



INTERNET
SECURITY
SYSTEMS™

RealSecure® 6.0

Last updated: 6/14/2001

1. What is RealSecure?

RealSecure™ is an enterprise threat management system. It provides end-to-end protection from internal and external threats against your network assets. With RealSecure, you can open your networks to enable e-business while maintaining accountability. RealSecure detection components monitor network and server activity for signs of malicious intent, such as denial of service attacks, unauthorized access attempts, and pre-attack reconnaissance probes. When RealSecure detects such activity, the system can respond in a variety of ways, including recording the event, notifying the network administrator immediately, and terminating the attack automatically. By providing a variety of detection and response modules as well as a sophisticated management system, RealSecure offers threat management for your enterprise network.

2. What components comprise the RealSecure system?

The RealSecure system uses a distributed client-server architecture and its components fall into two functional categories:

- a) **Sensors.** A class of modules that provide automated detection and response to threats. These modules are installed at strategic locations throughout the enterprise network and include:
 - A **Network Sensor** that monitors network traffic in real time for signs of malicious intent and responds automatically.
 - A **Server Sensor** that monitors both inbound and outbound network traffic directed at a single host as well as the operating-system log entries and key system files for indications of intrusion or unauthorized activity.
 - An **OS Sensor** that monitors operating-system log entries and key system files for indications of unauthorized activity and responds automatically.
- b) **Workgroup Manager.** A module that provides for configuration of the Sensors as well as detailed management and storage of the threat data generated by the Sensors. All management of RealSecure Sensors is accomplished across secure communications channels. Workgroup Manager modules include:
 - A **Console** that allows centralized control of remote Sensors and provides for centralized display of alerts and reporting.
 - **Event Collectors**, which collect data from many Sensors in real-time and send it to the Enterprise Database and Console.
 - An **Enterprise Database** that stores the Sensors' event data.
 - An **Asset Database** that contains information about assets including Event Collectors and Sensors

3. What kinds of threats does RealSecure recognize?

RealSecure recognizes two types of threats against the enterprise network:

Attacks Activity patterns indicating that someone may be engaged in malicious, unauthorized, or otherwise undesirable activity involving the systems and/or data on your network. Examples of these include Denial of Service attacks (such as WinNuke, SYN Flood, and LAND), Unauthorized Access Attempts (such as Back Orifice access and Brute Force login), Pre-Attack Probes (such as SATAN scans, stealth scans, and connection attempts to non-existent services), Suspicious Activity (such as TFTP traffic), attempts to install backdoor programs (such as rootkit or BackOrifice2000), attempts to modify data or web content, and attempts to stop services or kill programs.

Misuse Non-attack activity that violates stated security or appropriate use policy. Examples of these include abuse of administrator privilege (installation of inappropriate services), HTTP activity (who's surfing the net and where they're going), analysis of access to Windows shares (e.g., connections from engineering to accounting), and e-mail session decoding.

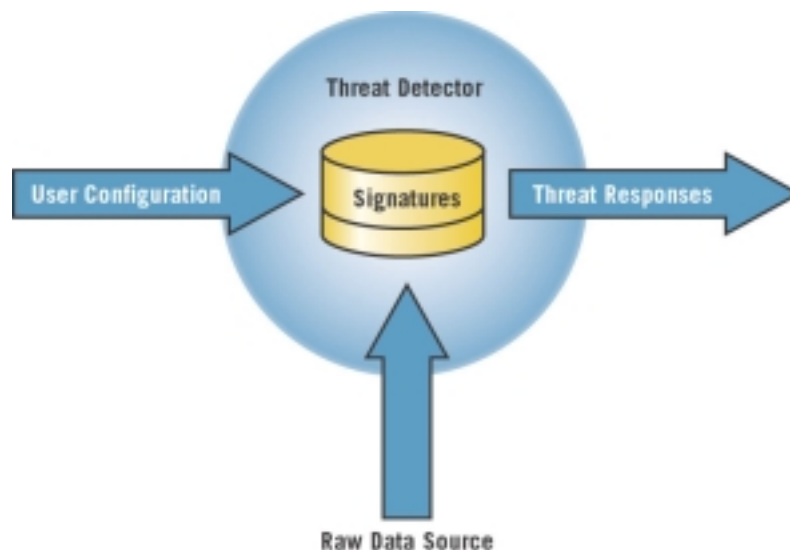
RealSecure is shipped with the most comprehensive set of threat detection signatures in the industry.

4. How do the RealSecure Sensors work?

RealSecure Sensors have a similar structure, although they vary considerably in what they detect and how they respond.

RealSecure Sensors are policy enforcement engines. The basic structure of a RealSecure Sensor can be viewed as a generic-processing module, as below. The inputs to the system include the user or administrator-specified configuration rules as well as the raw data source used to detect threats. For Network Sensors, this data source is raw network packets; for the OS Sensor, this data source is operating system log entries. The outputs of the system include the threat responses that the system initiates. The Sensor itself receives the data, compares it against the signature base, and, if there's a match, initiates the appropriate response. The signature base comes from ISS' X-Force research and development team and is the most comprehensive database of attack signatures in the industry today. Additionally, user-defined signatures are available in all Sensors, supporting further customization.

The RealSecure Network Sensor is installed on a host having a network adapter card. RealSecure puts the adapter card in promiscuous mode so that it receives all the traffic on the local network segment. If a packet meets the filter criteria currently in force, it is parsed through decode and attack recognition logic. Each active session is maintained and tracked, so that attack patterns that span many packets can be detected. This way, when an "interesting event" is detected, the appropriate actions can be taken.



The OS Sensor runs as a process on a server. When a new log file entry is generated by the operating system, the operating system interrupts the OS Sensor. The OS Sensor reads the new log entry, compares it against the signatures currently in force, and, if a match is found, initiates the appropriate responses. Some signatures span multiple log entries, so the OS Sensor also maintains the state of several user activity threads at one time. Unlike other products, the RealSecure OS Sensor does not simply rely on application logs, but it utilizes kernel-level audit data, so it cannot be fooled, spoofed, or bypassed.

5. Why do I need both Network Sensors and host-based Sensors?

Because the data that each type of Sensor generates is very complementary. Network-based intrusion detection is very good at providing early warning of attacks. By monitoring the traffic stream in real-time, a Network Sensor can see a threat and often neutralize it before it has a chance to do any damage.

However, Network Sensors cannot tell you whether an attack was successful or not. The information they manage is very network-centric. Host-based intrusion detection systems complement their network counterparts very nicely. Host-based Sensors provide confirmation of an attack's success or failure and they yield system-specific event data, such as user name and file name during an unauthorized access attempt.

Server Sensors work in network environments where it is either impractical or too costly to deploy Network Sensors. As networks get faster, it becomes more difficult to monitor all inbound and outbound traffic. In addition, more networks are becoming highly switched. A highly switched network means that more Network Sensors are required to get the same level of coverage as a single Network Sensor on a non-switched network.

Host-based Sensors are important for another reason. Local users can attack a system without being detected by the Network Sensor. For example, somebody who has access to the console can try passwords all day without a Network Sensor detecting it. A valid user running the hacker utilities "getadmin" or "sechole" to add him/herself to the administrator's group, or someone trying to open a file without permission or trojan a system file, can do so without being detected by the Network Sensor. All of these examples *are* detected by the RealSecure OS and Server Sensors.

By deploying both Network Sensors and host-based Sensors, you can have the best of both worlds: ultra-fast detection and response at the network level with rich, system-specific confirmation of events at the host level. In addition, the combination of Network Sensors and host-based Sensors is the most effective way to provide threat coverage to a switched network.

6. What additional capabilities does the Server Sensor have?

New threat management features that you won't find in other host-based intrusion detection systems:

- a) **Inbound and outbound network monitoring for intrusions.** The Server Sensor applies the same kind of signature analysis on network traffic that the Network Sensor does.
- b) **SecureLogic alert enhancement.** The Server Sensor allows you to define sophisticated rules to apply after an event is generated to perform additional analysis. Using SecureLogic rules means that false positives are reduced, fewer alarms are forwarded, and the information that is returned as part of the alarm is more complete.
- c) **Audit policy management.** The ability to manage an operating system's audit policy through the RealSecure manager user interface. This allows you to ensure that all of your critical servers have a consistent and effective audit policy including the management of true kernel-level audit, when available.
- d) **Firecell blocking.** The Server Sensor allows you to create blocking rules that are active at all times and that provide an extra level of protection. For example, Firecell rules can be utilized to allow application server traffic to a database server, but to block all traffic to the database server from any other machine on the network, or on the Internet.
- e) **Suspicious connection monitoring.** These are TCP listeners that make a port seem active – a listener on port 80 makes an attacker think you have a web server running on a given host. These are extremely useful for three reasons:
 - Attackers look for services on a host to penetrate. The suspicious connection monitors make inactive ports look active – so that the attacker wastes his time trying to penetrate a service that will gain him nothing. This is particularly useful for defeating automated tools.
 - An attempt to connect to a non-existent service is either an error or an intentional attack. The OS Sensor uses these connection attempts as another source of intrusion detection data. Persistent connection attempts to non-existent services yield high-level alarms.

- The administrator can set up banner responses to unauthorized connection attempts. These are purely optional and completely configurable. They might include a legal warning that trespassing is not allowed. Or they might include a fake connect string, like:

login:

Note that the suspicious connection monitors are not honeypot tools or tools meant to deceive an attacker for very long. They are merely intended to make an attacker take longer than ordinary to identify which ports are attackable, allowing the defender more time to catch him.

7. What impact do the Sensors have on the enterprise network?

The Network Sensor is completely unobtrusive. It only monitors the traffic on the local network segment and does not interrupt the traffic stream in any way. The Network Sensors do not add any delay to the network segment.

The impact of the OS Sensor on the server on which it is running is configurable. For single-user systems, the Server Sensor typically uses between 2-4% of the processor. However, this usage can increase on multi-user systems with large number of concurrent user sessions. This is why the RealSecure Manager provides complete control over the signatures that are currently active **as well as** the underlying system audit policy. By allowing you to specify how much auditing is done on the target server, the Server Sensor allows you to control how much processor time is spent defending the system from attack. The RealSecure Server Sensor also completely handles all of the tasks associated with auditing on a system. It manages the audit flags (so you don't have to worry if you have turned on too much or too little auditing) and it does data management for you so there is no worry about logs growing too large and wasting precious disk space.

8. How does RealSecure respond to attacks?

The actions taken upon detection of an attack or unauthorized activity are determined by the administrator and fall into three categories:

	Network Sensor	Host Sensors
Notification	Send alarm to console	Send alarm to console
	Send e-mail	Send e-mail
	Send SNMP trap	Send SNMP trap
	View active session	
	Notify Lucent Management Server	
Storage	Log summary	Log summary
	Record network session	
Active	Kill connection (TCP Reset)	Terminate user session
	Reconfigure Check Point Firewall-1	Disable user account
		Block network traffic (Server Sensor)
	Execute a user-specified program	Execute a user-specified program

The last option ("Execute a user-specified program") can be used to initiate any response that can be expressed in an executable binary (or batch file/shell script) form. Examples include initiating a pager call, playing a sound, or reconfiguring a network device that does not have an API for management.

The RealSecure system offers security administrators the widest variety of intrusion response options in the industry today.

9. How does the Check Point FireWall-1 response option work?

For this response, the RealSecure Network Sensor will send a message to the FireWall-1 management server instructing it to prevent the attacking source address, port and/or service from traversing the firewall boundary for a user-specified period of time. This period can range from one minute to forever. The FireWall-1 management server handles the clearing of this address block. The communication between the RealSecure Network Sensor and the FireWall-1 management server is done using SAMP, Check Point's suspicious activity monitoring protocol, which is part of the OPSEC framework. This communication can also be authenticated, if desired.

10. How does the Lucent Managed Firewall response option work?

In the current version of RealSecure, the LMF response option sends a real-time alarm to the Lucent Security Management Server (SMS) via a secure channel. This means that RealSecure alarms will appear on the same console used to manage your Lucent firewalls.

11. What types of protocols can the Network Sensor see and decode?

The Network Sensor can filter and monitor any TCP/IP protocol. RealSecure can interpret web, e-mail, file transfer, remote login, chat, talk and a host of other network services. In addition, the Network Sensor can monitor and decode Microsoft CIFS/SAMBA traffic for Windows networking environments. The range of services that RealSecure Network Sensors can analyze is extended regularly so be sure to check the ISS web site at <http://www.iss.net> for the latest status.

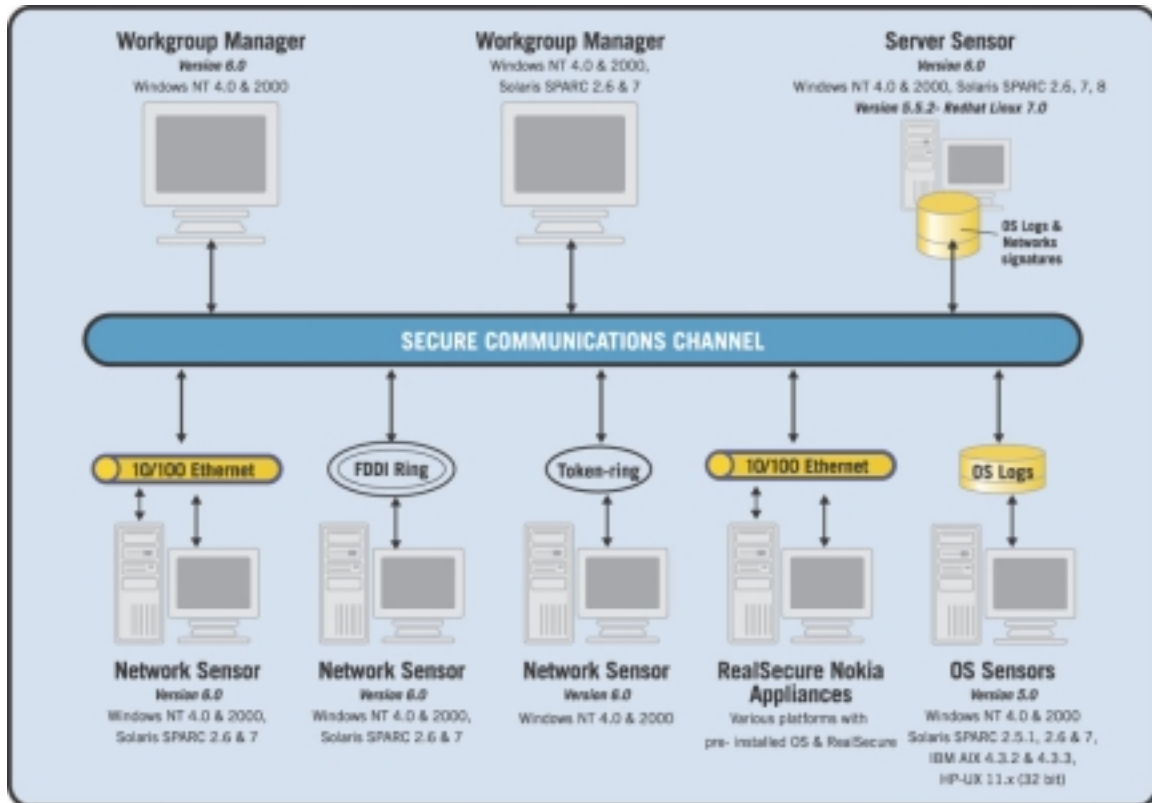
In addition to the predefined attack signatures that ship with RealSecure, the security administrator can define connection events based on protocol (TCP, UDP, ICMP), source port, destination port, source IP address, and/or destination IP address.

Finally, users can define character strings and context for custom pattern matching that does not sacrifice performance.

12. What networks can the Network Sensor monitor?

RealSecure operates on Ethernet networks (10 Mbps), Fast Ethernet networks (100Base-T only, 100 Mbps), FDDI (100 Mbps), and Token Ring networks (4 Mbps to 16 Mbps) on NT only.

The following diagram depicts the platforms and networks supported by RealSecure:



13. What platforms can RealSecure run on?

Currently, RealSecure is supported on the following platforms:

- Network Sensor: Windows NT® 4.0; Windows 2000, Solaris SPARC 2.6 & 7
- Server Sensor: Windows NT® 4.0; Windows 2000, Solaris SPARC 2.6, 7 & 8 (and SS 5.5.2 on Redhat Linux 7.0)
- OS Sensor: Solaris SPARC 2.5.1, 2.6 & 7, AIX 4.3.2/4.3.3, HP-UX 11.0 PA-RISC
- Workgroup Manager: Windows NT 4.0 & Windows 2000
- RealSecure CLI Manager: Windows NT 4.0, Windows 2000, Solaris SPARC 2.6 & 7

For complete up to date system requirements go to
http://documents.iss.net/literature/RealSecure/rs_sysreqs.pdf

Note that the OS/Server Sensor can also detect intrusions on other Unix servers through a) its ability to receive Unix syslog messages from remote servers and b) an extensive list of Unix-specific signatures.

14. What adapter cards does the RealSecure Network Sensor support?

RealSecure will operate over any adapter card that is capable of supporting promiscuous mode. Although most of the adapter cards on the market today support promiscuous mode, you should check the documentation for your network adapter to determine whether your card has this capability.

For Intel platforms, ISS recommends PCI-based, bus mastering adapters. ISS has tested RealSecure successfully with the following Fast Ethernet adapters:

- 3Com Fast Etherlink XL 3C905
- Compaq Netelligent 10/100 TX PCI
- Intel EtherExpress PRO/100
- SMC EtherPower 10/100
- Bay Networks NetGear FA310TX PCI

ISS has successfully tested RealSecure with the following Token Ring adapters:

- IBM PCI Token-Ring Adapter
- Madge Smart RingNode/BM2
- Olicom Token Ring PCI/II 16/4

ISS has successfully tested RealSecure with the following FDDI adapters:

- Network Peripherals NPI Adapter
- Osicom 2200 FDDI Adapter

15. What are the recommended specifications required for a host to run RealSecure?

- a) For the **Workgroup Manager** components, memory and disk space are the critical system resources. In addition, a monitor that supports 800 x 600 resolution and at least 256 colors is required for the Console user interface. Please refer to the latest system requirements document for detailed specifications.
- b) The **Server Sensor 6.0** will operate on NT 4.0 SP 6a, Windows 2000, and Solaris SPARC 2.6, 7 & 8 (and Server Sensor 5.5.2 on Redhat Linux 7.0). It will function on NT Server as well as NT Workstation. While it will consume some system resources in monitoring network traffic and analyzing log entries, the amount of system resources that it consumes are completely configurable by the user. The RealSecure Server Sensor is designed to run on existing systems, so purchase of new hardware is not required.
- c) The **OS Sensor** will operate on Solaris SPARC 2.51 host, AIX 4.3.2/4.3.3, HP-UX 11.0 PA-RISC. It will function on NT Server as well as NT Workstation. While it will consume some system resources in analyzing log entries, the amount of system resources that it consumes are completely configurable by the user. The RealSecure OS Sensor is designed to run on existing systems, so purchase of new hardware is not required.
- d) For the **Network Sensors**, system requirements depend on too many factors to describe with any accuracy: traffic rates, packet sizes, active signatures, selected responses, percentage of traffic stream that matches signatures, etc. Given a system with a "minimum recommended configuration", it's possible that prevailing conditions on a network would make that system unacceptable as the network grows and changes.

Therefore, you should adhere to the following guidelines when provisioning a system to run the RealSecure Network Sensor:

- Purchase as much processing power, RAM, and disk space as possible. While you may not require these resources all the time, it is vital that these resources be available during a structured attack. The Network Sensor is multi-threaded and can take advantage of multiple processors.
- Run the Network Sensor on a dedicated host. You should not use the host running RealSecure for other purposes. This is because RealSecure needs to be able to use as much processor capacity as possible in the face of a structured attack.

- Run the Network Sensor and the Workgroup Manager components on **separate** hosts. Since both components can be very processor- and memory-intensive, they need to operate on different machines. Running the engine and all the Workgroup Manager components on the same host is called the *localhost* configuration. The localhost configuration is appropriate for demonstration and evaluation purposes, but should never be used in a production environment.

16. How is RealSecure deployed across the enterprise network?


RealSecure uses a distributed architecture. The RealSecure Sensors perform the threat detection and response functions on critical network segments and servers. The Event Collector(s) collect events from the Sensors for storage in the Enterprise Database and the RealSecure Console displays alarms, consolidates engine data, provides report generation capabilities, and acts as a centralized engine management point.

The relationship between Sensors and managers is many-to-many. Several RealSecure Sensors can report to a single Event Collector. Up to 5 Event Collectors can send data to a single Enterprise Database. This is all independent of the number of Consoles used for reporting, command and configuration. This flexibility is useful for environments where there are geographical or organizational management boundaries.

With regard to placement of RealSecure Sensors, the best rule is to place a RealSecure Network Sensor on each segment where there is critical data to protect, or a set of users that should be monitored. Note that a RealSecure Network Sensor will only see the traffic that is on the local network segment. Since routers, bridges, switches, and firewalls prevent traffic from being copied to inappropriate segments, several RealSecure engines will be needed for complete coverage of your critical network resources.

You should also install a Server Sensor on all servers containing critical information. These include everything from internal file servers to external DMZ devices and communications servers.

The following figure shows a sample deployment of RealSecure on a distributed enterprise network. The light gray (yellow, if color) circles represent Network Sensors while the dark gray (burgundy, if color) circles are the Server Sensors.



- There is a RealSecure Network Sensor outside the firewall, between the DMZ and the Internet. This engine will detect attacks coming from the Internet and can be used to focus on defending the DMZ devices as well as the firewall itself.
- There is also a RealSecure Network Sensor just inside the firewall. This engine will detect unauthorized activity from the Internet that makes it through the firewall and can be used to validate firewall configuration.
- There is also a Network Sensor on the segment with the dial-up access server. Since the dial-up server represents another way for external traffic to access the network, the Network Sensor on this segment is also watching for external attacks. In general, you should place a RealSecure Network Sensor on every segment that is associated with your network perimeter.
- In addition, the network above has a trusted relationship with a business partner. Trusted relationships extend the network perimeter. A Network Sensor has also been placed behind the firewall of the corporate partner, ensuring that intruders do not compromise the partner's network and exploit the trusted relationship back to the home office.
- Network Sensors have also been placed on important internal segments to protect the vital data that is stored locally.
- In addition, the network above has Server Sensors installed on a variety of important systems: DMZ web servers, internal web servers, critical internal file servers, etc. These internal servers may be hosting such critical data as financial plans, sales databases, or engineering source code archives.
- In this example, all of these engines are reporting to a single Workgroup Manager. This console is located on the Intranet backbone. However, it might also be located at a headquarters across the Internet or at a Network Operations Center staffed by a service organization.

In this sample deployment, RealSecure Sensors have been deployed liberally across the enterprise network. This provides two distinct advantages:

- Each Sensor can be customized for the local area. For example, the Network Sensor on the dial-up access network can be configured to examine all inbound traffic from the dial-up server; The Server Sensor on the external web server can be configured to monitor the web pages as a key system resource.
- Sensors inside the network will also detect unauthorized activity that is initiated within the network, an engineer attempting to gain root access to the source code archive, for example.

17. How are updates handled? Can an administrator upgrade fifty Sensors across an enterprise (for example) without losing configuration settings?

RealSecure has X-Press Update capabilities, a technology that allows customers to take full advantage of ISS X-force research. From the ISS website, customers can download just the latest attack signatures and remotely distribute them to their Sensors. As always, our signatures are verified, QA tested and digitally signed for authentication within the product.

Improved remote management of Sensors. The RealSecure X-Press Update mechanism also supports the ability to completely upgrade Sensors to the new version. Once a RealSecure 6.0 Sensor is installed, customers can upgrade that Sensor to the next released version from a remote or centralized RealSecure Console. This makes it easier to manage more Sensors.

Simplifying management. X-Press Updates simplify the process by automatically pointing you to the latest, cumulative update for a Sensor. Its easy-to-use design ensures that only the software appropriate for your platform is installed.

Since configuration settings (i.e., filter rules, enabled attacks, Sensors being managed) are saved in separate configuration files, installation of new software will have no effect on the current settings.

Note that, when ISS provides signature updates, we also provide policy update files. These are the new-signature-specific additions to your Sensor configuration files that will automatically be added to your Sensor configuration files when you download the signature update. This means that your customized policies will be updated without erasing your customizations.

18. What impact do the host-based Sensors have on the machines that are running them?

Because we only enable the minimum auditing required to detect the kinds of attacks that you have configured, and because of today's efficient and powerful hardware, the host-based Sensors are not resource intensive. Without using RealSecure, enabling auditing in the past would cause an increase in disk usage, and sometimes logging would actually fill up all available disk space and cause the system to crash! Fortunately, the host-based Sensors do 100% event and disk management for you. Now there is no need to worry about an increase in disk usage.

As you'd expect, much of this is dependent on how the host-based Sensors are used and how much auditing is set up. Like most management or security products, the real answer is a resounding, "it depends". The user needs to be careful not to create rules that audit, "read" for everything (as an example) and choose something more specific/less common to keep the resource usage to a minimum. This type of logic is built into the existing signatures.

a) What affects RealSecure host-based Sensor memory utilization?

ISS has seen negligible memory utilization, at about 10MB. See the independent, 3rd party Security Focus RealSecure deployment whitepaper at <http://www.securityfocus.com/focus/ids/articles/issrealp1.html>

b) What affects RealSecure CPU utilization?

ISS has typically seen between 1%-5%, dependent on the following factors:

- Cost of Increased Processing Time
 - Least significant cost b/c auditing generally does not occur during computational-intensive tasks (image processing, complex calculations)
 - Single user workstations usually have plenty of CPU Cycles.
- Cost of Analysis
 - Proportional to amount of audit data collected.
 - Includes the time it takes to merge and review audit records, and archival.
- Cost of Storage
- Most significant cost. Amount of audit data depends:
 - Number of Users
 - Number of machines
 - Amount of use
 - Degree of security required
 - Beware of Full auditing, i.e. turning on read/write file auditing

<http://www.securityfocus.com/focus/ids/articles/issrealp1.html> for independent verification of less than 5%.

19. How does RealSecure perform in a heavily utilized segment?

To determine the performance of the RealSecure Network Engine, it is important to look at the following factors (some factors cannot be controlled):

- CPU type and speed
- Numbers of CPUs
- Amount of RAM
- The Policy enforced on RealSecure (signatures and responses that are selected)
- The number of packets on the wire & their size (small packets can overwhelm the NIC at lower bandwidths)
- The number of packets that match attack signatures which are enabled in RealSecure
- The bursting amount of the traffic

RealSecure's performance is enhanced by the following factors:

- Optimizing the filter rules

- Moving the response options for user-specified events out of the filter module
- Speeding up the packet driver
- Modifying the attack recognition logic for constant lookup time for signatures and data structures.

ISS is in the process of conducting a more sophisticated analysis where the variables discussed above will be used to understand each variable's impact on the performance of RealSecure.

20. How many RealSecure Sensors can a single RealSecure Workgroup Manager manage at one time?

The Console has been designed and tested to handle up to 50 Sensors. Technically, further scalability may be possible, however, there is a human limit that is reached well before any technical limitation. The real number that works for your organization will depend on how you configure the RealSecure Sensors and Event Collectors and how you have elected to handle real-time incident response within your organization.

RealSecure Sensors *can* generate quite a bit of data. Some organizations choose to establish a steady flow of real-time information from the engines to the Consoles so that the humans monitoring the events at the Consoles can respond to events immediately. In configurations like this, the number of Sensors sending alerts to a single Console is limited to the response capabilities of the humans at the Consoles. While this number can vary widely, it has typically been in the range of 10 to 20 Network Sensors.

Other organizations, however, emphasize post-facto forensics over real time response. In these situations, RealSecure Sensors are configured to generate a very small amount of real-time alerts and, instead, write event information into the Enterprise Database. In configurations like this, the number of Sensors that can be effectively managed by a single console depends largely on how much real-time analysis is required. While this number can also vary widely, it is typically in the range of 20 to 30 Network Sensors.

Still other organizations with very large deployments choose to perform virtually all of their administrative tasks from outside the Console, building management into their existing infrastructure utilizing the ISS Command Line Interface for Sensor management.

21. How many Consoles can a single Event Collector send alerts to at one time?

As with the previous question, this relates more to how humans choose to handle the real-time data and how much network bandwidth organizations are willing to dedicate to security management information exchange.

A single RealSecure Event Collector *can* send its alarm data to up to 50 distinct Consoles. However, this could (in case of a burst of activity) consume huge amounts of network bandwidth (and become a denial of service event in itself). For existing installations, the average number of Consoles to which a single Event Collector (previously a single Sensor) is sending event data is approximately two and, to date, we have not seen anything larger than four.

22. If multiple Consoles can change a Sensor's configuration, how do you prevent race conditions with Sensor configuration?

While any Console has the *capability* of changing a Sensor's configuration, only one may do so at any time. We call the console that has this capability for a given Sensor that Sensor's "Master Controller". For any Sensor that it manages, a single Console can request Master Controller status. This is granted by the Sensor on a first-come, first-served basis and is automatically granted to the console that starts a Sensor. A Console retains Master Controller status until:

- It explicitly relinquishes this status
- It is revoked by a user having the MasterStatusAdministrator role for that machine
- The Sensor is restarted

In all cases, Master Controller status is granted on a first-come, first-served basis to the next console that requests it.

It's also important to note that, in order to request Master Controller status, a manager must establish a secure communications channel with a Sensor before requesting said status. This secure channel involves authenticating the Console to the Sensor's managing Daemon.

23. How do the RealSecure Sensors communicate with the RealSecure Console?

Since RealSecure uses a distributed deployment scheme, the communications channel between Sensors, Event Collector(s) and Console(s) is a critical part of the architecture. RealSecure uses a secure channel for passing messages between all components. This channel guarantees the following:

- **Reliability.** Delivery is guaranteed with no additional action required by the user, subject to the availability of the communications path.
- **Privacy.** Data is securely encrypted to prevent unauthorized disclosure.
- **Integrity.** Data cannot be modified in, added to, or deleted from the data stream without the receiving entity detecting the corruption and aborting the session.
- **Authentication.** The party receiving the communications request is sure that the request is from a known peer and that there is no party in the middle proxying the data stream. Since communications channels are always initiated by RealSecure Managers and are never initiated by RealSecure Sensors, this ensures that the Managers authenticate themselves to the Sensors before issuing requests or retrieving data.

Data from the Sensors to the Consoles includes:

- Event messages – indications that something interesting has happened. These are passed up to the Console as they occur.
- Raw session data – the keystroke or data content of a session. This information is passed up to the Console as it occurs if the response associated with an event is "View Session".
- Detailed event and log file information. These are sent into the Enterprise Database by the Event Collector as it occurs

Data sent between the Console and the Sensors includes:

- Start and stop commands.
- Changes to Sensor configuration.
- Authentication Keys.
- Log files.
- Keep-alive checks.
- Software and signature updates (X-Press Updates).

All communications between the Sensor, Event Collector(s) and the Console(s):

- Use TCP only. No UDP is used.
- All TCP ports are user-configurable, allowing you to use existing firewall tunnels for engine-console communications.
- Are encrypted, using encryption technology from Certicom and RSA.
- On Windows NT/2000 hosts, encryption functions use the Microsoft Cryptographic API. This allows you to use the best encryption available on your system (40-bit to 128-bit symmetric encryption, 512-bit or 1024 bit public key encryption) or use a different cryptographic provider entirely. Microsoft's default cryptographic service provider uses RSA technology.
- On Unix hosts, the Certicom elliptic curve public-key encryption module must be used. ISS provides a Unix engine with 168-bit DESX or 3DEX symmetric encryption modules.
- Are authenticated with a public-private key exchange algorithm.
- Are verified, with cryptographic checksums appended and checked for each message.
- Are initiated from the Console to the Event Collectors and Sensors. This is particularly important for setting up Firewall rule sets. A rule can be created on the Firewall to allow TCP traffic to be initiated only on 2998 and 901/902/903 (default settings) from the Console to the Event Collectors and Sensors.

24. What authentication scheme do you use? What encryption scheme do you use?

The authentication scheme is based on public key exchange. Since Consoles are the controlling entity, RealSecure Consoles authenticate themselves to RealSecure Event Collectors and Sensors. During installation, you may select strong or no authentication. With no authentication, the RealSecure Sensor trusts the Console based on the Console's correct use of the communications protocol. With strong authentication, the RealSecure Sensor (or Event Collector) requires that the public key of the Console has been previously registered with the Sensor (or Event Collector) via any trusted means, typically a floppy or secure network file share or secure Key distribution from a user who has the KeyAdministrator role. Otherwise, the protocol and encryption strength are identical.

RealSecure provides encryption through several options. It may use Certicom's 168-bit DESX or 3DES symmetric encryption keys and Certicom's 239-bit elliptic curve public-key technology. Unix-based Sensors can only use Certicom. For strong encryption using Certicom, the Certicom CSP must be installed on NT/2000 platforms. (Note that the RS Manager is always on NT/2000, however this holds true if used on NT/2000 Sensors also.)

On NT/2000, you may also use encryption algorithms called through Microsoft's Cryptographic API, which will use whatever encryption technology is available through that API. During installation, you specify which cryptographic service provider (CSP) you wish to use for encryption. Microsoft's default CSP is based on RSA technology and provides for encryption using 40-bit or 128-bit symmetric encryption keys and 512-bit or 1024-bit public encryption keys.

25. How does RealSecure differ from a firewall? Don't they do the same things?

Firewalls and RealSecure use similar technologies to accomplish different things. Firewalls are *controlling* entities. They enforce general entry and exit rules for an entire network and aren't designed to look for attack patterns. Their main purpose is to keep the wrong kind of traffic off the network and their definition of "wrong kind of traffic" is usually based on IP address or protocol type.

RealSecure is not a product that controls network access. RealSecure does not interfere with the network traffic stream at all. Instead, it allows the traffic to go by and quietly watches for signs of unauthorized activity. RealSecure's definition of "unauthorized activity" is a sophisticated and customizable database of attack signatures.

Think of your network as a high-rise apartment building, the firewall as the doorman, and each RealSecure Network Sensor as a guard dog on a specific floor and each RealSecure OS/Server Sensor as bodyguard in each apartment. The doorman is generally pretty good about letting the right people in and keeping the wrong people out. However, a moderately clever criminal will be able to get past the doorman and into the building. The guard dog is better at knowing who's authorized to be on the floor and responding quickly to stop the intrusion. The bodyguard has a personal responsibility to defend the apartment in which he works. He knows the area well and monitors constantly for intruders.

26. Do I need firewalls if I have RealSecure?

Absolutely. RealSecure is an essential addition to, but not a replacement for, your firewall security. When firewalls are properly configured, they keep out most undesired traffic. However, in order to provide some level of access, firewalls have tunnels that can be exploited by would-be attackers. A good example is FTP. Many companies have an FTP server inside their network and associated tunnels through the firewall to allow access. A common attack is to attempt to gain root access to the FTP server. Once an attacker has access to a system inside the network, other systems become vulnerable. And although the firewall will not stop this type of attack, RealSecure will. By monitoring the traffic stream on the network behind the firewall, RealSecure can detect and terminate attempts to gain root access on the FTP server.

27. Do I need RealSecure if I have firewalls?

Yes. Threat management is a complementary technology to firewalls access control. While firewalls provide excellent packet-level protection for your network, they can never do it all:

- Firewalls have tunnels that allow packets through. Packets that pass through these tunnels are typically **not** analyzed by the firewall, but are passed through unexamined. This means that you are relying on the security of your internal devices to protect your network for these specific data streams. For example protocols like SSH can enable virtually any traffic to be tunneled through any open port on a firewall. Such circumstances can be dangerous to the integrity of the network. RealSecure can detect the use of SSH and monitor traffic streams and their system targets for malicious activity.
- Firewalls are frequently misconfigured. Does your firewall have a filter that prevents ICMP requests? What about the new video streaming protocol that's just been announced to the Internet? Has a vice-president requested a specific hole in the firewall's filtering rules because he needs to access his personal files from home? These are examples of configuration mistakes or omissions that weaken a firewall's effectiveness. RealSecure works in conjunction with your firewall (whatever its configuration) by monitoring **all** the activity on your network and servers for attempts to breach your security. Although firewall misconfigurations should ideally be fixed as soon as possible, having RealSecure Sensors inside the network can bring attention to much of the undesirable traffic that's leaking through. Even if you choose not to terminate these undesired connections, the sheer number of alarms that RealSecure will generate will quickly indicate that your firewall is not doing its job.
- Firewalls can be compromised by an external attacker. The security marketplace is a dynamic, evolutionary battleground between attackers and defenders. Attack methods that were popular six months ago have been defeated by today's defense systems, but have been replaced by new methods. There are now techniques to scan through firewalls that didn't exist one year ago. New Trojan horse programs and denial of service attacks are being discovered weekly. RealSecure can help you stay ahead of the curve with technology that sees all the packets and notes the suspicious anomalies. Moreover, if your firewall is compromised, RealSecure offers another processor dedicated to the defense of your network – doubling the work an external attacker must do to penetrate your enterprise.

28. What do I have to do to my network to run RealSecure?

Nothing. You don't have to buy new routers or bridges; you don't have to install any software on your servers; you don't have to reconfigure any devices. RealSecure works with your existing network infrastructure. To install Network Sensors, all you need to do is place a UNIX® or Windows NT® system with an adapter card on the segment to be monitored and install RealSecure. To install OS Sensors, all you need to do is install the software on important servers that you want to protect.

29. Will RealSecure run on a switched network?

Yes. There are three ways to support a switched network with RealSecure:

- a) **Strategic deployment of RealSecure Network Sensors.** In many cases, a careful look at your network reveals strategic locations where a switch can be placed that will provide excellent security coverage.

For example, if one switched port were connected to a router that connected to the Internet, then it would make sense to insert a small hub between the router and the switch and connect a RealSecure engine at that point. That would provide protection against attacks coming in from the Internet, regardless of what else was on the network.

If a mini-hub is unacceptable, you can also use a switch's administrative port or even a VLAN. Some switches (the Cisco Catalyst is an example) have a management port (sometimes also called the span port) that mirrors the traffic stream on one or more specified ports. You can attach a RealSecure Network Sensor to this management port and have the management port mirror one or more critical ports on the switch. If the switch is used to connect to an Internet router, then configure the management port to mirror the port to which the router is connected. If the switch is

used to connect several critical servers to the Intranet, then configure the management port to mirror the port that connects the switch to the rest of the Intranet.

This method can be used today.

There are usually two or three strategic locations where a RealSecure engine can be placed on a switched network to provide adequate security coverage.

- b) **Use of RealSecure host-based Sensors.** Of course, RealSecure Server Sensor can fill the gaps that the Network Sensors cannot reach. Many switched environments involve server farms, with a high density of hosts connected to one or more switches directly. In this environment, a RealSecure Server Sensor on each host will protect each host from attack or misuse. Since the Server Sensors are small and completely configurable, each one can be configured to monitor the key files and functions on each server.

A combination of Network Sensors and host-based Sensors is the most effective approach. Since all the Sensor data can flow back to one or more Workgroup Managers, the security administrator can have a global view of his enterprise threat status, even in a switched environment.

30. How much delay does RealSecure add to the network?

None. Unlike firewalls, which often store and evaluate data before forwarding it to the inner network, RealSecure is completely unobtrusive. RealSecure Network Sensors monitor the network traffic, copying packets as needed, but do not alter or delay the traffic at all. RealSecure Server Sensors only process the local traffic stream – while the data is still in memory, and at very high speed. The only time that RealSecure will have any impact the traffic flow is when it terminates connections in response to an attack and this won't be noticed, except by the attacker.

31. How much additional traffic does RealSecure add to the network?

It depends. In a distributed configuration (i.e., Sensors distributed around the enterprise network reporting data to the Workgroup Manager), the amount of additional traffic will depend on several factors:

- The number and frequency of network events reported from the Sensors to the Workgroup Manager.
- The frequency of database and log file uploads from the Sensors to the console.
- The size of the database and log file uploads from the Sensors to the Workgroup Manager. The administrator may choose not to upload everything, but may want a subset of the stored information instead.

Since you control some of these variables, you can control how RealSecure adds data to your network. Most customers upload databases and log files during periods of low network usage.

There are configurations in which RealSecure Sensors will generate **no** network traffic whatsoever.

32. Can RealSecure Network Sensors be completely transparent? Must they have an IP address?

Yes. RealSecure Network Sensors can be completely transparent. A RealSecure engine can use out-of-band communications to communicate with the Workgroup Manager. This is most commonly accomplished by using two separate adapter cards: one to monitor the local network segment and another to communicate with the console. This increases the security of the RealSecure engine by having it use a separate, possibly more secure, communications channel to send information back to the console.

In addition, the adapter card being used to monitor the local network segment does **not** require a protocol stack. Therefore, the RealSecure engine does not need an externally visible IP address or externally visible IP services. The RealSecure engine can be invisible from the network that it's monitoring.

The RealSecure engine does require TCP/IP to communicate with the Workgroup Manager, however. Consequently, it requires an IP address on the interface used to send messages to the Workgroup Manager.

33. Can RealSecure detect unauthorized activity in a Windows networking environment?

Yes. The current versions of RealSecure Network Sensor include the ability to decode SAMBA/CIFS protocols for Windows networking. The product also includes several attack signatures specific to the Windows networking environment. These include the ability to detect, among others:

- When one user attempts to copy a password file (.pwl) from a shared volume on a Windows 95 system.
- Remote registry access attempts.
- Null sessions.
- Attempts to read from or write to protected shares.
- SAMBA buffer overflow attack
- NetBIOS session grant decode
- NetBIOS session reject decode
- NetBIOS session request decode

In addition, the OS Sensors detect unauthorized activity at the host level in Windows environments. These include attempts to access unauthorized files, accesses to privileged services, and attempts to change a login's access rights.

34. Can RealSecure play back logged traffic data at a later date?

RealSecure cannot playback logged data as if it were being received from the adapter card. However, network events can be stored in log files and databases for retrieval later. RealSecure provides sophisticated reporting features that allow the administrator to sort and format event data by priority, source address, destination address, or network service over some period.

RealSecure Network Sensors also offer the ability to record the raw, binary content of an entire network session. This data is stored in a log file and can be replayed through the Workgroup Manager interface. It is played back exactly as it was received, keystroke for keystroke, so that the administrator can see how the attack or session unfolded.

35. How can I configure RealSecure specifically for my environment?

There are many ways to customize RealSecure.

- a) First, you can alter the actions that RealSecure takes when an attack or event is detected. These actions include the following:
- Terminate the attack automatically.
 - Terminate the user session.
 - Disable the user account.
 - Reconfigure a CheckPoint™ Firewall-1® to reject traffic from the attacking source address.
 - Send a secure real-time alarm to the Lucent Managed Firewall Security Management Server (SMS)
 - Send an alarm to the Console indicating that the event occurred.
 - Send an SNMP trap to an off-the-shelf management platform.
 - Log the event, including date, time, source, destination, description, and data associated with the event.
 - View the raw content of the session in real-time (or record for later playback).
 - E-mail a notification to the administrator.
 - Execute a user-specified program.

- b) Second, you can define your own signatures for the RealSecure host-based Sensors. The log monitoring component of the host-based systems allows you to specify a keyword or regular expression that, when found in appropriate operating system log file entries, will cause one or more of the responses above to be taken. Because we use detailed kernel-level audit data, there is virtually no signature that you cannot define yourself.

For example, you might want to monitor your main web page to be sure an attacker doesn't change it. RealSecure can watch your index.html file, notify you if someone attempts to change the file, and automatically replace the changed file with your original before any of your customers see the defaced web page.

- c) Third, you can define your own connection events for the RealSecure Network Sensor. A connection event is defined as an IP-based connection that matches any combination of the following criteria:
- Protocol
 - Source IP address
 - Destination IP address
 - Source port
 - Destination port

For example, you might want to log all traffic to and from the server that contains the financial data for the corporation. You would do this by defining a connection event that would catch all traffic to and from the IP address of the server. Note that any of RealSecure's response options can be applied to these connection events.

- d) Fourth, several of the RealSecure pre-defined attack signatures can be tuned to match the operation of your network. If you find, for example, that PointCast downloads are triggering the SYN Flood signature (which does happen on some networks), you may want to increase the thresholds for SYN Flood so that you reduce the number of false positives.
- e) Fifth, you can instruct the RealSecure Network Sensor's filtering logic to *ignore* certain types of traffic. By specifying a combination of protocol, source IP, destination IP, source port, or destination port, you can characterize precisely which traffic RealSecure should ignore. Network traffic matching these criteria will not be analyzed for pre-defined or user-defined signatures.

You can use this to fine-tune RealSecure's operation so that it's a better match for your network environment. For example,

- Suppose you have an older device on your network that is necessary, but that keeps triggering RealSecure attack signatures. You don't want to disable the attack signature, but you don't want to keep seeing the old machine keep triggering the signature, either. You can define a filter to prevent the egregious traffic from the older machine from being analyzed by RealSecure.
 - Suppose there is a segment of the network for which you have no administrative or security authority. You might not want to see any activity associated with that portion of the network. You can instruct RealSecure to ignore traffic that initiates from or is addressed to that portion of the network.
- f) Finally, you can create your own response options. Anything that can be launched from a command line (executable, batch file, shell script, etc.) can be invoked automatically by the RealSecure Network Sensor or Server Sensor in response to an attack.

36. Can the RealSecure Network Sensor be used for URL blocking?

Yes, but in a limited fashion. RealSecure is not designed to be a "network nanny" that enforces appropriate usage policies on a network. Its real function is security.

You can install filters that match certain web sites and you can instruct RealSecure to terminate connections that match these filters. For example, you might install a new filter that terminates all connections to 208.21.4.0 (sexkitten.com) and to port 80 (HTTP). However, RealSecure is not a product that is designed to be used in this manner and you will find that there are other, easier ways of blocking certain URLs on your LAN.

37. Can RealSecure be used to match my own string definitions?

Yes. You can define regular expression strings and the context with which to apply them for both types of Sensors. This allows you to confidently define a keyword or phrase to monitor for either on the network or in host logs, without having to worry about impacts on performance of your intrusion detection systems.

38. Can RealSecure log and flag the type and size of traffic or network service?

RealSecure can log and flag the type of traffic, but not the size. RealSecure is not designed to monitor or manage the performance of the network segment, but the security.

39. How does the RealSecure Network Sensor detect a SYN flood?

Some of the attacks that RealSecure detects involve more than just a single packet or single protocol type. Some involve variables that can be tuned for your network. SYN Flood is a good example. A SYN Flood is a denial of service attack. When a TCP connection is established, the initiator of the connection (the attacker, in this example) sends a SYN packet to the destination (the target system, in this example). The target system will acknowledge the connection and allocate memory to hold information about the connection. By establishing, but not using, many TCP connections, the attacker causes the target machine to use up all of its TCP buffers and renders it unable to communicate on the network for minutes at a time. A sustained SYN Flood can render a target system uncommunicative for hours.

RealSecure detects SYN Floods by monitoring the TCP connections that are established and by setting thresholds for the number of outstanding connections on a given machine at a given time. The network administrator may adjust the value of this threshold as appropriate for the network. There are several other attack signatures, like SYN Flood, that have tunable parameters.

40. Can RealSecure data be analyzed with a decision support system?

Yes, if the decision support system is capable of reading an ODBC database. The recommended Decision Support System that is compatible with RealSecure is SAFEsuite Decisions. More information is available at <http://www.iss.net>

41. Can multiple RealSecure Sensors run on a single host with multiple adapter cards?

This is supported for Solaris & Windows platforms. One may choose **multi-homed host** configuration, which allows a single host supporting up to 2 Network Sensors to monitor 2 different network segments. This configuration is only recommended for very low bandwidth segments. As well, one may install a Network Sensor along with the host monitoring components of the Server Sensor. This configuration aids in securing the Network Sensor's operating environment. Currently, there is a limit of 2 RealSecure Sensors per host.

42. This product gathers a lot of information about my network. How should the RealSecure hosts be configured in order to protect this product from misuse?

RealSecure is an amazingly powerful tool designed for network administrators. However, it could become a potentially dangerous tool in the wrong hands. It can grab user names, passwords, and even e-mail and file transfer content. Therefore, ISS recommends the following:

- Scan the Sensors and Workgroup Manager components with ISS' Internet Scanner and System Scanner to minimize the systems' vulnerability to attack.
- Use the RealSecure Network Sensor on a dedicated host. Do not run any other applications on the system.
- On the Network Sensor, disable all services except for TCP/IP. The RealSecure Sensor reads raw data link packets from the adapter card, but uses TCP to communicate with the Event Collector and Console.
- Ensure that nothing is listening on any of the ports on the Network Sensor host except for those TCP ports you have identified for engine-manager communications. (These are completely user-configurable.)
- Ensure that root or administrator access to the RealSecure Sensors is restricted. It is a good idea to disable unnecessary logins.
- For Windows NT hosts, install all of the latest service packs and operating system patches. Note that service pack 3 is required for operating anything later than RealSecure 2.0.

ISS provides guidelines for locking down your Sensor and console hosts. You should also see question 28 for a discussion of how to minimize your RealSecure Network Sensor host's exposure to attack.

43. How do the RealSecure host-based Sensors ensure that the operating system is logging the right things?

OS and Server Sensors can automatically clear/set audit flags when you make changes to the policy so the user does not have to know what flags to set. They can also enforce auditing to be sure auditing flags are not accidentally changed by users or programs.

Obviously, the host-based Sensors cannot detect a brute force login attack if the operating system is not recording failed login attempts. The RealSecure manager eliminates this problem by allowing you to control both the configuration of the Sensors and the configuration of the operating system log files.

The **detection policy** is the configuration of the host-based Sensors. It specifies which signatures are enabled and which are disabled; identifies how the system should respond to each signature match; and lists any user-defined signatures and actions.

The **audit policy** is the audit configuration of the underlying operating system. It specifies which events are detected and logged by the operating system; identifies how log files are handled; and enumerates any key system resources that should be monitored by the operating system.

The RealSecure manager allows you to specify both the detection policy of the agent as well as the audit policy of the operating system. You are prevented from trying to detect items in the detection policy that are not logged in the audit policy – the RealSecure manager won't let that happen.

In addition, because you can push detection policies and audit policies down to remote agents, you can use the RealSecure manager to *ensure* that all of your critical servers have a *consistent* and *effective* audit policy.

44. If I have RealSecure already, can I use the host-based Sensors as part of the maintenance agreement?

No. While they are part of the same family of products, the RealSecure Sensors are licensed separately. You can add OS or Server Sensors to an existing deployment of RealSecure Network Sensors for an additional license fee. All host-based Sensors are new modules and are not included under existing maintenance agreements.

45. Isn't Server Sensor just an upgrade to OS Sensor?

No. Server Sensor is not just OS Sensor with some enhancements. It is a new product with major features that are not included in the OS Sensor. Firecell rules are not available in the OS Sensor. Network-based intrusion detection signatures are not available in the OS Sensor. The SecureLogic scripting subsystem is not available in the OS Sensor. The Server Sensor is focused on protection, not just detection, and it provides much more than log analysis.

46. How do I get a copy of RealSecure?

Call ISS at 1-800-776-2362 (in North America) or at +1-404-236-2600 (outside North America) for instructions on how to download RealSecure from our web or FTP site.

47. Whom do I contact for technical support?

You can send e-mail to support@iss.net. Or you can download our tech support FAQ from the ISS web site at <http://www.iss.net>. You can also search for solutions to known issues on the technical support knowledge base at http://www.iss.net/customer_care/knowledgebase/. Finally, you can call ISS technical support directly at 1-888-447-4861 or +1-404-236-2700. Technical support operates 24 hours a day, 7 days a week.

48. Is there an archive of technical papers and utilities for RealSecure?

You can download free tech notes, unsupported utilities and other useful RealSecure information from the RealSecure resource center at http://www.iss.net/customer_care/resource_center/realsecure_tech_center/

49. Whom do I contact with product suggestions?

Send your enhancement request to enhancements@iss.net and it will be recorded.

About Internet Security Systems

Internet Security Systems (ISS) (Nasdaq: ISSX) is the leading global provider of security management solutions that protect digital assets. By offering the industry-leading SAFEsuite®

Security management platform, managed security services, strategic consulting and education offerings, ISS is the trusted security provider for its customers, protecting online resources and ensuring safe, uninterrupted business operations. ISS supports more than 8,000 customers worldwide, including 21 of the 25 largest U.S. commercial banks, 10 of the largest telecommunications companies and all major government agencies. Founded in 1994, ISS is headquartered in Atlanta, GA, with additional operations throughout North America, Asia, Australia, Europe, Latin America and the Middle East.

Copyright © 1996-2001, Internet Security Systems, Inc. All rights reserved worldwide.

Internet Security Systems, the Internet Security Systems logo, RealSecure and the Power to Protect are trademarks and service marks, and SAFEsuite a registered trademark, of Internet Security Systems, Inc. Other trademarks and trade names mentioned are marks and names of their owners as indicated. All other trademarks are the property of their respective owners and are used here in an editorial context without intent of infringement. Specifications and content are subject to change without notice.