert, attacks featuring large numbers of packets could potentially flood the alert mechanism causing system congestion. Adaptive Aggregation will limit the action set of particular filters that fire more than x times in the last minute. This is not to be confused with "Alert Aggregation".

"Block & Notify" and "Block & Notify & Trace" action sets are reduced to "Block" when Adaptive Aggregation triggers. The same results occur for "Permit" action sets as well. System logs will note entering and leaving this condition. Adaptive Aggregation will stop when the offending filter fires fewer than x/3 times in a minute.

Physical Connections

The IPS is placed in-line between two network elements (i.e. between 2 routers or switches).

The IPS doesn't act as a network element in the sense that it does not route traffic – it simply inspects the traffic. Because the IPS is an in-line device, the physical interfaces must match the segment in which it will be placed. Currently, the IPS only supports Ethernet interfaces, optical and copper, 10/100/1000 Mbps.

Note: For 10/100Mbps connections, it is standard practice to **DISABLE AUTO NEGOTIATION** on the IPS interfaces and the interfaces surrounding the IPS. Auto-negotiation has been known to cause problems after a router reboots. It is always best practice to hard code all endpoints and the IPS to the same speed and duplex.

For Gigabit connections, it is recommended to leave the interfaces on auto-negotiation.

Note: There is a known bug where the ports on the IPS revert to auto-negotiate after a TOS upgrade. This bug will be fixed in a future release.

Cabling Requirements

The IPS ships with the following cables:
2 AC power cable for the redundant power supplies
3-Meter Multimode LC-SC Fiber Patch Cables (4 or 8, depending on model)
Null modem cable (DB-9 FM - DB-9 FM) for (COM) port

Note: The IPS can use Lucent Connector (LC) fiber-optic cables in single-mode or multi-mode with the appropriate SFP for each cable type. The module also supports Category 5 Ethernet cable for the 10/100/1000 Ethernet connections.

Fiber-Optic Connection Guidelines

The IPS can use fiber-optic connectors. The connector type is a Small Form-Factor Pluggable (SFP) fiber optic connector that is LC-Duplex compatible. The IPS also supports the following fiber-optic media:

Multi-Mode Short Reach Fiber (MMSRF)
Single-Mode Intermediate Reach Fiber (SMIRF)
Single-Mode Long Reach Fiber (SMLRF)

ZX Fiber Interfaces

The ZX fiber interface is not officially supported by TippingPoint. Users of this type of interface should make sure that the patch cables used are certified for this type of use. Additionally, use of short patch cables may lead to intermittent failure, due to the SFP's being over-biased by the short length. This can be avoided by installing a trap to increase the attenuation so that the signal won't be too strong.

IPS Rack Clearance

The manuals state 3" around ventilation openings. On the 50/100, the air inputs are on the side with exhaust out the rear. On the 400, the air input is through the front of the unit and out the back.

IPS Deployment Considerations

When designing your IPS deployment, consider a defense-in-depth strategy where in addition to your network border, you also subdivide areas of your internal network into separate “attack domains” (also known as “security broadcast domains”); this not only contains outbreaks within your LAN, but also allows continued IPS protection if one unit is bypassed for maintenance. In most cases user traffic can pass through as many as three IPS’s before any cumulative latency is noticed.

Stacking SYN Proxies as it relates to Advanced DDoS

Stacking syn-proxy devices will increase latency and potentially cause client packets to be dropped during connection initiation. For a single syn-proxy device, the client-side connection is established first to verify it is a legitimate connection. Once that is done, the server side connection is initiated. During this time the client may start sending traffic. From the client perspective the connection is up. Until the server-side connection is established however, the syn-proxy must buffer client packets. Since buffering is not unlimited, some packets may even be dropped if the server-side connection takes too long.

If you add a second syn-proxy device, the issue is compounded from the client point of view. The client will not get a response from the server until the second syn-proxy device completes the server-side connection, which is additional latency on top of the completion of the connection with the first syn-proxy device.

Note that using more than one syn-proxy does not cause 100% failure. It may work some of the time but it most likely will increase the number of client disconnects.

DDoS performance on the 100E

The 100E is not meant to mitigate attacks that approach or exceed 60,000 SYN’s per second.

The 100E is optimal for 10Mbps worth of traffic or less (when it comes to DDoS).

10Mbps is usually plenty when the IPS is placed directly in front of an Internet facing server.

Note: Any additional performance requirements should be addressed with the TippingPoint 5000E platform.

System Administration

How to Recover the IPS SuperUser Password

You cannot recover the SuperUser password for the IPS, but you can reset it to a new value, or create a new login with SuperUser privileges.

Preparation: Make sure you are connected to the IPS serial port with serial settings of 115200 bps, 8 Data Bits, Parity None, Stop Bits 1.

Note: This procedure requires a reboot operation, which will disrupt traffic!

1. As the IPS is booting, watch for the word “Loading”, this comes up after the TippingPoint banner.
2. Type the word **mkey** within 3 seconds.
3. Specify the new security level, SuperUser login and password.

How to Reset an IPS to Factory Settings

Issue the following command from the IPS Command Line Interface (CLI).

```
debug factory-reset
```

Note: If your IPS was originally shipped with a version of the TippingPoint OS (TOS) older than v.2.1, you may need to contact TAC to perform the factory reset procedure.

Configuring the Management Port of the IPS

For enhanced security of the management port, configure the management port on the IPS to use a non-routed IP from private (RFC 1918) address space. You may also wish to do this with your SMS. Many VPNs can be configured to allow outside access to private IPs within your network if necessary.

Use the management port IP filter feature to limit access to the management port. For example, issue the following command to limit management port access to one host.

```
configure terminal host ip-filter permit ip 111.222.33.44
255.255.255.255
```

It is best practice to connect the management port of the IPS and SMS on separate physical network links outside of the in-line IPS links.

Note: The management port of the IPS is set to auto-negotiation, and cannot be changed. Make sure that the switch port where the management port is plugged into is set for auto-negotiation to prevent potential duplex mismatch issues.

Using a Non-IE Browser

Officially, only Microsoft Internet Explorer 6.0+ is supported. However, much of the functionality does work if you use another browser such as Firefox. To enable access from a browser other than IE, you need to disable the browser check on the IPS. From the CLI:

```
configure terminal server no browser-check
```

How to Turn Off SMS Management on the IPS

Issue the following command at the command line:

```
configure terminal no sms
```

You can also turn off SMS management using LSM via Configure -> NMS Management.

Subsequently, you may re-enable SMS management by issuing the following command:

```
configure terminal sms
```

System Upgrades

Remember that performing a TOS upgrade with the unit inline can disrupt those segments for as long as 15 minutes during the TOS upgrade reboot.

Always execute a 'debug disk stat' before performing a TOS upgrade. If you see any errors, abort and contact support.

Always be connected to the console port (if feasible) during a TOS upgrade so that you can watch the status and catch any errors.

Before rebooting, or troubleshooting any problems, first execute a 'configure terminal ramdisk force-sync all' from within the CLI/console. This will ensure that any log information is captured.

Always update the DV after a TOS upgrade since the TOS will include a DV from the time period in which the TOS file was generated.

Reports That Can Be Obtained via LSM

The IPS itself has basic reporting capabilities via the LSM interface.

Apart from the top level display that shows the number of attacks by severity, there are several useful displays for indicating general information such as relative amounts of TCP/UDP/ICMP as well as specific security reports.

The following are the types of reports available via the LSM.

- **Top Ten Attacks**
- **Attacks by Severity**
- **Attacks by Action**
- **Attacks by Protocol**
- **Attacks by Port**
- **Traffic Profile by Transmission Type** (Unicast, Multicast, Broadcast)
  This report can be very useful when troubleshooting – if you see a disproportionate amount of multicast and broadcast traffic, you may find that you are just seeing router control traffic, such as OSPF or EIGRP multicast HELLO traffic, RIP updates or Cisco discovery protocol (CDP).

  In one scenario, the IPS appeared to be working normally, and there were traffic of all types (unicast, multicast, broadcast). A test was conducted by logging on to an IM client with the IM filters activated, but there were no alerts, which indicated that the IPS was not in the path. By comparing some of the graphs available via the LSM, it turned out that we were connected to the second of two proxy servers, which was the backup rather than the master. Looking at the large amount of multicast traffic (strong indicator of control traffic), the large number of very small packets (typical of control traffic), and that the amount of traffic was small and only going through one port, it showed what the problem was.

- **Traffic Profile by Protocol** (TCP, UDP, ICMP)
- **Traffic Profile by Frame Size**
  According to CAIDA (http://www.caida.org/) the average packet size on the Internet is very small, in part due to the dominance of TCP traffic, which generates a large number of very small packets (e.g. SYN and SYN+ACK for the connection establishment). The mean is around 420 and median around 80-90 bytes. In a real network with a reasonable amount of HTTP and FTP, the number of large packets will be substantial. If only small packets are present, there is usually something wrong.
- **Traffic Profile by Port**

Traces and Email Notifications

When troubleshooting, use packet traces and email notifications sparingly and remember to revert back to normal after your “forensics” analysis is finished.

System Tuning

Alerting without Blocking

There are two ways to alert without blocking.

1 - The inline “Permit + Notify” mode

With this one, you decide which filters will be active, then (using SMS) sort them by action type, select all those which are set to **block** or **block+notify** and change the action to **permit+notify**. This way we do not block anything, but you see all the alerts that the IPS would generate. **Please make sure that you do not change an entire category setting to “Permit + Notify” (see note “big difference” below).**

Note: Permit + Notify should be used sparingly and with caution. Permit + Notify could have the effect of causing deep inspection on an abnormally large amount of traffic, leading to system performance degradation, Adaptive Filter Configuration and Layer 2 Fallback.

2 - The "IDS" mode.

You set up a SPAN or Mirror port on the Ethernet switch or router and have the information copied to the IPS. Only a single port on the segment of the IPS is connected to the IPS, which sees both parts of the conversation (the port must be set up to mirror both sides of the SPAN/Mirror port to the destination).

When the IPS is ready to be deployed in “Prevention” mode, the filters can be switched to "Recommended" settings.

Note: There is a **big difference** between switching from Recommended to Permit+Notify at the Category level and just changing the action of active filters from Block + Notify to Permit + Notify. The former activates all the filters in the category and sets them to notify (including ones that are disabled as part of recommended settings) while the latter just disables blocking, and only enables the recommended filters. Make sure you just flip the active filters in B+N to P+N to avoid potential false positives and excessive information.

Note: Misuse and Abuse Filters do not have a “Permit + Notify” settings. This is because the types of traffic under this category (i.e. – P2P file sharing and IM) generates a massive amount of events and would degrade the performance of the IPS. Use rate-limiting or blocking feature for this category.

Note: The TippingPoint IPS has small and fast buffers and is built to perform like a network element; not an IDS. The IPS expects to control the flow of traffic. Even though it only introduces <215us of latency, it still controls the flows. If you insert a span port into a segment and that span port bursts at a high bandwidth rate (i.e. - more than 600Mbps), the flow cannot be controlled and some of the burst traffic might get dropped. An IPS won’t buffer all the traffic and parse through the packets at its leisure as an Intrusion Detection System (IDS) will.

IPS in TAP/SPAN Mode

We do “work” in span/tap mode, but with some caveats. When a segment is put into a span mode you can not use the second port for another span (1 span per segment). When in span, if you push a block profile, we may report the events as alerts since the device knows that it isn’t blocking.

Note: Once again, we remind you that the TippingPoint IPS was designed to perform like a network element. (See note in previous section.)

Vulnerability versus Exploit Filters

If a vulnerability filter and an exploit (or multiple exploit filters) exist that act on the same vulnerability, use the vulnerability filter since it will catch the other exploits and new ones as well.

Action Sets

When creating an action set, name it something that will make sense to other users, since action sets are shared by all.

Creating an Exception to a Filter

If you discover that a connection is being blocked but you decide that the connection is legitimate, you may want to create an exception within the filter. Exceptions can be made on an individual filter by filter basis both from within LSM and SMS.

One thing that you may encounter when creating exceptions is that the connection that you created an exception for may not begin working immediately. This is likely due to the fact that the flow may still be in the TSE's blocked streams table.

In the SMS, Click on "Devices", then expand the tree of the device that you want to see the blocked flows on. Click on "Device Configuration", and then click on the "Threat Suppression Engine" tab. The "Blocked Streams" tab will display the blocked flows on that particular IPS.

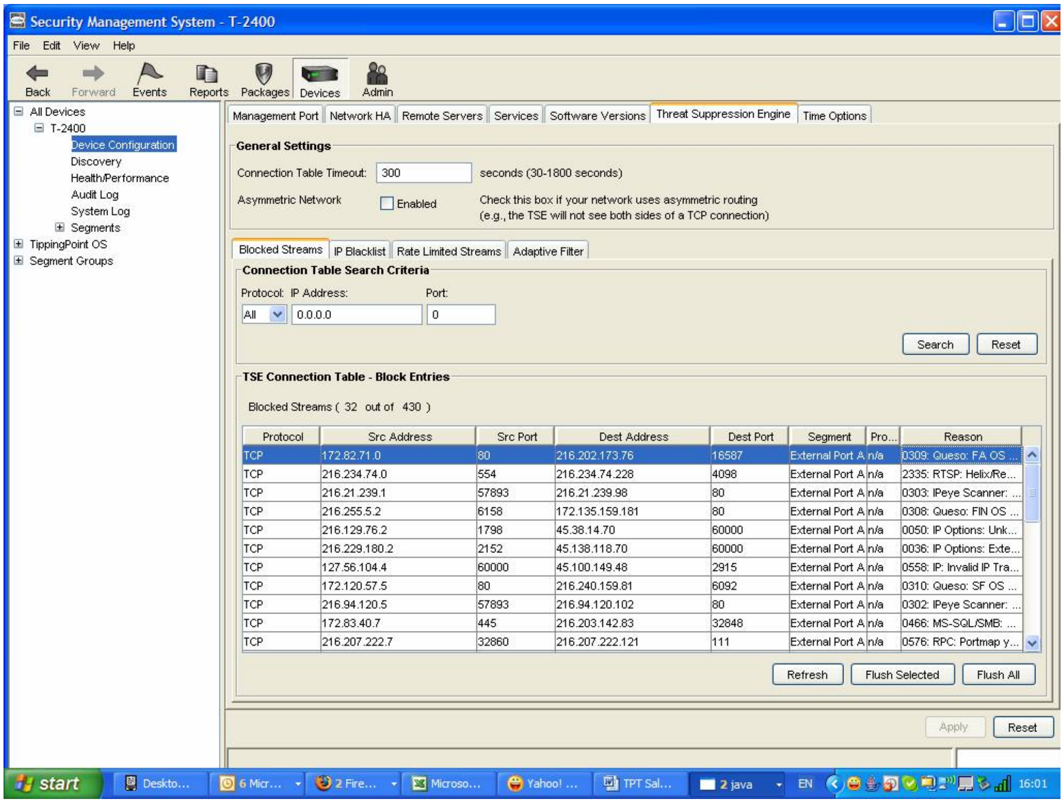

General Settings

Connection Table Timeout: 300 seconds (30-1800 seconds)

Asymmetric Network □ Enabled Check this box if your network uses asymmetric routing (e.g., the TSE will not see both sides of a TCP connection)

Connection Table Search Criteria

Protocol: All IP Address: 0.0.0.0 Port: 0

TSE Connection Table - Block Entries

Blocked Streams (32 out of 430)

Protocol	Src Address	Src Port	Dest Address	Dest Port	Segment	Pro...	Reason
TCP	172.82.71.0	80	216.202.173.76	16587	External Port A	n/a	0309: Queso: FA OS ...
TCP	216.234.74.0	554	216.234.74.228	4098	External Port A	n/a	2335: RTSP: Helix/Re...
TCP	216.21.239.1	57893	216.21.239.98	80	External Port A	n/a	0303: IPeye Scanner: ...
TCP	216.255.5.2	6158	172.135.159.181	80	External Port A	n/a	0308: Queso: FIN OS ...
TCP	216.129.76.2	1798	45.38.14.70	60000	External Port A	n/a	0050: IP Options: Unk...
TCP	216.229.180.2	2152	45.138.118.70	60000	External Port A	n/a	0036: IP Options: Exte...
TCP	127.56.104.4	60000	45.100.149.48	2915	External Port A	n/a	0558: IP: Invalid IP Tra...
TCP	172.120.57.5	80	216.240.159.81	6092	External Port A	n/a	0310: Queso: SF OS ...
TCP	216.94.120.5	57893	216.94.120.102	80	External Port A	n/a	0302: IPeye Scanner: ...
TCP	172.83.40.7	445	216.203.142.83	32848	External Port A	n/a	0466: MS-SQL/SMB: ...
TCP	216.207.222.7	32860	216.207.222.121	111	External Port A	n/a	0576: RPC: Portmap y...

In the LSM, go to Configure -> TSE Config -> Blocked Streams.

In the IPS CLI, the command "show tse connection-table" will display the same information.

When a flow is blocked and becomes listed in the TSE Block Entries table, that connection will remain blocked for at least 30 minutes. In order for the exception that you just created to take effect immediately, you must either wait for the 30 minutes to elapse, or else you must flush the flow from the TSE Block Entries table.

"Management interface under attack" and How to Avoid It

This message appears when too much of the traffic sent to the management port wasn't meant for the management IP address - too much broadcast traffic for instance.

To mitigate this condition, try putting your management interface on a less busy VLAN. You may also want to add an ACL to only allow traffic from specific subnets (has to be added via the CLI).

Disabling IP Fragmentation on the IPS

The following command exists for bypassing IP fragmentation on a per segment or src/dst IP address pair:

```
conf t tse ip-frag add|remove <sourceAddr> <destAddr> [-segment
[<slot>]:<segment>]
```

To disable fragmentation you must add a rule.

When disabling by segment it must also contain a CIDR block. For example:

```
conf t tse ip-frag any any -segment 3:1
```

If you don't specify a CIDR block, you will get an error.

Note: This command enables you to bypass traffic processing by the IPS as a workaround for problems with specific applications, such as customized or older backup applications that produce high volumes of fragmented packets. NFS and SMB traffic may also generate high volumes of fragmented packets. For each source/destination pair specified, IP fragments are passed through the device without inspection. This renders the corresponding paths not secure, so add ip-frag rules sparingly.

How Rate Shaping Works

Rate shaping – when using a rate limiting action set – is implemented in the Agere chipset, specifically the RSP (Route Switch Processor.) This uses a leaky bucket algorithm. Simply put, if a packet arrives and there is no token, then the packet is dropped. If there is a token, the packet is transmitted. Tokens are refreshed at a rate consistent with the bandwidth specified in the action set.

For system wide settings – e.g. a 200 v. a 1200 v. 2400 – the speed is an overall Agere chipset setting. The RSP will transmit packets only as fast as it is set. There are some queues/buffers so that bursts can be handled, but if these are full then the packet is dropped. Bear in mind TCP should scale back its transmission rate accordingly; UDP doesn't guarantee delivery.

Note: A common misconception among users is to assume that a single rate limiting action can be assigned to multiple filters, and that each filter will be able to utilize the full throughput specified by the rate limit action. In reality, the rate limit will be shared across all filters that have been assigned that rate limit action. For example:

1. A rate limit action called "10Mbps" is created.
2. Two filters (HTTP and FTP) are modified and are both assigned the 10Mbps rate limit action.
3. When the total bandwidth of traffic matching both filters reaches 10Mbps, that traffic will be rate limited at 10Mbps.

In the example above, if the intent was to rate limit each filter to 10Mbps *each*, the best thing to do would have been to create a specific rate limit action for each type of filter. For example, a rate limit action of “10Mbps – HTTP” and another called “10Mbps – FTP” assigned to the respective filters would have resulted in both filters becoming rate limited at 10Mbps each.

IPS – SMS Communication

The diagram below illustrates the ports and protocols used between the IPS, SMS and associated network devices.

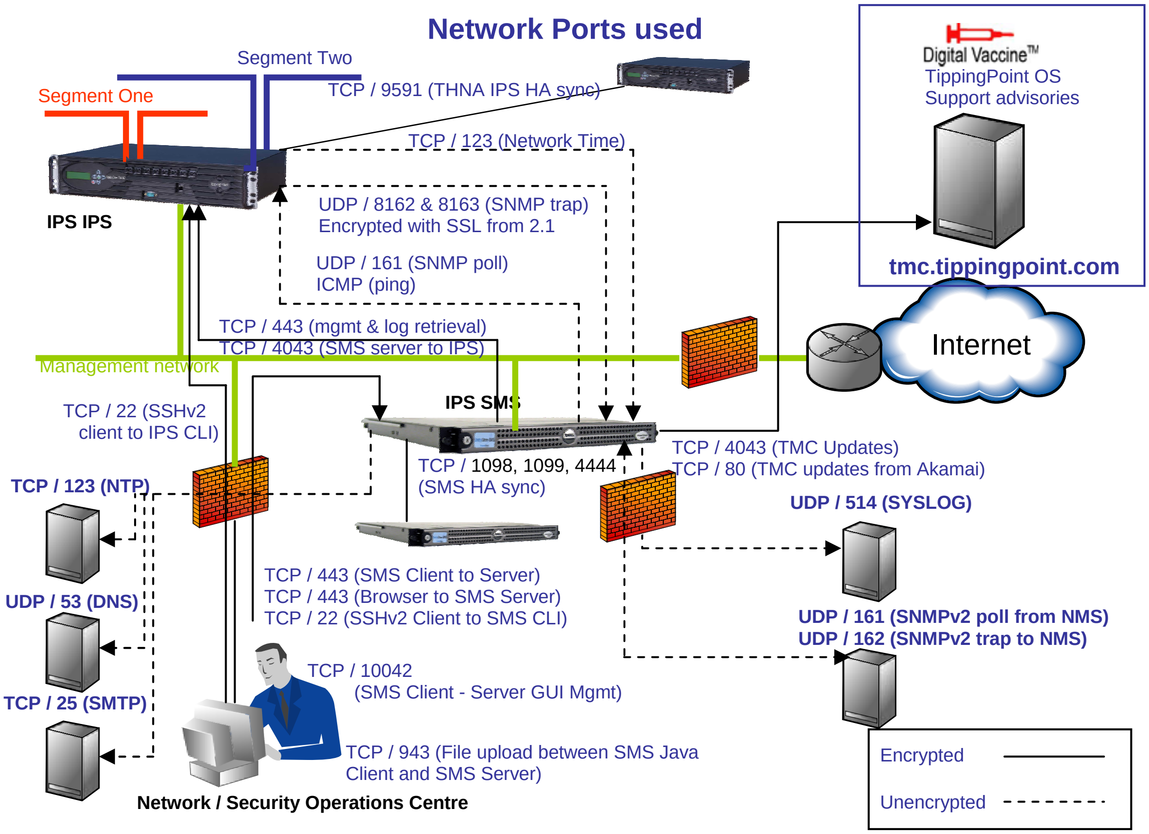
Network Ports used
Segment Two
Segment One
TCP / 9591 (THNA IPS HA sync)
TCP / 123 (Network Time)
IPS IPS
UDP / 8162 & 8163 (SNMP trap)
Encrypted with SSL from 2.1
UDP / 161 (SNMP poll)
ICMP (ping)
TCP / 443 (mgmt & log retrieval)
TCP / 4043 (SMS server to IPS)
Management network
Digital Vaccine™
TippingPoint OS
Support advisories
tmc.tippingpoint.com
Internet
TCP / 22 (SSHv2
client to IPS CLI)
IPS SMS
TCP / 4043 (TMC Updates)
TCP / 80 (TMC updates from Akamai)
TCP / 1098, 1099, 4444
(SMS HA sync)
TCP / 123 (NTP)
UDP / 514 (SYSLOG)
UDP / 53 (DNS)
TCP / 443 (SMS Client to Server)
TCP / 443 (Browser to SMS Server)
TCP / 22 (SSHv2 Client to SMS CLI)
UDP / 161 (SNMPv2 poll from NMS)
UDP / 162 (SNMPv2 trap to NMS)
TCP / 10042
(SMS Client - Server GUI Mgmt)
TCP / 25 (SMTP)
TCP / 943 (File upload between SMS Java
Client and SMS Server)
Network / Security Operations Centre
Encrypted
Unencrypted

Note that the default behavior with 2.1 is to use non-encrypted UDP/8162 and UDP/8163 to communicate between the SMS and IPS, which ensures backwards compatibility.

High Availability

TippingPoint High Availability Implementations

TippingPoint intrusion prevention systems can be implemented in a variety of high-availability scenarios. These implementations allow a range of redundancy based on budget, network architecture, and security requirements.

The IPS is designed as a piece of network hardware. It is not a server sitting off a SPAN port whose failure would not impact traffic flow. Thus, the tolerances, performance, and reliability had to meet the requirements of other network devices; such as routers, switches, and firewalls. Because the IPS does not make a path determination, due to its transparency, TippingPoint has designed other redundancy features into the line to provide even greater protection.

Intrinsic High Availability

TippingPoint IPS has been designed as a redundant device from the ground up. Internal low level system monitoring maintains operation. Upon system failure the device will fail open, bypassing internal security processing, but allowing network availability. This is termed Layer 2 Fallback and can be configured to fail open or closed on a per segment basis. The default configuration is open.

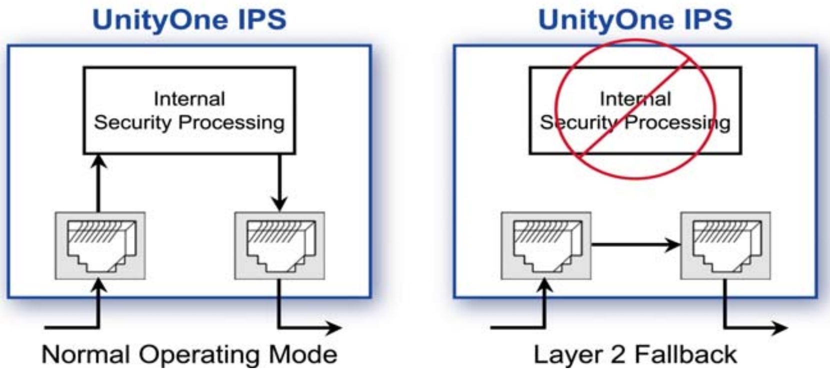

Fig. 1

This action is a powered event. Each IPS device, except for the TippingPoint 50 and 100E, comes standard with dual power supplies to ensure this operation. The appliances also have redundant power feeds.

This level of protection allows for a single device deployment while still allowing for high availability when budget prevents the deployment of redundant units, or such is not warranted. When in Layer 2 Fallback no attack protection or detection occurs but the device is not blocking traffic flow.

Zero Power HA

The Layer 2 Fallback operation requires power and allows traffic flow on both copper and fiber segments. It prevents an internal device failure from stopping all traffic flow but requires power. Zero Power HA is an external copper unit designed to pass traffic around the IPS at layer 1 upon power failure. It consists of 20 RJ45 interfaces, five groups of four interfaces, to cover five segments. These interfaces are wired pin to pin through switches that are open when receiving power and close when not powered. The unit as a whole receives power through a USB cable connected to the IPS. When the ZPHA device loses power traffic flows without being impeded but no security processing is accomplished as represented in Fig. 2. In this diagram the dark lines represent the cabling of such a deployment.

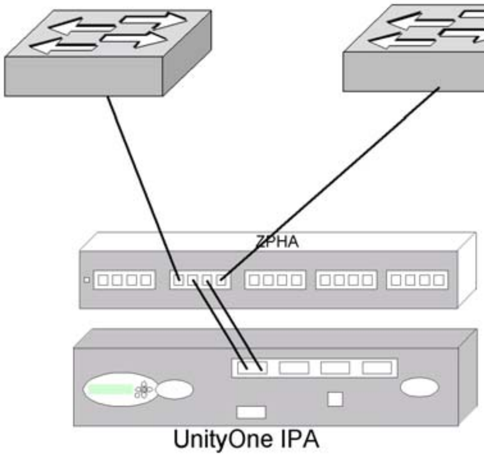

Fig. 2

This optional product was designed for deployments where a redundant unit could not be used due to budget constraints or where it was not warranted but higher redundancy requirements were in place. The use of this solution must be weighed against the security benefits gained from a redundant implementation where your network is always protected.

Note: Two ZPHA models are offered; one with fixed copper ports and another with modular slots for copper or fiber ports.

Note: Any given ZPHA cannot be shared between 2 or more IPS's, because the ZPHA is powered by a single USB cable connected to a single IPS. However, a single ZPHA can protect multiple segments.

Transparent High Availability Configurations

The final scenario is the most encompassing. The TippingPoint platforms can all be implemented in a redundant configuration, either active-active or active-passive. The units can be configured to maintain state across the two devices so that attacks are blocked on both sides. This configuration is shown in figure 3.

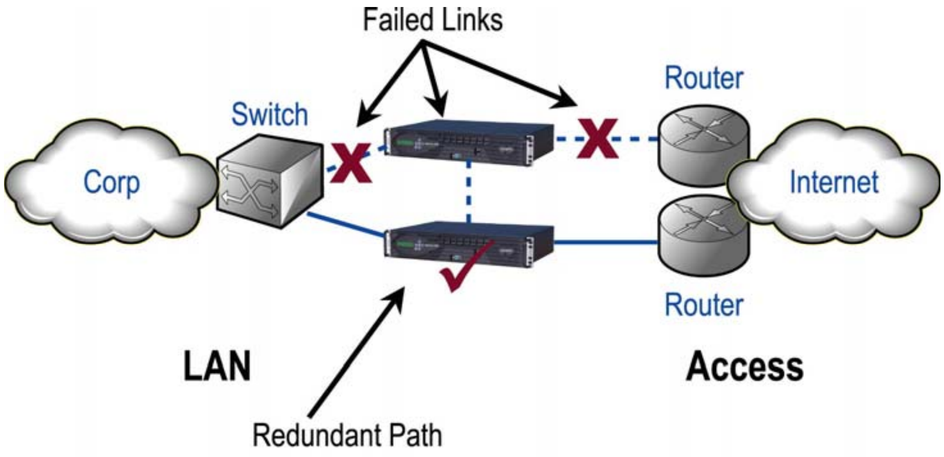

Fig. 3

Link Down Synchronization

When using these methods of path determination it is possible for the redundancy protocol to not detect the path down due to the physical connection being lost on the far side of the IPS. The link from the IPS to the router/switch/firewall is maintained. When only monitoring link status this will occur and defeats the purpose of the redundant architecture. TippingPoint has developed a work around solution to this as part of TOS 2.1.0 called Link-Down Synchronization.

Link-Down Synchronization, also called Sympathetic HA, allows you to configure the IPS to force both ports down on a segment when the device detects a link state down on one of the ports. When Link-Down Synchronization is enabled, the IPS monitors the link state for both ports on a segment. If the link goes down on either port, both ports on the segment are disabled. This functionality propagates the link state across the IPS. In the case of router A and Router B, if the link to Router A goes down, then both ports are disabled, resulting in the link to Router B going down, which Router B detects.

If the link goes down on one port, Link-Down synchronization can be configured to maintain the partner port. You can also select to have the system disable the partner requiring manual restarting to bring them both up or automatically bringing up the ports when the link comes back up.

In addition to the ability to enable Link-Down synchronization for each segment, you can change the amount of time after detecting a link is down before forcing both ports down on a segment. The default is one second and con be configured from 0 to 240 seconds. Once you enable Link-Down synchronization for a segment, monitoring of that segment begins only after link up is detected on both ports.

Note: Testing has shown that it can take up to 4 seconds for the partner link to be shut down even if the timer is set to less than 4 seconds.

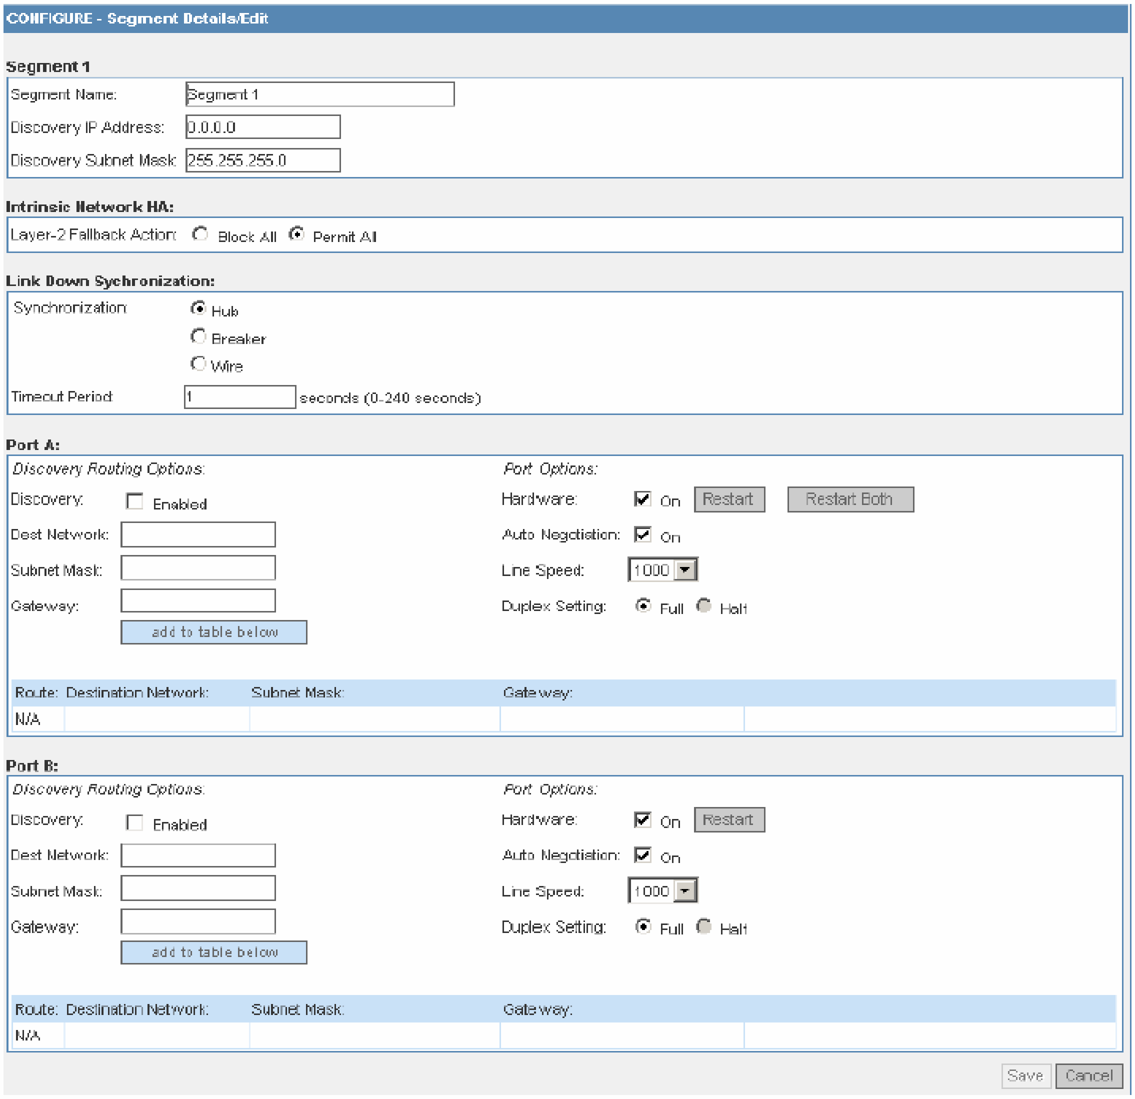

CONFIGURE - Segment Details/Edit

Segment 1

Segment Name:	Segment 1
Discovery IP Address:	0.0.0.0
Discovery Subnet Mask:	255.255.255.0

Intrinsic Network HA:

Layer-2 Fallback Action: ○ Block All ● Permit All

Link Down Sychronization:

Synchronization: ● Hub ○ Breaker ○ Wire

Timeout Period: 1 seconds (0-240 seconds)

Port A:

Discovery Routing Options:

Discovery: □ Enabled

Dest Network:

Subnet Mask:

Gateway:

Port Options:

Hardware: ☑ On

Auto Negotiation: ☑ On

Line Speed: 1000

Duplex Setting: ● Full ○ Half

Route:	Destination Network:	Subnet Mask:	Gateway:	
N/A				

Port B:

Discovery Routing Options:

Discovery: □ Enabled

Dest Network:

Subnet Mask:

Gateway:

Port Options:

Hardware: ☑ On

Auto Negotiation: ☑ On

Line Speed: 1000

Duplex Setting: ● Full ○ Half

Route:	Destination Network:	Subnet Mask:	Gateway:	
N/A				

Fig. 4 Link-Down Synchronization LSM configuration page

Example #1 Redundant router protocol

HSRP or VRRP are protocols designed to allow redundant router implementations. These protocols allow two routers to operate as one to provide redundancy in the event that one fails. Two IPS's can be implemented behind two routers running one of these protocols to provide continuous protection in the event of a failure within the network. The IPS does not participate in layer 2 or layer 3 protocols and will pass VRRP/HSRP packets transparently.

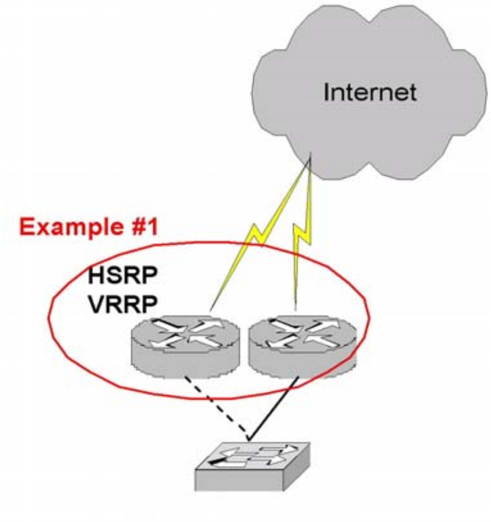

Fig. 5

Example #2 Redundant routing protocol

Often times areas of the network are connected with multiple connections to provide backup routes should one route be down. This could occur due to a WAN connection outage, a device failure, or many other scenarios. With this design a routing protocol such as RIP2, OSPF, BGP, or EIGRP can be utilized to leverage both paths for redundancy. These routing protocols operate at Layer 3 and can use one path as a primary with another as backup or can use multiple paths load sharing across them. They can even take into consideration factors such as load and capacity when making these decisions. The IPS can be implemented in these various scenarios to provide protection regardless of the implementation. The IPS can also operate without issue in an asymmetrical environment allowing a seamless integration into an existing network topology without configuration changes.

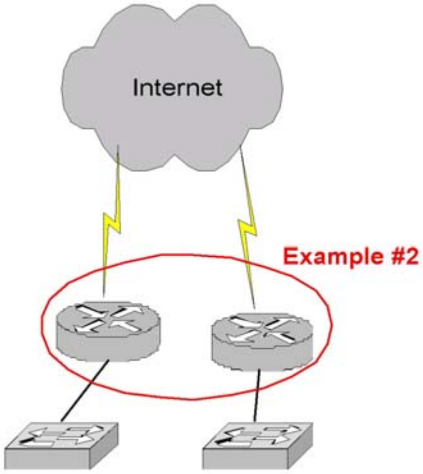

Fig. 6

Example #3 Redundant Firewall configuration

Firewalls by their very nature have become a critical element of the network border. They provide access controls at critical network interconnections. They also often times participate in routing at some level as they divide Layer 3 networks. Because of this firewall vendors have developed various redundant implementations to provide high availability and device scaling. Implementation range from a simple active-passive solution where the virtual MAC and IP addresses are shared across the two units, to complex active-active solutions where load balancers determine which sessions are sent to which firewalls. The IPS can be implemented transparently into any of these scenarios so that protection is always provided and state is maintained across the redundant architecture. With two units in place continuous protection is provided in the event of a failure within the network. In this case the firewalls or load balancers determine which path the traffic should take and the IPS transparently handles the change.

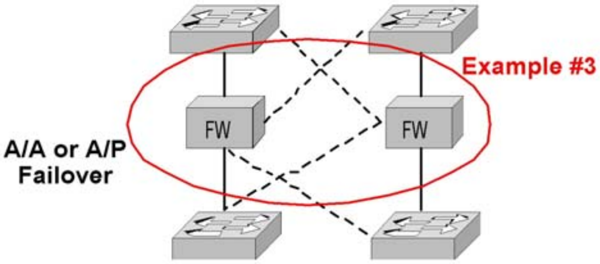

Fig. 7

Example #4 Redundant physical link with Spanning Tree

Spanning Tree Protocol is designed to allow multiple physical paths between switches running STP whereby only one path is active. When a path goes down, either due to a device failure or a cable failure, an alternate path can be used. The IPS can be implemented in this scenario with a device covering active links and a redundant device covering inactive links. These two units can maintain state across the pairs of links in the event of a failure within the network. This is a Layer 2 implementation.

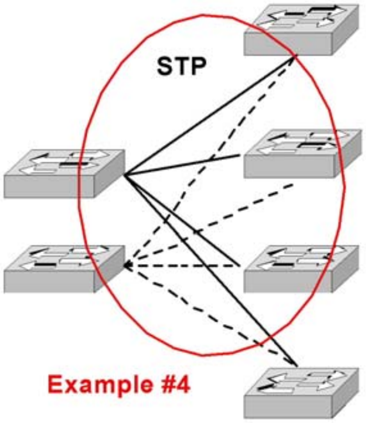

Fig. 8

Advanced Health Checks and Troubleshooting

The TippingPoint IPS is a complex device that operates impeccably in the default mode 'Recommended Settings' deployment. However, here are a few examples to help troubleshoot a problem when it occurs, or to detect a problem even before it happens. Most of these are CLI commands to be used on SSH/telnet/serial consoles, unless noted otherwise.

1. show np general statistics

This command generates an output in this format:

```
General Statistics:
-------------------
Incoming = 4892238694
Outgoing = 3668589038
Congestion = 1222036384
Deep = 262916030
Matched = 413337
Blocked = 1164690
```

Take 2 snapshots of this output over an interval you want to measure. For example, if you want to get the average traffic flowing through IPS over the last 1 minute, take 2 readings a minute apart. The difference will give you the packet throughput over the elapsed time. The numbers reset to zero on a reboot, so the actual numbers represent traffic that has passed through the IPS since the last reboot.

The important numbers to note here are:

- Congestion should be 0.1% or less of Incoming. Action: Contact TippingPoint TAC (congestion)
- Deep should be 10% or less of Incoming. Action: Contact TippingPoint TAC (too much traffic going deep)
- (Incoming minus Outgoing) should be less that 10% of Incoming. Action: Check your network. This is indicative of an unusually large percentage of attack traffic.
- Matched and Blocked are not important.

2. show health

This command generates an output in this format:

```
Temperature :
        Current: 49 degrees (C)
        Health: Normal

Memory :
        Current: 41 percent in use
        Health: Normal

Disk Partitions:
Partition       total(Mbytes) Used(Mbytes) % Used Health
----------      --------------- -------------- --------   ---------
```

```
/boot          1,003        25     2      Normal
/opt           1,002        26     2      Normal
/usr           1,002        21     2      Normal
/log           25,583       1      0      Normal
```

Watch out for:

- Temperature should normally be below 65 degrees C
- Memory should not keep climbing up steadily. If it approaches 80%, contact Tippingpoint TAC.
- Disk Partitions are almost never a cause of concern.

3. debug info ticks

This command generates an output in this format:

```
23%
```

Watch out for:

- This counter should be interpreted in conjunction with "Congestion" from "Show np general statistics."
- Ideally "debug info ticks" should be less than 90% at all times. However, a value of 100% together with congestion indicates overload.
- Action: You should turn filters to block (instead of permit+notify); or to Block only (from block+notify), turn off packet tracing, or reduce the number of overrides you have selected. If this does not work, contact TippingPoint TAC.

4. show np protocol-mix

This command generates an output in this format:

```
                   Packets     Bytes
================== ========    ========
EthType:
     ARP           585         37480
     IP            5330416048  4325965293694
     Other         52940       11738211
IpVersion:
     IPv4          5330416048  4325965293694
     IPv6          0           0
     Other         0           0
IpProtocol:
     TCP           5002822158  4268528147306
     UDP           309643586   54707082666
     ICMP          7769994     733554523
     Other         10010892    1824106533
```

Again, like "show np general statistics" these counters are reset to zero at reboot. You may want to take a difference (delta) based on 2 reading taken a minute apart.

Things to watch out for:

- Watch out for protocol balance – i.e. the ratio of TCP, UDP and ICMP on your network. Usually the ratios are 90% TCP, 9% UDP, 1% ICMP.
- Of course, each network is unique and you should baseline what is 'normal' for your network.
- If the ratios digress too much, it could indicate an attack currently going on in your network. You could also use "Traffic Thresholds" feature of the IPS for this monitoring.

5. show interface ethernet

This command generates an output in this format. There is a stanza for each physical port (An inline segment consists of 2 physical ports).

Example output:

```
Slot/Port                3/1
        Type             Ethernet
        MTU              1500
        Link             down(2)
        Speed            1000
        Duplex           Half(2)
        RX Unicast Pkts  2842395549
        RX Multicast Pkts 32249
        RX Broadcast Pkts 541
        RX Error Pkts    1
        RX Discards      0
        RX Unknown Protocols 0
        RX Total Pkts    2842428339
        TX Unicast Pkts  1654020637
        TX Multicast Pkts 15302
        TX Broadcast Pkts 310
        TX Total Pkts    1654036249

2004-08-10 15:24:31 <WARN> [DRV] Port:3/5 duplex(Half) does NOT match
Port:3/6 duplex(Full)
```

Watch out for:

- Rx Error Pkts should not steadily increase and should =0. If it does increase, contact TippingPoint TAC.
- Also check link negotiation issues, like half duplex when you should have full-duplex, line speed, etc.
- RX Error Pkts and Discards are not incrementing as you run the command repeatedly.
- The two ports on a segment should also match each other in speed and duplex. If not you may see problems and a system log message indicating this condition will be written.

6. show np rule-stats

Below is a sample output.

```
EXAMPLE OUTPUT
===============
device04# show np rule-stats
  Filter      Flows  Success  % Total  % Success
    2021     706346        0       39       0.00
    2438     407462        0       22       0.00
    2434     139529        0        7       0.00
    2769      73777        0        4       0.00
    3210      69204        0        3       0.00
    3003      67429        0        3       0.00
    2667      53714        0        3       0.00
    2402      30275    10745        1      35.49
    3174      30254        0        1       0.00
    2926      25946        0        1       0.00
    1612      17463        0        0       0.00
    1627      17463        0        0       0.00
    3236      17463        0        0       0.00
    2435      17370        0        0       0.00
    2913      10725        0        0       0.00
    2285       7697        0        0       0.00
    2419       7692        0        0       0.00
    2441       7692        0        0       0.00
    2445       7692        0        0       0.00
    2443       7692        0        0       0.00
Total of 1774958 flows
```

What does this all mean? The basic idea of the command is to give you a view into what effect a filter has on the device. The first column tells you the filter – the first one listed is 2021. Run another command to get the name:

```
device04# show filter 2021
2021: HTTP: IE Frame Flood Exploit
1 instance found
Any segment    State: Enabled    AFC: Enabled    Category: Recommended
```

The second column is "Flows". This is the number of flows that have come to Tier 3 (the processor). The third column is the Success number. In this case we received 706,346 flows and not one of them matched filter 2021. The "% Total" is the ratio of the filter’s flows to the total number of flows that went to Tier 3. The “% Success” is the number of Successes divided by the number of Flows.

If the filter does have any success, it would be prudent to leave it on. If it doesn't match – i.e. 0% success – then some performance might be regained by turning that particular filter off.

Note: Care should be taken when disabling filters. Read the filter description to make sure that your systems are patched for that particular vulnerability. If your systems have been patched, it should be safe to disable that filter.

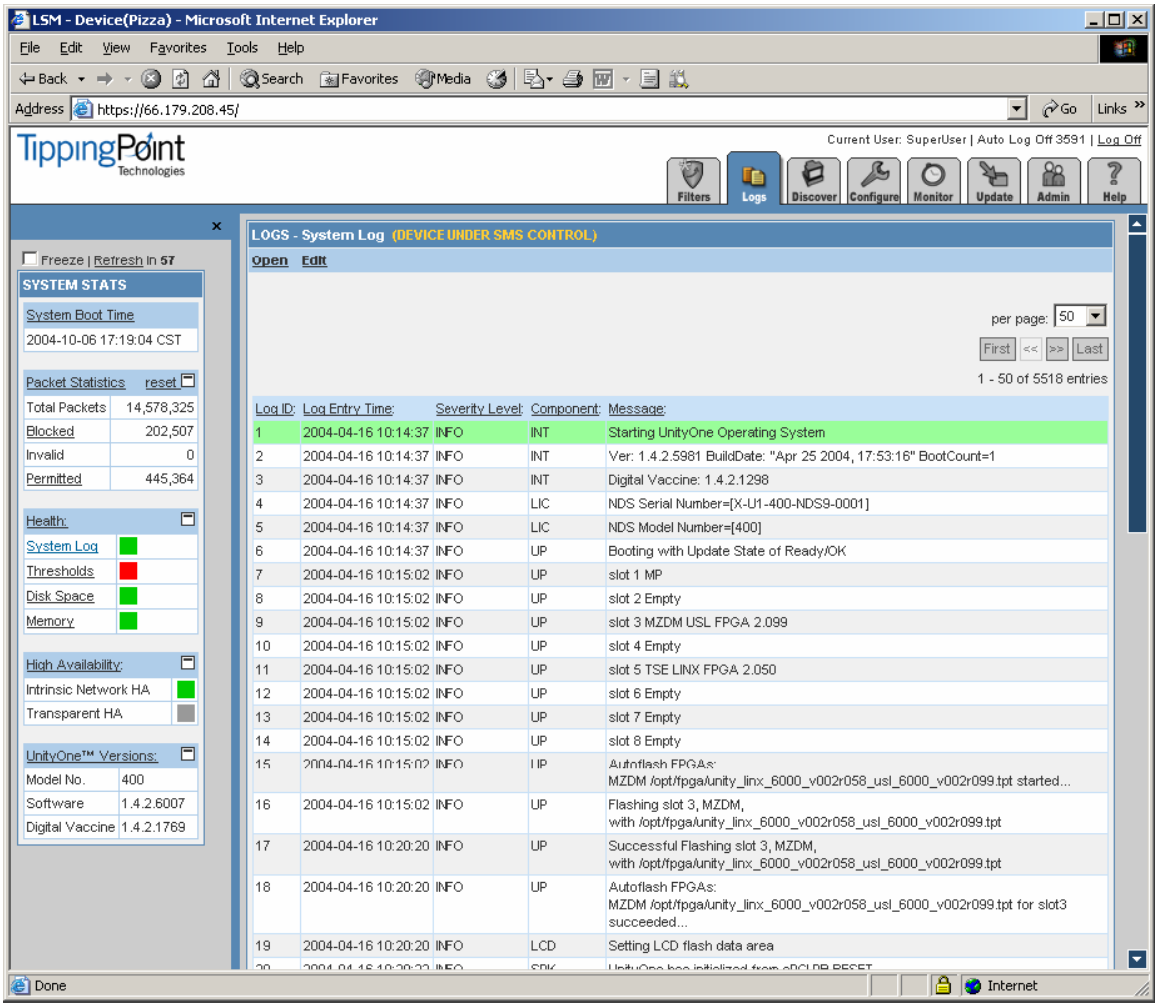
LSM - Device(Pizza) - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address https://66.179.208.45/
Go Links
TippingPoint Technologies
Current User: SuperUser | Auto Log Off 3591 | Log Off
Filters
Logs
Discover
Configure
Monitor
Update
Admin
Help
Done
Internet

□ Freeze | Refresh in 57

SYSTEM STATS

System Boot Time
2004-10-06 17:19:04 CST

Packet Statistics	reset
Total Packets	14,578,325
Blocked	202,507
Invalid	0
Permitted	445,364

Health:	
System Log	🟩
Thresholds	🟥
Disk Space	🟩
Memory	🟩

High Availability:	
Intrinsic Network HA	🟩
Transparent HA	▨

UnityOne™ Versions:	
Model No.	400
Software	1.4.2.6007
Digital Vaccine	1.4.2.1769

LOGS - System Log (DEVICE UNDER SMS CONTROL)

Open Edit

per page: 50

First << >> Last

1 - 50 of 5518 entries

Log ID:	Log Entry Time:	Severity Level:	Component:	Message:
1	2004-04-16 10:14:37	INFO	INT	Starting UnityOne Operating System
2	2004-04-16 10:14:37	INFO	INT	Ver: 1.4.2.5981 BuildDate: "Apr 25 2004, 17:53:16" BootCount=1
3	2004-04-16 10:14:37	INFO	INT	Digital Vaccine: 1.4.2.1298
4	2004-04-16 10:14:37	INFO	LIC	NDS Serial Number=[X-U1-400-NDS9-0001]
5	2004-04-16 10:14:37	INFO	LIC	NDS Model Number=[400]
6	2004-04-16 10:14:37	INFO	UP	Booting with Update State of Ready/OK
7	2004-04-16 10:15:02	INFO	UP	slot 1 MP
8	2004-04-16 10:15:02	INFO	UP	slot 2 Empty
9	2004-04-16 10:15:02	INFO	UP	slot 3 MZDM USL FPGA 2.099
10	2004-04-16 10:15:02	INFO	UP	slot 4 Empty
11	2004-04-16 10:15:02	INFO	UP	slot 5 TSE LINX FPGA 2.050
12	2004-04-16 10:15:02	INFO	UP	slot 6 Empty
13	2004-04-16 10:15:02	INFO	UP	slot 7 Empty
14	2004-04-16 10:15:02	INFO	UP	slot 8 Empty
15	2004-04-16 10:15:02	INFO	UP	Autoflash FPGAs: MZDM /opt/fpga/unity_linx_6000_v002r058_usl_6000_v002r099.tpt started...
16	2004-04-16 10:15:02	INFO	UP	Flashing slot 3, MZDM, with /opt/fpga/unity_linx_6000_v002r058_usl_6000_v002r099.tpt
17	2004-04-16 10:20:20	INFO	UP	Successful Flashing slot 3, MZDM, with /opt/fpga/unity_linx_6000_v002r058_usl_6000_v002r099.tpt
18	2004-04-16 10:20:20	INFO	UP	Autoflash FPGAs: MZDM /opt/fpga/unity_linx_6000_v002r058_usl_6000_v002r099.tpt for slot3 succeeded...
19	2004-04-16 10:20:20	INFO	LCD	Setting LCD flash data area
20	2004-04-16 10:20:22	INFO	SPK	UnityOne has initialized from cPCI PB RESET

Search block log on IPS to see if there are blocks from or to an IP:

show log block -match <ip address>

```
2004-10-22 09:50:41 <INFO> [BLK] Block v3 2 [00000002-0002-0002-0002-
000000001473] 4 00000001-0001-0001-0001-000000001473 udp
172.5.60.11:4308172.5.60.92:14341 3:1 1098456601 0430384304 1 pt0 0 00
```

Block log can be searched in LSM by going into that log, selecting Edit -> Search.

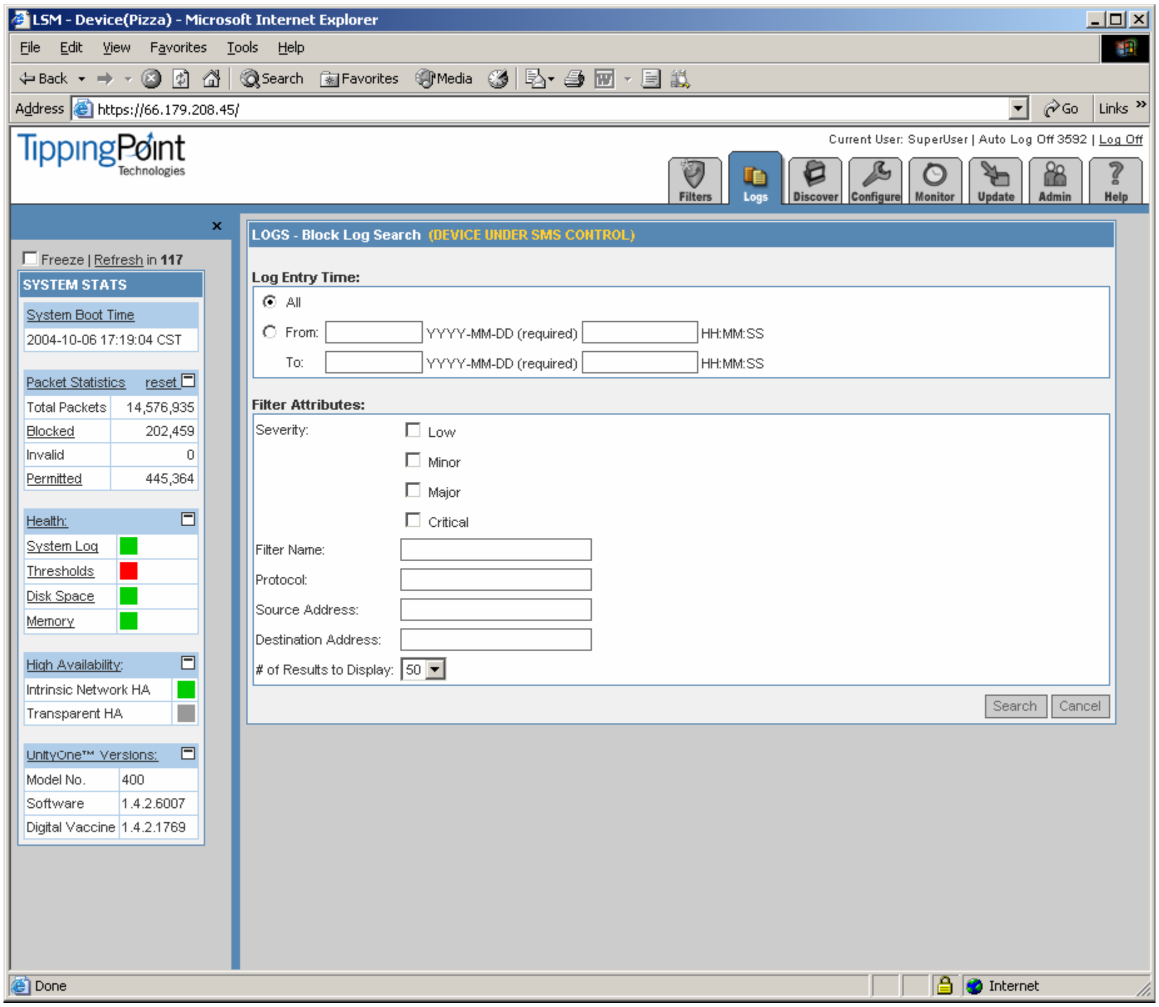
LSM - Device(Pizza) - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address https://66.179.208.45/
Go Links
TippingPoint Technologies
Filters Logs Discover Configure Monitor Update Admin Help
Done
Internet

Current User: SuperUser | Auto Log Off 3592 | Log Off

□ Freeze | Refresh in 117

SYSTEM STATS

System Boot Time
2004-10-06 17:19:04 CST

Packet Statistics	reset
Total Packets	14,576,935
Blocked	202,459
Invalid	0
Permitted	445,364

Health:	
System Log	🟩
Thresholds	🟥
Disk Space	🟩
Memory	🟩

High Availability:	
Intrinsic Network HA	🟩
Transparent HA	▨

UnityOne™ Versions:	
Model No.	400
Software	1.4.2.6007
Digital Vaccine	1.4.2.1769

LOGS - Block Log Search (DEVICE UNDER SMS CONTROL)

Log Entry Time:

● All

○ From: ______ YYYY-MM-DD (required) ______ HH:MM:SS

To: ______ YYYY-MM-DD (required) ______ HH:MM:SS

Filter Attributes:

Severity: □ Low □ Minor □ Major □ Critical

Filter Name: ______

Protocol: ______

Source Address: ______

Destination Address: ______

# of Results to Display: 50

Search Cancel

All events can be searched in SMS event viewer as well.

View blocked streams table to see if there are blocks from/to IP that are not logged:

show tse connection-table blocks

```
Blocked Connections

Protocol Src IP          Src Port        Dest IP         Dest Port
-------- --------------- --------        --------------- ---------
UDP      24.16.56.147 1214               172.236.117.143 2049
UDP      128.57.187.119 1924             172.76.244.121 1214
UDP      24.5.36.223 1214                172.110.167.120 1214
UDP      172.254.37.60 4308              172.254.37.141 1434
UDP      66.229.154.220 1662             172.175.125.67 1214
UDP      24.75.201.202 1214              172.253.207.43 1214

6 of 6 blocked connections shown.
```

Blocked Streams Table can be searched in the LSM at Configure tab -> Open -> TSE Config -> Blocked Streams

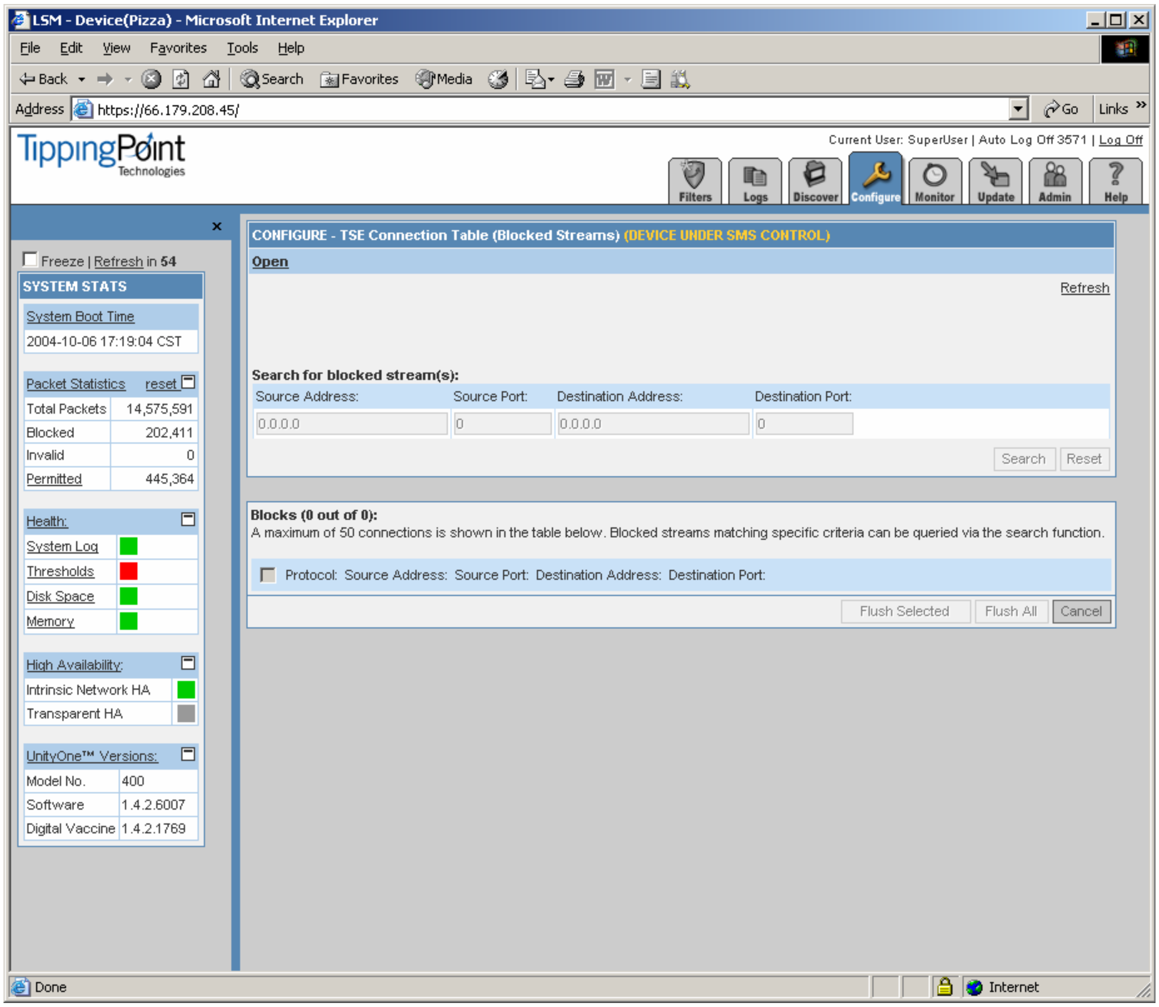
LSM - Device(Pizza) - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address
Go Links
TippingPoint Technologies
Filters Logs Discover Configure Monitor Update Admin Help
Done
Internet

Address: https://66.179.208.45/

Current User: SuperUser | Auto Log Off 3571 | Log Off

□ Freeze | Refresh in 54

SYSTEM STATS

System Boot Time
2004-10-06 17:19:04 CST

Packet Statistics	reset
Total Packets	14,575,591
Blocked	202,411
Invalid	0
Permitted	445,364

Health:	
System Log	🟩
Thresholds	🟥
Disk Space	🟩
Memory	🟩

High Availability:	
Intrinsic Network HA	🟩
Transparent HA	▨

UnityOne™ Versions:	
Model No.	400
Software	1.4.2.6007
Digital Vaccine	1.4.2.1769

CONFIGURE - TSE Connection Table (Blocked Streams) (DEVICE UNDER SMS CONTROL)

Open

Refresh

Search for blocked stream(s):

Source Address:	Source Port:	Destination Address:	Destination Port:
0.0.0.0	0	0.0.0.0	0

Search Reset

Blocks (0 out of 0):
A maximum of 50 connections is shown in the table below. Blocked streams matching specific criteria can be queried via the search function.

□ Protocol: Source Address: Source Port: Destination Address: Destination Port:

Flush Selected Flush All Cancel

This can also be searched on SMS under the device configuration of that device.

Place unit in Layer 2 Fallback

- This will remove the unit from the equation, thus passing all traffic through un-inspected.
  In CLI: **high-availability force <fallback> <normal>**
- In LSM click High Availability on left stats pane, select Layer 2 Fallback, then Apply.

Physical interface problems will still exist if present so interfaces should be checked for incrementing RX errors.

7. show log sys (LSM/SMS can be used for this)

Search system log for alerts:

```
show log system -logLevel WARN
show log system -logLevel ERR
show log system -logLevel CRIT
```

System log can be viewed in logs section of LSM and in SMS under that device.

Watch out for:

- AFC (Adaptive Filter Configuration) messages in the logs. Action: Contact TippingPoint TAC.
- Adaptive Aggregation means that a filter is firing too often. Try to debug your network to fix the problem, or set the filter to block only (instead of block+notify).
- Any CRIT or WARN messages.

Security Management System (SMS)

User Administration

Recovering a Lost SMS SuperUser Password

Preparation: You will need to use a VGA console and a PS/2 keyboard (no mouse required). A serial connection is not sufficient.

1. Reboot the SMS
2. When the LILO prompt appears, press the **TAB** key
3. Type **Recover**

The SMS will run a recovery script that will ask for the new SuperUser password.

4. When the SMS is ready to login, specify;
Login: **SuperUser**
Password: **SERIAL-NUMBER-OF-THE-SMS (Press ALT-F12 to see this)**
Press Alt-f1 to get back to the main login screen.
Once in, you can change the password via the "getpasswd" command.

SuperUser versus Administrator/Operator

SuperUser accounts should be assigned sparingly, and only to those users that need to be able to upgrade the device, view audit logs and push DV's (these last three functions can only be performed by users who are "SuperUsers").

It is considered good practice to apply Administrator and/or Operator roles to your users so you can lock down control of certain devices, profiles and segment groups in order to give your users more granular control. This prohibits users from stepping on one another while using the SMS.

RADIUS Integration Tips

- You still need to have a user account and authorization (ACL) information in the SMS database that matches the user in RADIUS. RADIUS is only used to validate the user's authentication (username/password) and all authorization information (security level, profile access, device access, segment group access) is stored in SMS.
- If the user has a password defined in the SMS, that password is ignored and never checked against the one presented by the user.
- New users created after setting up RADIUS authentication will not allow a password to be entered.
- If RADIUS goes down, there is one designated "fallback user" (defaults to the SuperUser account created during OBE, but can be set to another via CLI).
- The User Table changes when switching from RADIUS to non-RADIUS.

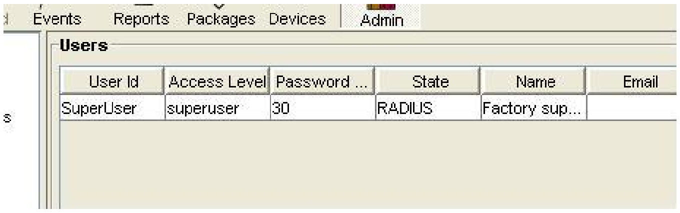

User Id	Access Level	Password ...	State	Name	Email
SuperUser	superuser	30	RADIUS	Factory sup...	

System Administration

Resetting the SMS to Factory Defaults

Run the following sequence of commands to reset an SMS to its original factory settings.

```
set db.reinit
set pwd.service-enable=1 (wait for completion)
set repos.reset=1 (wait for completion)
shutdown (will restart SMS with setup wizard)
```

After the reboot, you may need to run the "setup" command to initiate the setup wizard.

How to Enable Pinging of the SMS

Connect to the SMS CLI using a secure shell connection.

Type the following command:

```
 set svc.ping-enable=1
```

Lock Down the SMS Webpage

Lock down access to the Web Interface on the SMS. This will keep unauthorized users from being able to get the client, view reports, download backups, etc. Edit -> Preferences -> Security -> Require Login for Web Access

Protecting Reports

Use password authentication when creating reports to lock down access to those reports. When scheduling a report, under the Permissions and Remote Export, do not check the boxes labeled: “Allow anonymous users to view results” or “Allow all SMS users to view results.”

Number of IPS’s that an SMS Can Manage by Default

Licensing

Pre 1st May 2005
An SMS license covered 5 IPS’s, with each new license covering 5 new IPS’s.

Post 1st May 2005
An SMS license covers 25 IPS’s, with each new license covering 25 new IPS’s. Each new group of 5 IPS’s requires an additional license key to be purchased.

IPS and SMS in Different Time Zones

The SMS client uses the local time zone for displaying date/time. We are aware of some areas that are not correctly handling the time zone and plan to fix in 2.5.

Bandwidth Consumption between the IPS and SMS

We have observed traffic consumption of anywhere from 10-15K per IPS to SMS.

In addition, the following are some data with regards to HA.

Transparent HA (TRHA) uses 2 IPSs, which synchronize state over a SSL channel that is established across the management network. The IPS’s synchronize state 10x/second. The following information is sent over the TRHA channel:

- Connection information for rate limited flows (src/dest ip, src/dest port, protocol)
- Connection information for blocked flows (ditto above)
- Flush of all rate limited streams
- Flush of all blocked streams
- If rate limited streams flushed individually, a flush command for each flow cleared
- If blocked streams flushed individually, a flush command for each flow cleared

The last 4 only occur when the user manually flushes the stream (via the SMS or LSM), so for all practical purposes, the first 2 represent the bulk of the bandwidth across the management network. The data sent for the first 2 updates is minimal: 14 bytes per blocked or rate-limited flow (13 bytes for the quadruple and 1 byte for the command).

How much bandwidth does this translate into? In most enterprises we see 1-10 blocked flows per second. The most extreme case we've seen is 14 million blocked flows in a day, which corresponds to 162/sec. Allowing for a 10x factor for burstiness, this corresponds to 1600 blocked flows/sec. At 14 bytes of data per flow, we get:

1600 flows/sec * 14 bytes/flow * 8 bits/byte = 0.18 Mbps.

50K flows/sec corresponds to 5.6 Mbps on the TRHA link.

General Backup Best Practices

- Take IPS Snapshots and copy them to the SMS periodically.
    - If you don't use the "Copy to SMS" feature, the IPS snapshot will not be backed up with the SMS backup.
- Take SMS backups regularly and push them to an external file share.
- The events table can be large, so backup with care.
- You can only restore a backup to the same version of the SMS. This is because the DB schemas must match.
- Microsoft Windows will try to save the sms.bak file as a ZIP. Don't let it, however, you can use this fact to open up the sms.bak file and inspect what is being saved.
- During a restore operation, all SMS users will be kicked off until the restore action is complete

Backing up an SMS and Restoring to a New Machine

When upgrading from the Dell 1750 to a Dell 1850, the operation is not quite as simple as it appears with release 2.1m1. There are two known issues at the time of writing.

A. The platform is not the same (1750 vs 1850), and the Ethernet drives are different. If you restore the 1750 backup over the 1850, then it overwrites one file, which breaks the network configuration. The workaround is to prevent this file from being overwritten.

B. The license file on the new SMS will be overwritten. The easiest way to work around this is to make sure you get the license file for the serial number of the SMS ahead of time, so it can be reinstalled afterwards. Otherwise you will need to copy the files in service mode for the SMS.

Here are the steps to follow:

1. New SMS. Connect the keyboard and screen, and run the setup program, giving it the same network profile (IP address, mask, default gw etc) as the old SMS.
2. Old SMS – make a backup of the SMS. It's easiest to write this to a Microsoft file server (SMB) share, but you can also download the backup file from the SMS itself.
3. Old SMS – Unmanage the IPS's, because you will need to re-manage them from the new IPS, and an IPS can only be managed by one SMS at a time.
4. Open the backup file with WinZIP. Now open the **manifest** file:

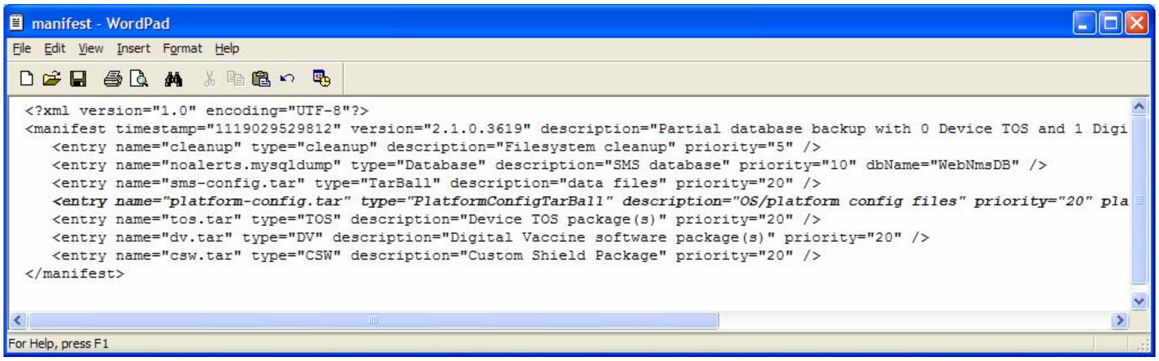

```
<?xml version="1.0" encoding="UTF-8"?>
<manifest timestamp="1119029529812" version="2.1.0.3619" description="Partial database backup with 0 Device TOS and 1 Digi
   <entry name="cleanup" type="cleanup" description="Filesystem cleanup" priority="5" />
   <entry name="noalerts.mysqldump" type="Database" description="SMS database" priority="10" dbName="WebNmsDB" />
   <entry name="sms-config.tar" type="TarBall" description="data files" priority="20" />
   <entry name="platform-config.tar" type="PlatformConfigTarBall" description="OS/platform config files" priority="20" pla
   <entry name="tos.tar" type="TOS" description="Device TOS package(s)" priority="20" />
   <entry name="dv.tar" type="DV" description="Digital Vaccine software package(s)" priority="20" />
   <entry name="csw.tar" type="CSW" description="Custom Shield Package" priority="20" />
</manifest>
```

Remove the line about Platform-config and save the **manifest** file.

5. New SMS. Restore the backup from the old SMS.
6. New SMS. Verify the network configuration.
7. Old SMS & new SMS. Disconnect the old SMS from the network (or change its IP address) and connect the new SMS.
8. New SMS. Manage the IPS's.
9. New SMS. Push the policy (to synchronize the IPSes with the SMS).
10. New SMS. Verify that the SMS can see the TMC. If not, call the TAC with the serial number of the new SMS and make sure that it is authorized to use the TMC.

An alternative (requiring service access) when the restore is done is to change the /etc/modules file so that it correctly reflects the Ethernet drivers.

Filters, Segments and Profiles

Naming Segment Groups and Profiles

Use good naming conventions for Segment Groups and Profiles since these labels will show up in reporting and in the event viewer.

Action Sets

When using an action set for the first time, make sure to check what the action set does since it could be misnamed or have been modified.

Using the SMS to Pull Existing Profiles from the IPS

If you have an IPS that has already been deployed and tuned optimally, but is not under SMS management, you can preserve the filter and action set settings that have been created and applied in two ways:

First Method

1. Click on the “Devices” icon in the navigation bar.
2. Expand the tree for the IPS that you want to import the profile from.
3. Click on “Segments”.

4. Highlight the segment that you want to import the profile from and right-click on it.
5. Select “Import Profile”.

Second Method

1. Follow steps 1 and 2 above.
2. Expand the “Segments” tree.
3. Highlight the segment that you want to import the profile from.
4. Choose “File -> Import -> Import File from Device…”

Using either of the methods above, the newly imported profile will be named in the format [device name]_[segment].

Differences between LSM and SMS Traffic Management Filters

When creating a Performance Protection Traffic Management filter in the LSM, the default settings for ICMP Attributes, Code and Type, at the bottom of the create screen are blank, so when the filter is created, LSM fills it in as ANY ANY. In the SMS, however, the defaults for ICMP Attributes, Code and Type, are 0. When the filter is created, ICMP type 0 (echo reply) and code 0 are what get filled in. Thus, if you are creating a traffic management filter that deals with ICMP traffic in the SMS and the Attributes field doesn’t get changed, you will get a different filter than if you created the filter in the LSM and don’t change the attribute field.

The inconsistency is easy to miss and will cause dramatically different behavior between the filters. Please make sure that if you are creating the filter in the SMS and you want to affect all ICMP and not just echo reply, delete the 0’s that are filled in for you automatically before hitting “Create”. If you are creating the filter in the LSM, this is done for you automatically.

This inconsistency has been fixed as of release 2.1 M1.

System Upgrades

TOS Versions of IPS and SMS

You must always be on the same or later version of SMS to manage IPS devices. For example, managing a 1.4.2, 2.0, or 2.1 IPS with 2.1 SMS is fine. Managing a 2.1 IPS with 2.0 SMS is not. If you do the latter, it will not work correctly but you can fix it by upgrading the SMS to 2.1 and redistributing the profiles.

Always upgrade your SMS first.

You will see the message “Failed to create device distribution package” on the SMS if you upgrade the IPS to 2.1.0.6305 without first upgrading the SMS to 2.1.3619. (See section “Which version of TOS should I run for SMS and IPS?” above.)

Other possible reasons for this message are:

- There is an action set with an email contact and the IPS does not have an email server defined.
- There is an action set with a rate limit not supported by the particular model of IPS.

If the items above are addressed and the message still reoccurs, try downloading the latest DV from the TMC (even if you have already downloaded the latest), and make it active.

Upgrading from SMS 1.4.2 to 2.1

There are known issues with upgrading straight from 1.4.2 to 2.1. We highly recommend that the SMS be upgraded to 2.0 first.

High Availability

HA Between Two Disparate Dell Platforms (1750 and 1850)

HA can indeed be configured on two SMS's on different Dell Platforms. On further note, the IPS license key will be determined by the system that has the larger number of IPS licenses. For example, if a Dell 1850 with a 25-IPS license is put in HA with a Dell 1750 with a 10-IPS license, the maximum number of IPS's supported will be 25, not 35.

SMS Authorized IP Address

When SMS's are put in HA mode, a virtual address is assigned to the HA SMS system. It is best practice to not put this address in the "SMS Authorized IP Address" field of the managed IPS's. This is due to the fact that the active SMS will not use the virtual HA address to communicate to the IPS's it is managing. Instead of specifying an address for the aforementioned field, use the keyword "any".

Filters and the Digital Vaccine (DV)

Filters and Their Descriptions

The Library -> Digital Vaccines section of the Threat Management Center (http://tmc.tippingpoint.com/) has an exhaustive listing of all filters and their descriptions that have been released to date.

Filters to Combat Various Threats

Here are documents that explain the filters relating to attachments, BOTnets, Instant Messaging, Peer to Peer, and Microsoft vulnerabilities.

https://tmc.tippingpoint.com/TMC/library/product_documentation/digital_vaccine/attachment_coverage.html/document_preview
https://tmc.tippingpoint.com/TMC/library/product_documentation/digital_vaccine/botnet_coverage.html/document_preview

https://tmc.tippingpoint.com/TMC/library/product_documentation/digital_vaccine/im_coverage.html/document_preview
https://tmc.tippingpoint.com/TMC/library/product_documentation/digital_vaccine/p2p_coverage.html/document_preview
https://tmc.tippingpoint.com/TMC/library/product_documentation/digital_vaccine/microsoft_coverage.html/document_preview

Threat Management Center (TMC)

Diagnosing TMC Accessibility Problems

There are two things to check – the network access and the validity of the customer number.

Network Access

The SMS requires the following outbound network ports to be active to reach the TMC correctly:

Source IP	Source Port	Destination IP	Proto/Dest. port
SMS IP	Dynamic	tmc.tippingpoint.com	TCP/4043
SMS IP	Dynamic	Akamai IP servers	TCP/80

If everything appears to be correctly configured, check the firewall logs with the source IP address of the SMS, checking which firewall policy applies to the subnet or IP address of the SMS.

If a proxy is being employed, make sure that it is configured in the TMC. If the proxy requires authentication, try to get an exception to this, otherwise with release 2.1 or earlier, it will not operate.

Support

Troubleshooting the IPS

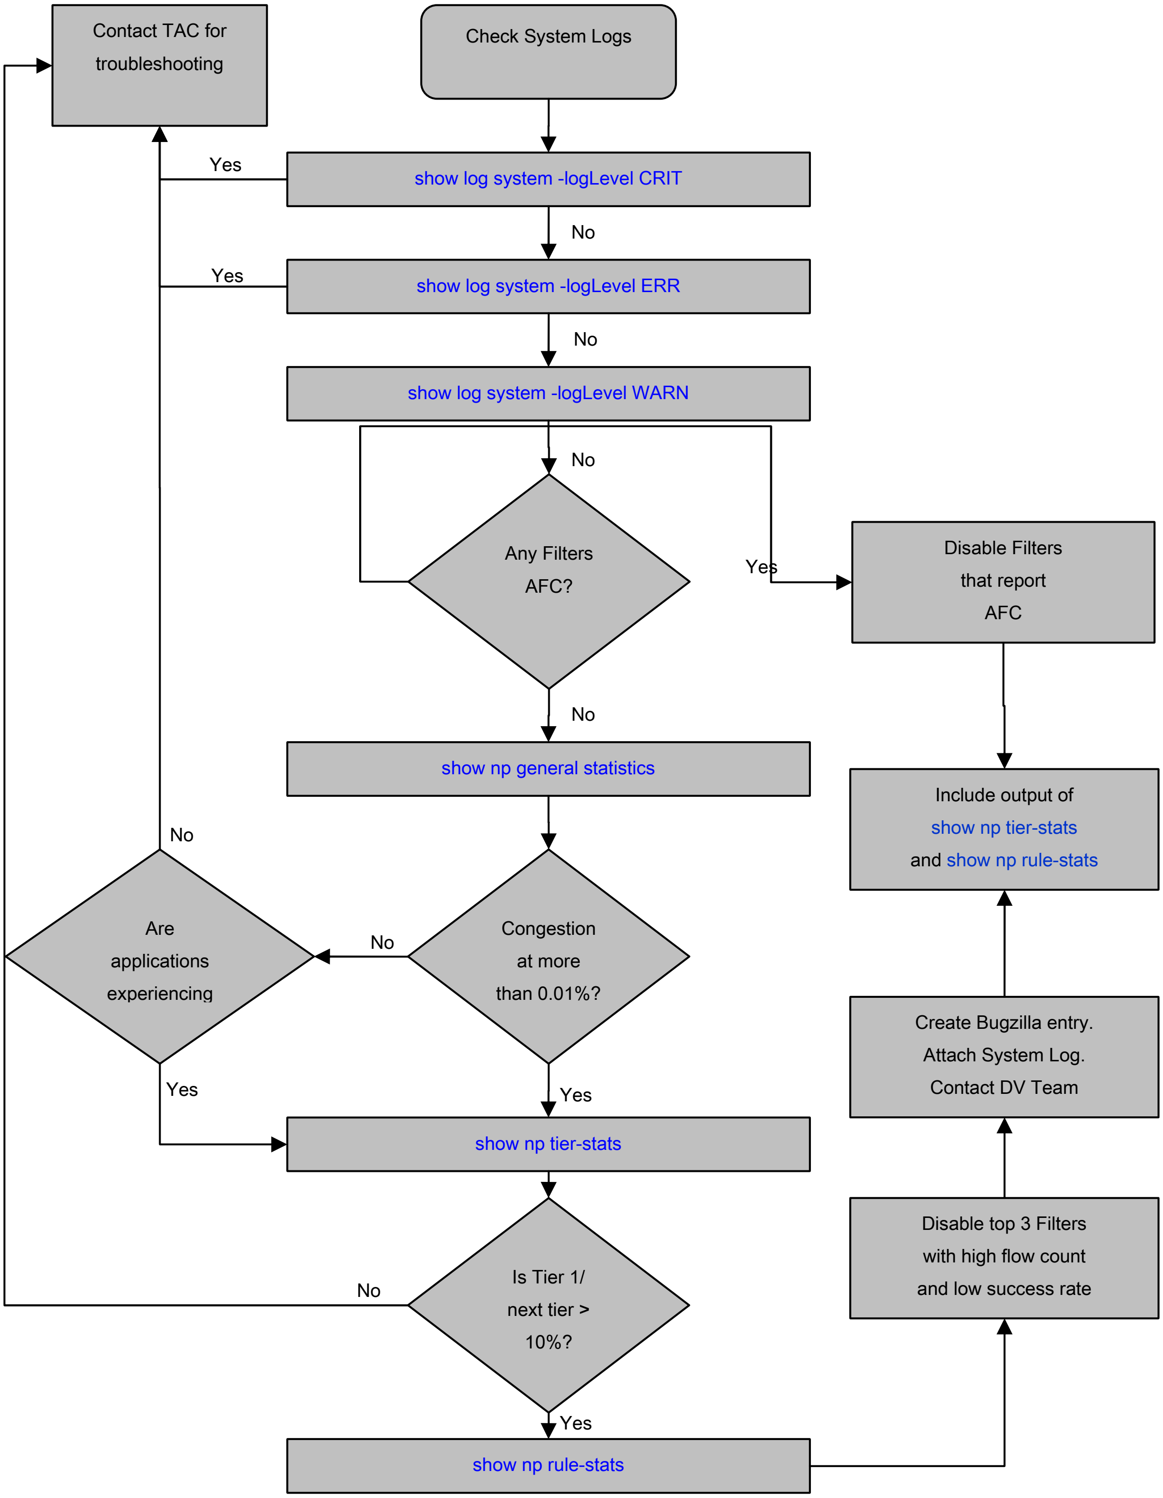
Contact TAC for troubleshooting
Check System Logs
show log system -logLevel CRIT
Yes
No
show log system -logLevel ERR
Yes
No
show log system -logLevel WARN
No
Any Filters AFC?
Yes
Disable Filters that report AFC
No
show np general statistics
Include output of show np tier-stats and show np rule-stats
No
Are applications experiencing
No
Congestion at more than 0.01%?
Create Bugzilla entry. Attach System Log. Contact DV Team
Yes
Yes
show np tier-stats
Disable top 3 Filters with high flow count and low success rate
No
Is Tier 1/ next tier > 10%?
Yes
show np rule-stats

How to Contact Technical Assistance Center (TAC)

E-mail

Please include your customer ID in all support requests sent to support@tippingpoint.com

Phone

The phone number for TippingPoint Technical Support is +1(866) 681-8324. Please be prepared to give your customer ID to the support technician. If you are outside the U.S., please call +1(512) 681-8524 for support.

The hours of operation for Standard Maintenance customers are 8AM to 5PM Central Time, Monday through Friday, with a two hour callback response time.

Digital Vaccine and Premium Maintenance customers are entitled to phone support 24 hours a day, seven days a week.

Please have the following information ready when contacting technical support:

(1) Customer number
(2) Model number of the product, OS version and Digital Vaccine Version. This can be obtained by logging in to the command line interface and typing ‘show version’

```
IPS # show version
Serial: IPS2400C-0146-4804
Software: 1.4.2.6024  Build Date: "Oct 11 2004, 19:03:21"
Digital Vaccine: 2.0.0.1888
Model: 2400
Rev: B
```

TippingPoint Users Group

There is a TippingPoint users group open to all TippingPoint end users. Visit the URL http://lists.unc.edu/read/all_forums/subscribe?name=tippingpoint to subscribe.