# Evaluating an Intrusion Detection Solution

*A Strategy for a Successful IDS Evaluation*


**6600 Peachtree-Dunwoody Rd. East**
**300 Embassy Row, Suite 500**
**Atlanta, GA 30328**
**Tel: 678.443.6000**
**Toll-free: 800.776.2362**
**Fax: 678.443.6477**
**E-mail: sales@iss.net**

**TABLE OF CONTENTS**

**SUCCESSFULLY TESTING AN IDS ..... 4**

INTRODUCTION & SCOPE ..... 4

**PHASE 1: REQUIREMENTS DEFINITION ..... 5**

**PHASE 2: IDS EVALUATION METHODOLOGY ..... 6**

GOAL ..... 6
PROCEDURE ..... 6
TESTING CONSIDERATION ..... 8
RESULTS ..... 9

➢ **APPENDIX A - NETWORK-BASED INTRUSION DETECTION REQUIREMENTS ..... 10**

INSTALLATION AND DEPLOYMENT ..... 10
SECURITY ..... 10
INCIDENT DETECTION ..... 11
INCIDENT RESPONSE ..... 12
CONFIGURATION ..... 13
EVENT MONITORING ..... 14
DATA MANAGEMENT ..... 14
PERFORMANCE ..... 15
ARCHITECTURE ..... 16
PRODUCT UPDATES, TECHNICAL SUPPORT, AND INDUSTRY RESEARCH ..... 16
OTHER ..... 17

❖ **APPENDIX B - HOST-BASED INTRUSION DETECTION REQUIREMENTS ..... 18**

INSTALLATION AND DEPLOYMENT ..... 18
SECURITY ..... 18
INCIDENT DETECTION ..... 19
DECOY SERVICES ..... 21
INCIDENT RESPONSE ..... 21
CONFIGURATION ..... 22
EVENT MONITORING ..... 22
DATA MANAGEMENT ..... 23
PERFORMANCE ..... 24
PRODUCT UPDATES, TECHNICAL SUPPORT, AND INDUSTRY RESEARCH ..... 24
OTHER ..... 25

# Successfully Testing an IDS

## Introduction & Scope

This document describes a process by which you can evaluate network-based and host-based intrusion detection products. It will help you ask the right questions, help you understand how best to explore the product candidates, and help you make buying decisions.

It is important to note that intrusion detection is best implemented by the combined use of network-based and host-based solutions. An intrusion detection solution is comprised of Sensors (network-based and/or host-based) and Management consoles. Working in tandem, these two product categories will give you a thorough, real-time understanding of the risks to your enterprise's security.

### Network-based Intrusion Detection Sensors

There are certain types of attacks and network or system misuse that can be most easily discovered by monitoring network traffic. Network-based intrusion detection software monitors and evaluates network traffic for two types of problems - 1) Attacks, either from outside or inside the network, and 2) network misuse as it is defined by your organization's security policy. Network-based Intrusion Detection is deployed at key points within your enterprise network. For example, it is usually deployed at access points to your internal network where it works with your firewall(s) to protect the perimeter. Strategic deployment of an IDS would also include crucial LAN and WAN segments. Because it is installed on relatively few, dedicated hosts and is used to monitor entire network segments, it is easier to deploy than host-based solutions. In this document, paragraphs that refer specifically to network-based products will have a unique bullet. For example:

- This paragraph refers to network-based intrusion detection solutions.

### Host-based Intrusion Detection Sensors

On the other hand, there are certain types of attacks and misuse that can only be detected by software that actually resides and runs on the target system. For example, a host-based intrusion detection product can detect an intruder trying to logon unsuccessfully at the keyboard. This category of product is also more appropriate when you need to provide unique defenses for different mission-critical systems on your network. Host-based products are usually designed for use with your most critical systems. In this document, paragraphs that refer specifically to host-based products will have a unique bullet. For example:

- This paragraph refers to host-based vulnerability detection solutions.

Finally it is important to note that these products should provide a range of functions associated with security monitoring. By this we mean that "intrusion" problems actually consist of a) external attacks, b) internal attacks, and c) network usage that is out of compliance with an organization's security policy. Your intrusion detection solution should be able to monitor your organization's network and systems in such a way that all three of these areas are addressed.

### Management Consoles

An IDS Management Console is an important piece of an effective IDS Solution. The console represents a central management point for the intrusion detection solution. From a console, you should be able to minimally do the following: receive alarms from all the sensors (Network and Host based) connected to the console, control the sensors and sensor configurations, remotely upgrade the sensors, collect data from the sensors, and generate reports on network activity.

### IDS Solution evaluation

A typical evaluation process consists of three phases:

Phase 1: IDS Requirements Definition – Identifying your specific requirements for Intrusion Detection for complete coverage of your critical assets and compliance to your security policy.
Phase 2: IDS solution Evaluation - Selecting the right product to meet your needs.
Phase 3: IDS Deployment - Deploying the solution in your organization effectively.

At each stage of the evaluation process, we recommend assistance by the vendor(s) of the intrusion detection products to ensure that you get an accurate assessment of the product's capabilities.

Through the structured approach provided by this document, you will be able to identify your needs, objectively evaluate LAN-based intrusion detection products, select the right product, and implement it successfully in your organization.

## Phase 1: Requirements Definition

The Requirements Definition Phase is the most important phase in gathering information to select an IDS best for your environment. To determine specific requirements necessary, first determine where your critical assets are located. A complete understanding of your company's business model is extremely helpful in understanding what needs to be protected. In some cases, companies are deploying a complete e-commerce solution where several levels of protection are necessary. These levels can include perimeter protection, application protection, e-business protection, key server protection, policy compliance, and preventing legal liability. Make sure that you have clearly defined and prioritized your requirements for an Intrusion Detection Solution before embarking on a full-scale IDS product evaluation.

## Phase 2: IDS Evaluation Methodology

### Goal

During the evaluation program, you will become familiar with the issues of both network-based and host-based intrusion detection and the solutions provided by both categories of product. By the end you will:

- Have installed and run each IDS Solution.
- Understand how well each IDS detects attacks and misuse on your networks.
- Understand how well the products prioritize and explain the attacks and occurrences of network and host misuse, including how false positives might be generated.
- Understand the reporting capabilities of the product, including its completeness, its flexibility, and its applicability to technical as well as managerial people.
- Have a better sense of the effort required to integrate the products into your organization's operational process and to use it on an on-going basis.

### Procedure

In order to successfully evaluate an IDS, it is important to get a good understanding of its features, abilities, and limitations. Since no two IDS offerings are exactly alike, it is important to comprehend the differences and their value.

Try to resist too many "mid-stream" changes to your Requirements list; this makes it more difficult to do an "apples-to-apples" comparison of products. If you find the requirements *are* changing, then go ahead and finish the tests, but then go back, re-work your requirements, and re-evaluate the solutions. Make sure that any requirement changes are based on your environment needs, and not based on the vendor's suggestions.

When evaluating an IDS, areas to consider are:

1) Installation and Deployment:
   a) IS a network and host sensor solution provided from a single vendor?
   b) How easy is it to deploy the IDS sensors and console?
   c) How much security knowledge is needed to guarantee a successful implementation?
   d) How long does it take to install, setup, and configure?
2) Security:
   a) Is the IDS sensors and console secure?
   b) How easily can it be subjected to attack?
   c) Are communications between the sensors and console secured? Are all communications authenticated and encrypted?
   d) Can you deploy the network sensor so that it cannot be detected by would-be attackers?
   e) Can the host sensors be deployed securely?
   f) What prevents the management console from being 'taken over' by a would-be attacker?
3) Incident Detection:
   a) Is the incident determined in real or near real time to enable speedy response to the attack?
   b) How does the IDS determine if an attack is happening?

c) Does the IDS have strong attack detection capabilities - Provided Signatures, String matching, User Defined ports for non-standard port utilization IE. HTTP on 8080 or 8000?
d) Does the IDS Provide the strongest set of attack signatures that do not require Protocol, Security, or Coding knowledge, but also allow the user to create a variety of string based signatures?
e) Are their capabilities to reduce the number of alerts to ones that represent an attack thereby eliminating the need to respond to too manyh alerts?

4) Product Update Capability:
   a) How frequently is the product updated with attack information?
   b) How often does the IDS Vendor provide product updates?
   c) What capability is there to distribute the updates across the installed sensor base?
   d) What are the security mechanisms to remotely distribute updates?
5) Incident Response:
   a) How does the IDS respond to an attack?
   b) What are the response options? (For Example: SNMP traps, User-Defined Actions, Mail, Console, Log to Database, Log Raw Data, Reconfigure Checkpoint Firewall)
   c) What capability is there to reduce false positives? Does the IDS product provide tunable parameters to help minimize data, events, and false positives? Can you specify a propagation frequency for each signature?
6) Ease of Operational Management and deployment:
   a) Does the IDS fit with your existing supported Operating systems and network environment?
   b) Can the link to your existing SNMP management systems?
   c) Does the IDS offering support integrated host and network IDS management and reporting from one console?
   d) Does the IDS have an intuitive interface?
   e) Can you easily display events, drill down for more detail, etc.?
   f) When a likely attack pattern is recognized, how does the product respond?
   g) How easy is it to view event information?
   h) How easily can the IDS be configured and reconfigured?
   i) Can you easily retrieve Policies from any Console?
   j) Does the IDS provide the ability to have viewer/alerting consoles across several sensors? Can multiple consoles connect to multiple sensors? Is there a hierarchical relationship?
   k) How useful is the Help file?
   l) How easy is it to deploy policies across all IDS sensors?
   m) How many sensors can be easily managed from a console?
7) Data Management:
   a) How does the product maintain attack information?
   b) Does the IDS provide an integrated database?
8) Reporting Capability:
   a) Are predefined reports provided with the product?
   b) Are the reports easy to use?
   c) Does the IDS provide built-in, graphical and text based Reports with the ability to create your own reports?
9) Performance

a) How does the product perform under the typical utilization of your network?

10) Architecture
    a) Does the product's design fit into your network? What network topologies are supported? (examples: Ethernet, Fast Ethernet, Token Ring, FDDI)
    b) What operating systems are supported? For example, NT, Solaris, AIX, HP/UX?
    c) Does the Vendor support an open architecture for their IDS offering? Do they provide Software Development Kits for you to extend the use of the IDS product?
    d) Does the Vendor provide an Open API (Application Programming Interface) to leverage existing management solutions?
11) Third Party Support:
    a) Are there any 3rd party integration efforts - such as integration with HP OpenView and/or Tivoli?
12) Correlation Reporting:
    a) Is there integration capability with a decision support reporting application? (SAFEsuite Decisions allows the customer to correlate and consolidate Intrusion Detection data, Vulnerability Assessment Data, Firewall Data, and other security data to provide an overall security assessment picture of their environment.)
13) Technical Support:
    a) What is the quality of Technical Support, do they know the product?
    b) Does the Technical Support Staff understand the functionality of key network components?
    c) Does the Technical Support Staff understand computer Security?
14) Industry Research:
    a) What does Industry Analysts say about the product and the vendor?
    b) What do the reviews written by security experts who have used the product say?
15) Security Services
    a) Does the Company provide professional consulting services that are security experts who can assist you in deployment of your IDS solution?

❖ In addition to some of the Items above, host-based IDS' have some unique criteria to be evaluated on, additional areas to consider are:

1) Incident Detection – Are system specific attacks recognized (i.e. brute-force)?
2) Types of Incident Detection – What are the mechanisms used to detect attacks? Log analysis, Kernel Level Auditing, etc. What kind and which attacks can be recognized?
3) Incident Response – Can system specific attacks be responded to by system specific responses (i.e. disable a user login)?
4) Decoy Services – Does the IDS have the ability to create decoys to occupy and frustrate unwelcome intrusions?

## Testing Consideration

- If the IDS offering has both host and network sensors, take this into account when creating the test environment. They should be evaluated separately and as a solution. Host sensors possess features that are not found in network sensors and vice versa. It is important to understand how each component of the IDS functions.
- Once the configuration of the IDS is determined, implement the configuration in a lab or isolated environment. Ideally this environment will best represent the current network on which the IDS will reside once in production.

- For evaluation of the network sensor, Establish network performance baselines utilizing different levels of realistic network noise or traffic (For example: set test utilization levels of 0%, 10%, 25%, 40%, 50% and 65%). A great temptation here is to generate unrealistically high network noise. Networks operating at greater than 65% utilization will typically have problems dropping even legitimate connections. A network can not be expected to function properly at unreasonably high rates of utilization, neither should an IDS. ➢
- For the evaluation of a host sensor, establish differing levels of system CPU utilization baselines (i.e. 10%, 25%, 40%, 50%, and 65%). ❖
- Utilize a scripted set of attacks for each IDS sensor being tested. This is most easily accomplished utilizing a network monitor. Capture a series of attacks from an assessment tool or from the vendor and then transmit the captured network information back on to the network. Independent tests can be conducted afterwards, however a common set of scripts will ensure unity within the testing phase.
- After setting the utilization baseline in the previous steps from both host and network sensors, use a series of predefined attack scripts to test the IDS. These attack scripts can be obtained off the internet, or usually through the vendor. These scripts should exercise the full range of functionality of both the network and host-based components.
- Appendix A and Appendix B are provided as a courtesy expanding upon the points listed above in the Procedure. Items in this list are intended to provoke thoughts and ideas regarding the implementation of an IDS within you enterprise network.
- Repeat this process for each vendor IDS you are evaluating.

## Results

When this process is completed, you will have a series of documents and notes regarding the performance of each IDS under similar conditions as well as information gathered from informal review of an IDS. A thorough understanding of each IDS and its ability is a ***must***. In addition, you should be able to comprehend the steps required and benefits from implementing a particular IDS within you environment.

## ➢ Appendix A - Network-based Intrusion Detection Requirements

Product Name: _____
Evaluation Date: _____
Evaluated By: _____

### Installation and Deployment

| | | |
|---|---|---|
| ❑ | Product installs quickly and easily. An unskilled operator can install a single engine within 30 minutes. | |
| ❑ | Product engines run on Windows NT hosts as well as Solaris hosts. | |
| ❑ | The installation follows standard installation procedures for Windows NT and UNIX. | |
| ❑ | The installation instructions are adequately documented. | |
| ❑ | Product management consoles run on Windows NT hosts. | |
| ❑ | Product engines fully support 10Base-T, 100Base-T, Token Ring (4Mbps or 16Mbps), and FDDI networks. | |
| ❑ | All software required to operate the product is provided by the vendor. There is no additional third-party software to purchase or install. | |
| ❑ | Product is a software package and does not require the installation of special hardware. | |
| ❑ | Product adds no delay or central points of failure to the network being monitored. | |
| ❑ | Product can be remotely and securely updated easily with new signature updates or full product updates. | |
| ❑ | The installation of the product does not require changes to the network infrastructure or affect the MTBF of the network in any way. | |

### Security

| | | |
|---|---|---|
| ❑ | Product uses separate communications channels for control data and for event data. | |
| ❑ | These communications channels use TCP, are connection oriented, and use ports that can be specified by the network administrator, allowing for simple passage through firewalls | |
| ❑ | The data carried in these communications channels must be | |

| | | |
|---|---|---|
| | encrypted, authenticated, and verified using standard export-approved technologies. | |
| ❑ | Vendor provides adequate instructions for hardening the operating system of the host on which the product is running. | |
| ❑ | The product can be configured in stealth mode, so that it does not betray its presence on the network. In other words the products existence on the monitored network can be concealed. | |
| ❑ | The product system itself is protected against attacks and uses no services on the host that might make it vulnerable to attack. | |
| ❑ | The product console monitors its connections to the engines and will detect when an engine goes off line unexpectedly. | |
| ❑ | Product is capable of using out-of-band communications for either or both of its communications channels. | |

### Incident Detection

| | | |
|---|---|---|
| ❑ | Product monitors the network traffic on the local LAN segment for signs of attack, unauthorized access attempts, and misuse – as uniquely defined by the customer. | |
| ❑ | Product detects incidents that originate from inside the network perimeter, as well as from outside the network perimeter. | |
| ❑ | Product detects incidents based on patterns in network traffic that indicate malicious intent (pattern-based signatures). | |
| ❑ | Product's pattern-based signatures have a strong sense of context, so that false positives are minimized. | |
| ❑ | Product detects incidents based on a number of occurrences over a specified period of time (frequency or threshold-based signatures). | |
| ❑ | Product detects and can be configured to stop Denial of Service attacks. | |
| ❑ | Product detects and can be configured to stop Unauthorized Access Attempts. | |
| ❑ | Product detects and can be configured to stop Pre-Attack Probes. | |
| ❑ | Product detects and can be configured to stop Suspicious | |

| | | |
|---|---|---|
| | Activity. | |
| ❑ | User-specified signatures can be created based upon content; i.e. string matching. | |
| ❑ | Product detects active content on the network, including Java, ActiveX, and Shockwave. | |
| ❑ | Product allows users to modify the engine filtering logic such that it detects incidents related to a subset of the network traffic (specific IP addresses, for example). | |
| ❑ | Product's signatures can be tuned to match the operational requirements of the customer network so that false positives are minimized. | |
| ❑ | Product's Help system describes the incidents in adequate detail, providing sufficient information about: 1) the incident, 2) the potential damage, 3) possible false positives, 4) the systems affected, 5) how to respond immediately upon detection of the incident, and 6) how to remove the vulnerability associated with the incident. | |
| ❑ | Product, when combined with the output of a vulnerability assessment scan, can be configured to focus on the incidents that pose the greatest risk to the customers network. | |
| ❑ | Product detects the attacks and the network misuse that represent risk to the customer. | |

### Incident Response

| | | |
|---|---|---|
| ❑ | Product can send alarms to the management console, or to multiple management consoles, upon detection of an incident. | |
| ❑ | Product can send an SNMP trap to the network or systems management console of your choice upon detection of an incident. | |
| ❑ | Beyond SNMP, product is capable of native integration with popular Network Management | |
| ❑ | Product can notify an administrator via e-mail of an attack or misuse. | |
| ❑ | Product can log a summary of an incident to persistent data storage. | |
| ❑ | Product can record the entire binary content of a network session and write it to persistent data storage for future analysis or forensics. | |

| | | |
|---|---|---|
| ❑ | Product can copy the entire binary content of a network session up to the management console in real time so that it may be viewed as it is happening. | |
| ❑ | Product can terminate a TCP session by issuing TCP Reset packets to each end of the connection. | |
| ❑ | Product can prevent TCP, UDP, and IP access to a network by automatically reconfiguring a firewall or router to prevent certain traffic from crossing the firewall boundary for a user-specified period of time. | |
| ❑ | Product can respond to an incident by executing one or more user-specified programs. These can be batch files, command line scripts, executables, etc. | |

### Configuration

| | | |
|---|---|---|
| ❑ | Remote product engines can be configured from the management console using a point-and click-interface. | |
| ❑ | Product provides configuration templates that describe an engine configuration (i.e., active pre-defined signatures, and responses). These templates can be customized, applied to many engines at the same time, saved for future use, and exchanged among management domains. | |
| ❑ | Product's help system is integrated into the console's policy configuration, providing a detailed description of the attack signature that is selected. | |
| ❑ | The priority level for each pre-defined signature can be configured from the management console. | |
| ❑ | The interface allows attack signatures to be activated or deactivated via check-box selection. | |
| ❑ | The administrator, from the management console, can specify the response to each pre-defined event. | |
| ❑ | The administrator from the management console can specify multiple responses to each attack or misuse. | |
| ❑ | The pre-defined signatures can be tuned such that false positives are minimized. | |
| ❑ | Product includes the ability to tune event propagation. | |
| ❑ | Product can be configured such that attack signature and traffic analysis focus only on specified hosts, specified protocols, or specified services. | |

| | | |
|---|---|---|
| ❑ | New Services (as defined by TCP/IP port number) can be specified by the administrator. New signatures can then be based upon that new, user-defined Service. | |

### Event Monitoring

| | | |
|---|---|---|
| ❑ | Product graphically depicts both suspicious activity and normal network activity. | |
| ❑ | The graphical interface can be used effectively by NOC operators with limited experience and training and requires no special technical knowledge. | |
| ❑ | The graphical interface uses an iconic display to alert operators to important occurrences. Varying shapes and colors (red, yellow, and green) are used to guide problem resolution. | |
| ❑ | The graphical interface can display summary information sorted by source address (initiator), destination address (target), or event type. | |
| ❑ | The graphical interface supports a "drill down" mechanism so that the operator may obtain additional information about an event. This information includes action(s) that were taken by product in response to the event. | |
| ❑ | The graphical interface consolidates multiple event occurrences into a single alarm. | |
| ❑ | Events from any product engine can be monitored from a single, authorized management console. | |
| ❑ | A single product engine can report attack and misuse data to multiple management consoles simultaneously. | |
| ❑ | An HP OpenView management console can receive product alarm data. | |
| ❑ | A Tivoli TME 10 NetView management console can receive product alarm data. | |
| ❑ | Operator can launch the product from the Tivoli TME 10 NetView management console, or from an HP OpenView management console. | |

### Data Management

| | | |
|---|---|---|
| ❑ | Data from many product engines is assimilated on a single management console. This includes event summary data as well as the binary content of logged sessions. | |

| | | |
|---|---|---|
| ❑ | Data on the management console is stored in an ODBC database. This database is built-in and requires no installation of third-party software. | |
| ❑ | The ODBC database can be exported to a database of your choice or to a delineated text file. | |
| ❑ | The database structure is completely open and is published in the product documentation. Third-party management tools can easily access this database, if desired. | |
| ❑ | Product provides built-in report generation capability. | |
| ❑ | Product provides at least 12 pre-defined reports. | |
| ❑ | The vendor includes the templates for the pre-defined reports, so that custom reports can be generated using the standards reports as a starting point. | |
| ❑ | Product provides multiple forms of reporting suitable for all technical levels. | |
| ❑ | Product reports are configurable and customizable. | |
| ❑ | Product reports can be exported to different formats, such as CSV or a Word document. | |
| ❑ | Product's data management capabilities provide critical information required for risk assessment and decision-making. | |

### Performance

| | | |
|---|---|---|
| ❑ | Product engines can monitor network traffic and take action autonomously, without a console running. | |
| ❑ | Product can process network traffic at a rate that is acceptable to you with all of the attack signatures active. | |
| ❑ | Product's performance scales well with the number of attack signatures and filters active. Increasing the number of predefined or custom signatures does not significantly impact the performance of the system. | |
| ❑ | Product handles traffic bursts gracefully, switching to sampling mode until the traffic levels return to a consistent level. | |
| ❑ | Product engines function effectively on an Intel-based system with a 266 MHz Pentium II processor and 128MB RAM or a SPARC-based system with 128MB RAM. | |

## Architecture

| | | |
|---|---|---|
| ❑ | Product's architecture adapts well to higher network speeds and switched network topologies. | |
| ❑ | Product's architecture allow the attack recognition and response modules to be integrated into other network devices, such as firewalls and switches. | |
| ❑ | Product's architecture allows for the use of off-the-shelf components, significantly reducing the cost of the product’s deployment. | |
| ❑ | Product's architecture allows for the capability of remotely and securely updating installed sensor base. | |

## Product Updates, Technical Support, and Industry Research

| | | |
|---|---|---|
| ❑ | Vendor updates its attack signature database at least, monthly. | |
| ❑ | Vendor provides major new major product releases at least two times per year. | |
| ❑ | Vendor notifies you automatically through e-mail about the availability of new signatures and new product releases. | |
| ❑ | Vendor makes new attack signatures and new major software releases available for download from its Web site. | |
| ❑ | Vendor provides technical support via telephone, e-mail, and fax to you during business hours in your time zone (i.e. 8 a.m. to 8 p.m. eastern time. Monday through Friday). | |
| ❑ | Vendor provides 24-hour technical support or it is available to those that require it. | |
| ❑ | Vendor supports a research and development team for compiling and understanding new attack signatures and new system vulnerabilities. | |
| ❑ | Vendor notifies the industry about newly discovered attack signatures and system vulnerabilities periodically through an e-mail service. | |

## Other

| | | |
|---|---|---|
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |

## ❖ Appendix B - Host-based Intrusion Detection Requirements

Product Name: _____
Evaluation Date: _____
Evaluated By: _____

### Installation and Deployment

| | | |
|---|---|---|
| ❑ | Product installs quickly and easily. An unskilled operator can install a single detector within 15 minutes. | |
| ❑ | The installation follows standard operating system installation procedures. | |
| ❑ | The installation instructions are well documented, and are made available to the customer in both hard and soft form. | |
| ❑ | Product management consoles run on Windows NT and Solaris hosts. | |
| ❑ | Product engine supports any TCP/IP-based network, including 10Base-T, 100Base-T, Token Ring (4Mbps or 16Mbps), and FDDI. | |
| ❑ | All software required to operate the product is provided by the vendor. There is no additional third-party software to purchase or install. | |
| ❑ | Product is a software package and does not require the installation of special hardware. | |
| ❑ | Product adds no delay or central points of failure to the network being monitored. | |
| ❑ | The installation of the product does not require changes to the network infrastructure or affect the MTBF of the network in any way. | |

### Security

| | | |
|---|---|---|
| ❑ | Product uses separate communications channels, between the console and the distributed engines, for control data and for event data. | |
| ❑ | These communications channels use TCP, are connection oriented, and use ports that can be specified by the network administrator, allowing for simple passage through firewalls. | |
| ❑ | The data carried in these communications channels is encrypted, authenticated, and verified using standard export- | |

| | | |
|---|---|---|
| | approved technologies. | |
| ❑ | Vendor provides adequate instructions for hardening the operating system of the host on which the product is running. | |
| ❑ | The product can be configured so that it does not betray its presence to an intruder. | |
| ❑ | The product system itself is protected against attacks and uses no services on the host that might make it vulnerable to attack. | |
| ❑ | The product console monitors its connections to the engines and will detect when an engine goes off line unexpectedly. | |

### Incident Detection

| | | |
|---|---|---|
| ❑ | Product detects incidents that originate from inside the network perimeter, as well as from outside the network perimeter. | |
| ❑ | Product detects incidents based on entries and patterns in log files that indicate malicious intent (pattern-based signatures). | |
| ❑ | Product's signatures have a strong sense of context, so that false positives are minimized. | |
| ❑ | Product detects incidents based on a number of occurrences over a specified period of time (frequency or threshold-based signatures). | |
| ❑ | Product detects and can be configured to stop Unauthorized Access Attempts. | |
| ❑ | Product detects and can be configured to stop Pre-Attack Probes. | |
| ❑ | Product detects and can be configured to stop Suspicious Activity. | |
| ❑ | Product's signatures can be tuned to match the nuances, and operational requirements of the customer critical systems, so that false positives are minimized. | |
| ❑ | Product can detect brute force login attempts made against multiple TCP/IP services. | |
| ❑ | Product can detect modifications made to key files and the termination of executables. | |
| ❑ | Product can detect the questionable use of Administrator privileges, such as changes made to security policy, | |

| | | |
|---|---|---|
| | add/removal of trusted domain, and clearing of the security log. | |
| ❑ | User-specified signatures can be created based upon content; i.e. string matching. | |
| ❑ | User-specified signatures can be created based upon the frequency of events. | |
| ❑ | User-specified signatures can be based upon the correlation of two or more events. | |
| ❑ | Product's correlation signatures can be based upon combinations of attacks; deviations from policy, specific log events, and decoy service access. | |
| ❑ | Product's Help system: 1) describes the incidents in adequate detail, 2) provides sufficient information about the incident, 3) the potential damage, 4) possible false positives, 5) the systems affected, 6) how to respond immediately upon detection of the incident, and 7) how to remove the vulnerability associated with the incident. | |
| ❑ | Product, when combined with the output of a vulnerability assessment scan, can be configured to focus on the incidents that pose the greatest risk to the customers network. | |
| ❑ | Product detects the attacks and the network misuse that represent risk to the customer. | |
| ❑ | Product monitors the host's log files for signs of attack, unauthorized access attempts, and misuse – as uniquely defined by the customer. | |
| ❑ | Product monitors all three Windows NT logs - security, application, and system. | |
| ❑ | Product will watch for instances where there is a deviation from the security policy that was established for that host. | |
| ❑ | Product can be configured to monitor remote syslog messages from UNIX hosts. | |
| ❑ | Product watches for matches against NT message IDs and regular expressions. | |
| ❑ | Product has the ability to detect that an intruder has compromised a host. | |

## Decoy Services

| | | |
|---|---|---|
| ❑ | Product can simulate a TCP/IP service on an unused port and make it appear active. | |
| ❑ | The console can specify the port to be decoyed, its timeout parameters, and the selected response(s) for that port. | |
| ❑ | It is possible to turn a decoy service on or off from the management console. | |
| ❑ | Product will allow up to 24 ports to be decoyed. | |

## Incident Response

| | | |
|---|---|---|
| ❑ | Product can send alarms to the management console, or to multiple management consoles, upon detection of an incident. | |
| ❑ | The management console has fully integrated capability to manage and receive alerts from BOTH host and network-based intrusion detection software. | |
| ❑ | The product can be configured to initiate a single response, or multiple responses, to an event directed at the host | |
| ❑ | Product can send an SNMP trap to the network or systems management console. | |
| ❑ | Product can notify an administrator via e-mail of an attack or misuse. | |
| ❑ | Product can write a message into the Windows NT application log. | |
| ❑ | Product can log a summary of an incident to persistent data storage. | |
| ❑ | Product can respond with a user-definable banner. | |
| ❑ | Product can terminate the intruder's login in response to an attack or misuse. | |
| ❑ | Product can break the host's connection to the network. | |
| ❑ | Product can respond to an incident by executing one or more user-specified programs. These can be batch files, command line scripts, executables, etc. | |

## Configuration

| | | |
|---|---|---|
| ❑ | The product can be configured from the central management console using a point-and click graphical user interface. | |
| ❑ | The management console has fully integrated capability to configure both host and network-based intrusion detection engines. | |
| ❑ | Product provides configuration templates that describe a host configuration (signatures, responses, etc.). These templates can be customized, applied to many engines at the same time, saved for future use, and exchanged among management domains. | |
| ❑ | Product's help system is integrated into the console's policy configuration, providing a detailed description of the attack signature that is selected. | |
| ❑ | The user from the management console can configure the priority level for each pre-defined signature. | |
| ❑ | The interface allows attack signatures to be activated or deactivated via check-box selection. | |
| ❑ | The administrator from the management console can specify the response(s) to each pre-defined event. | |
| ❑ | The administrator from the management console, if necessary, to each attack or misuse can specify multiple responses. | |
| ❑ | The pre-defined signatures can be tuned such that false positives are minimized. | |

## Event Monitoring

| | | |
|---|---|---|
| ❑ | Product graphically depicts both suspicious activity and normal network activity. | |
| ❑ | The graphical interface can be used effectively by NOC operators with limited experience and training and requires no special technical knowledge. | |
| ❑ | The graphical interface uses an iconic display to alert operators to important occurrences. Varying shapes and colors (red, yellow, and green) are used to guide problem resolution. | |
| ❑ | The graphical interface can display summary information sorted by source address (initiator), destination address | |

| | | |
|---|---|---|
| | (target), or event type. | |
| ❑ | The graphical interface supports a "drill down" mechanism so that the operator may obtain additional information about an event. This information includes action(s) that were taken by product in response to the event. | |
| ❑ | The graphical interface consolidates multiple event occurrences into a single alarm. | |
| ❑ | Events from any product engine can be monitored from a single, authorized management console. | |
| ❑ | A single product engine can report attack and misuse data to multiple management consoles simultaneously. | |
| ❑ | An HP OpenView management console can receive product alarm data. | |
| ❑ | A Tivoli TME 10 NetView management console can receive product alarm data. | |

### Data Management

| | | |
|---|---|---|
| ❑ | Data from many product engines is assimilated on a single management console. This includes event summary data as well as the binary content of logged sessions. | |
| ❑ | Data on the management console is stored in an ODBC database. This database is built-in and requires no installation of third-party software. | |
| ❑ | The ODBC database integrates host and network-based intrusion detection events, allowing for combined reporting. | |
| ❑ | The ODBC database can be exported to a database of your choice or to a delineated text file. | |
| ❑ | The database structure is completely open and is published in the product documentation. Third-party management tools can easily access this database, if desired. | |
| ❑ | Product provides built-in report generation capability. | |
| ❑ | The vendor includes the templates for the pre-defined reports, so that custom reports can be generated using the standards reports as a starting point. | |
| ❑ | Product provides multiple forms of reporting suitable for all technical levels. | |
| ❑ | Product reports are configurable and customizable. | |

| | | |
|---|---|---|
| ❑ | Product's data management capabilities provide critical information required for risk assessment and decision-making. | |

### Performance

| | | |
|---|---|---|
| ❑ | Product engines can monitor network traffic and take action autonomously, without a console running. | |
| ❑ | Product can process host traffic at a rate that is acceptable to you with all of the attack signatures active. | |
| ❑ | Product's performance scales well with the number of attack signatures and filters active. Increasing the number of predefined or custom signatures does not significantly impact the performance of the system. | |

### Product Updates, Technical Support, and Industry Research

| | | |
|---|---|---|
| ❑ | Vendor updates its attack signature database at least monthly. | |
| ❑ | Vendor provides major new major product releases at least two times per year. | |
| ❑ | Vendor notifies you automatically through e-mail about the availability of new signatures and new product releases. | |
| ❑ | Vendor makes new attack signatures and new major software releases available for download from its Web site. | |
| ❑ | Vendor provides technical support via telephone, e-mail, and fax to you during business hours in your time zone (i.e. 8 a.m. to 8 p.m. e.s.t. Monday through Friday). | |
| ❑ | Vendor provides 24-hour technical support or it is available to those that require it. | |
| ❑ | Vendor supports a research and development team for compiling and understanding new attack signatures and new system vulnerabilities. | |
| ❑ | Vendor notifies the industry about newly discovered attack signatures and system vulnerabilities periodically through an e-mail service. | |

## Other

| | | |
|---|---|---|
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |
| ❑ | | |