

IBM i
Version 7.2

Security
Intrusion detection



Note

Before using this information and the product it supports, read the information in [“Notices” on page 31](#).

This edition applies to IBM i 7.2 (product number 5770-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

This document may contain references to Licensed Internal Code. Licensed Internal Code is Machine Code and is licensed to you under the terms of the IBM License Agreement for Machine Code.

© **Copyright International Business Machines Corporation 2002, 2013.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Intrusion detection.....	1
What's new for IBM i 7.2.....	1
PDF file for Intrusion detection.....	1
Intrusion detection concepts.....	2
Intrusion detection system initialization.....	2
Intrusion detection system operation.....	4
Intrusion detection and prevention.....	6
Real-time intrusion and extrusion detection notification.....	7
Intrusion and extrusion types.....	7
Using the Intrusion Detection System GUI.....	12
Using the IDS GUI in IBM Navigator for i.....	12
Setting up intrusion detection system properties.....	13
Setting up e-mail and message notification.....	13
Starting the intrusion detection system.....	14
Stopping the intrusion detection system.....	14
Creating intrusion detection policies.....	15
Creating a set of default intrusion detection policies.....	15
Creating an attack policy.....	15
Creating a scan policy.....	16
Creating a traffic regulation policy.....	16
Creating an intrusion detection policy based on another policy.....	17
Managing intrusion detection policies.....	18
Changing an intrusion detection policy.....	18
Changing the priority of an intrusion detection policy.....	19
Deleting an intrusion detection policy.....	19
Enabling an intrusion detection policy.....	19
Disabling an intrusion detection policy.....	20
Backing up the intrusion detection policy file.....	20
Writing intrusion detection programs.....	21
Displaying intrusion detection events.....	21
Filtering intrusion detection events.....	22
Intrusion monitor audit record entries.....	22
Examples: Intrusion detection.....	25
Example: Traffic regulation policy.....	25
Example: Restricted IP options policy.....	25
Example: Perpetual echo policy.....	26
Example: E-mail notification.....	27
Example: Intrusion detection scan policy.....	27
Example: Variable dynamic throttling for scan events.....	28
Example: Variable dynamic throttling for traffic regulation events.....	29
Related information for Intrusion detection.....	30
Notices.....	31
Programming interface information.....	32
Trademarks.....	32
Terms and conditions.....	33

Intrusion detection

The intrusion detection and prevention system (IDS) notifies you of attempts to hack into, disrupt, or deny service to the system. IDS also monitors for potential extrusions, where your system might be used as the source of the attack. These potential intrusions and extrusions are logged as intrusion monitor audit records in the security audit journal and displayed as intrusion events in the Intrusion Detection System graphical user interface (GUI). You can configure IDS to prevent intrusions and extrusions from occurring.

Important: The term *intrusion detection* is used two ways in IBM® i documentation. In the first sense, intrusion detection refers to the prevention and detection of security exposures. For example, a hacker might be trying to break into the system using a user ID that is not valid, or an inexperienced user with too much authority might be altering important objects in system libraries. In the second sense, intrusion detection refers to the intrusion detection function that uses policies to monitor suspicious traffic on the system.

Intrusion detection involves gathering information about attacks arriving over the TCP/IP network. *Intrusions* encompass many undesirable activities, such as information theft and denial of service attacks. The objective of an intrusion might be to acquire information that a person is not authorized to have (information theft). The objective might be to cause a business harm by rendering a network, system, or application unusable (denial of service), or it might be to gain unauthorized use of a system as a means for further intrusions elsewhere. Most intrusions follow a pattern of information gathering, attempted access, and then destructive attacks. Some attacks can be detected and neutralized by the target system. Other attacks cannot be effectively neutralized by the target system. Most of the attacks also make use of *spoofed* packets, which are not easily traceable to their true origin. Many attacks make use of unwitting accomplices, which are machines or networks that are used without authorization to hide the identity of the attacker. For these reasons, a vital part of intrusion detection is gathering information, and detecting and preventing system attacks.

The IDS GUI allows you to configure and manage intrusion detection policies, and start and stop IDS. You no longer have to edit the IDS policy configuration file directly. You can use the IDS GUI to display the intrusion events that have been logged in the audit journal. Security administrators can analyze the audit records that IDS provides to secure the network from these types of attacks. In addition, you can use the IDS GUI to manage IDS on your IBM i systems.

IDS does not monitor for viruses, Trojan horse programs, or malicious e-mail attachments.



What's new for IBM i 7.2

Read about new or significantly changed information for the Intrusion detection topic collection.

Miscellaneous updates have been made since the previous publication.

How to see what's new or changed

To help you see where technical changes have been made, the information center uses:

- The  image to mark where new or changed information begins.
- The  image to mark where new or changed information ends.

In PDF files, you might see revision bars (|) in the left margin of new and changed information.

To find other information about what's new or changed this release, see the [Memo to users](#).

PDF file for Intrusion detection

You can view and print a PDF file of the intrusion detection information.

To view or download the PDF version of this document, select [Intrusion detection](#).

You can view or download these related topic PDFs:


- [Planning and setting up system security](#), which discusses techniques for detecting other types of intrusions.

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF link in your browser.
2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you want to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

You need Adobe Reader installed on your system to view or print these PDFs. You can download a free copy from the [Adobe Web site](http://www.adobe.com/products/acrobat/readstep.html) (www.adobe.com/products/acrobat/readstep.html) .

Intrusion detection concepts

An *intrusion detection policy* defines the parameters that the Intrusion Detection System (IDS) uses to monitor for potential intrusions and extrusions on the system. If a potential intrusion or extrusion is detected, an *intrusion event* is logged in an intrusion monitor record in the security audit journal.

Before IDS can monitor for potential intrusions, you need to use the Intrusion Detection System GUI to create a set of intrusion detection policies that cover various types of intrusions. Once the intrusion detection policies have been created and IDS has been started, the TCP/IP stack detects potential intrusions and extrusions based on those policies.

You can create any of the following:

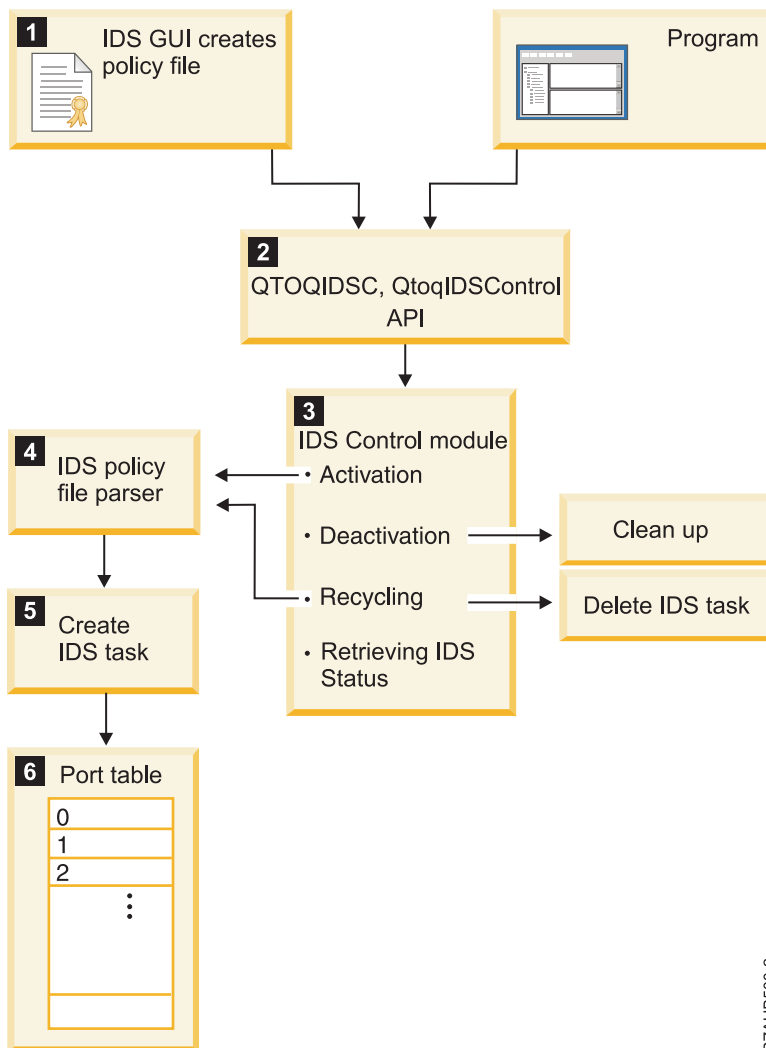
- A set of default intrusion detection policies that monitor the entire system. Your default policies include attack, scan, and traffic regulation policies.
- Attack policies.
- Scan policies.
- Traffic regulation policies.

Use the **Intrusion Detection Events** page to display the intrusion events that have been logged on the system, as well as to view details about each event.

Intrusion detection system initialization

If the intrusion detection system (IDS) is active, it monitors intrusions when the system is IPLed as well as when the system is running. When you use the IDS GUI to create intrusion detection policies, IDS creates a set of conditions and actions based on the information in the policies.

The following graphic shows how IDS is initialized when you create an intrusion detection policy using the IDS GUI or a program.



RZAU500-2

1. When you create an intrusion detection policy, the IDS GUI builds the IDS policy file and activates IDS using the Control Intrusion Detection and Prevention (QTOQIDSC, QtoqIDSCControl) API.

Note: After you create a new policy, IDS is automatically stopped and restarted for the policy to take effect.

2. The QTOQIDSC API sends the policy information to the IDS control module.

3. The IDS control module has four functions:

- Starting IDS. If IDS is started or recycled, IDS control reads the policy file and sends it to the IDS policy file parser.
- Stopping IDS. If IDS is stopped, IDS control performs internal cleanup functions.
- Recycling (stopping and restarting) IDS. If you delete an IDS policy, IDS control deletes the IDS task associated with that policy.
- Retrieving IDS status. This status indicates whether IDS is stopped or active.

4. The IDS policy file parser creates the IDS task.

5. The IDS task creates the port table with the condition and action lists.

6. The IDS port table represents TCP ports 1 through 65 535. This table also has a port 0 provision which applies to all ports. Conditions are assigned to ports using the IDS GUI. Actions are assigned to conditions using the IDS GUI.

Intrusion detection system operation

While IDS is active, it reports the suspected intrusions and extrusions that are defined by the enabled IDS policies. The production and service stacks detect these intrusions and extrusions. When an intrusion or extrusion event occurs that exceeds user-defined or default thresholds, IDS writes an intrusion monitor record to the audit journal, and optionally sends a notification to a message queue and e-mail message.

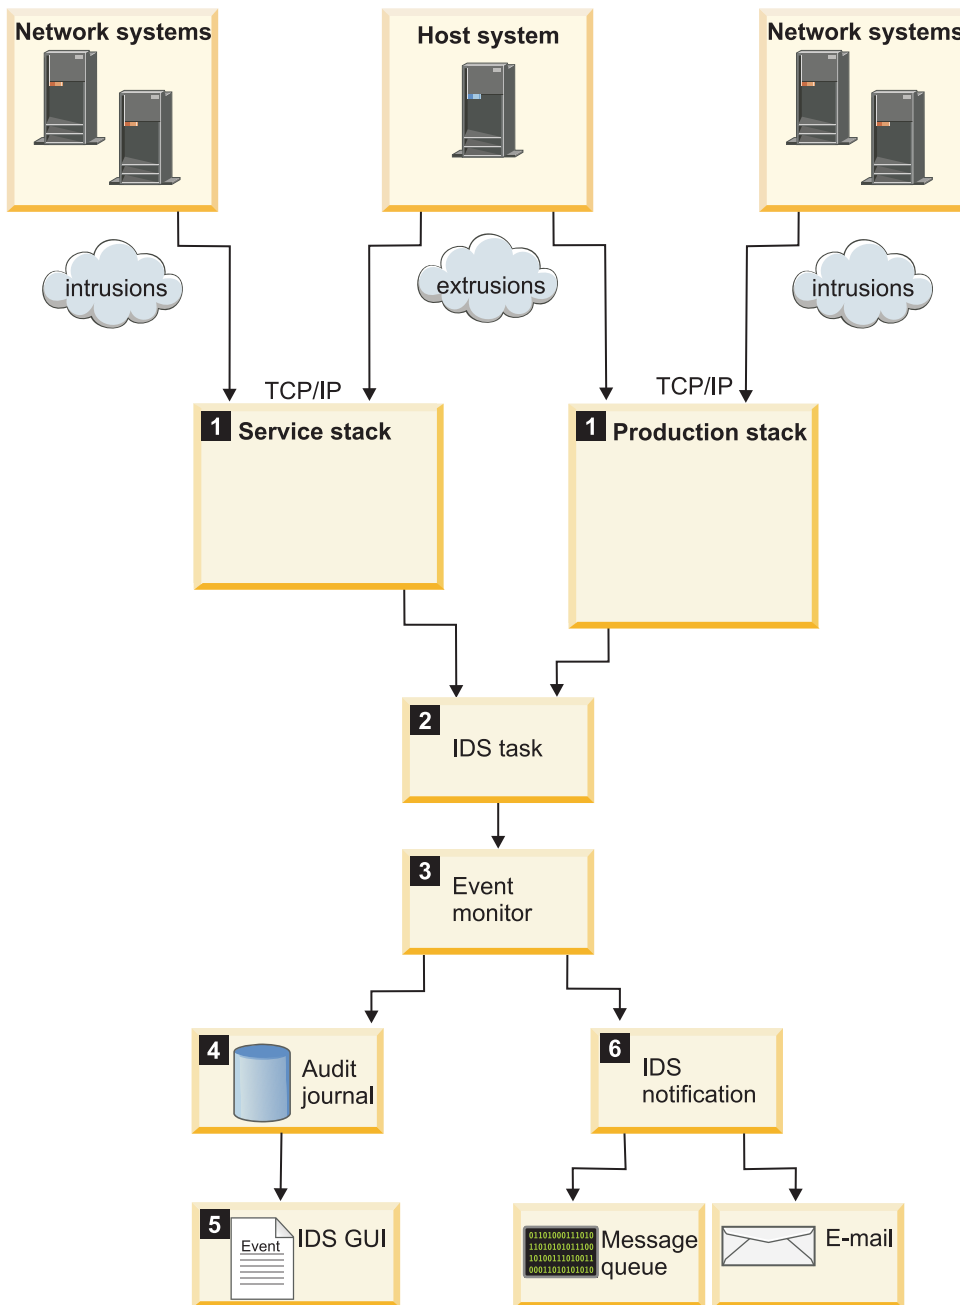
The *production stack* consists of the TCP/IP modules involved in most of the network operations on the IBM i platform. The *service stack* consists of the TCP/IP modules involved in service and support of the IBM i platform.

The service stack comes up first and remains until the next IPL. The production stack comes up after the service stack and remains until TCP/IP is ended. After an IPL, the service stack checks to see if IDS was active before the IPL. If so, IDS is reactivated. Any intrusions and extrusions that are detected by the service stack are logged either by VLOG or Intrusion Monitor records. At this stage, IDS does not send notification to a message queue or an e-mail address. Once the policy file is available, both stacks work with IDS in the same way.

The TCP, UDP, and IP support within the stack, detect the potentially malicious situation. Even if you do not have any intrusion detection policies defined, the service stack detects certain types of intrusions, such as traffic regulation or scan events, using a set of default values. When you define a set of intrusion detection policies, the production stack starts checking for potential intrusions.

The service stack detects only IPv4 intrusions and extrusions, while the production stack detects both IPv4 and IPv6 intrusions and extrusions.

The following graphic shows how IDS detects and reports suspected intrusions and extrusions.



1. When the production or service stack detects a suspected intrusion or extrusion, it sends an event to the IDS task.
2. The IDS task takes each event off the queue one at a time, and matches each event with a condition (from the port table). The IDS task also keeps statistics about the intrusion and extrusion events.
3. IDS signals events for intrusions and extrusions that exceed set thresholds in the policy files.
4. If an event is signaled, the intrusion monitor record is created in the audit journal.
5. The IDS GUI displays the intrusion events from the intrusion monitor audit records.
6. If you have set up e-mail and message notification on the IDS Properties page, IDS notification sends an e-mail to the specified e-mail addresses and a message to a message queue.

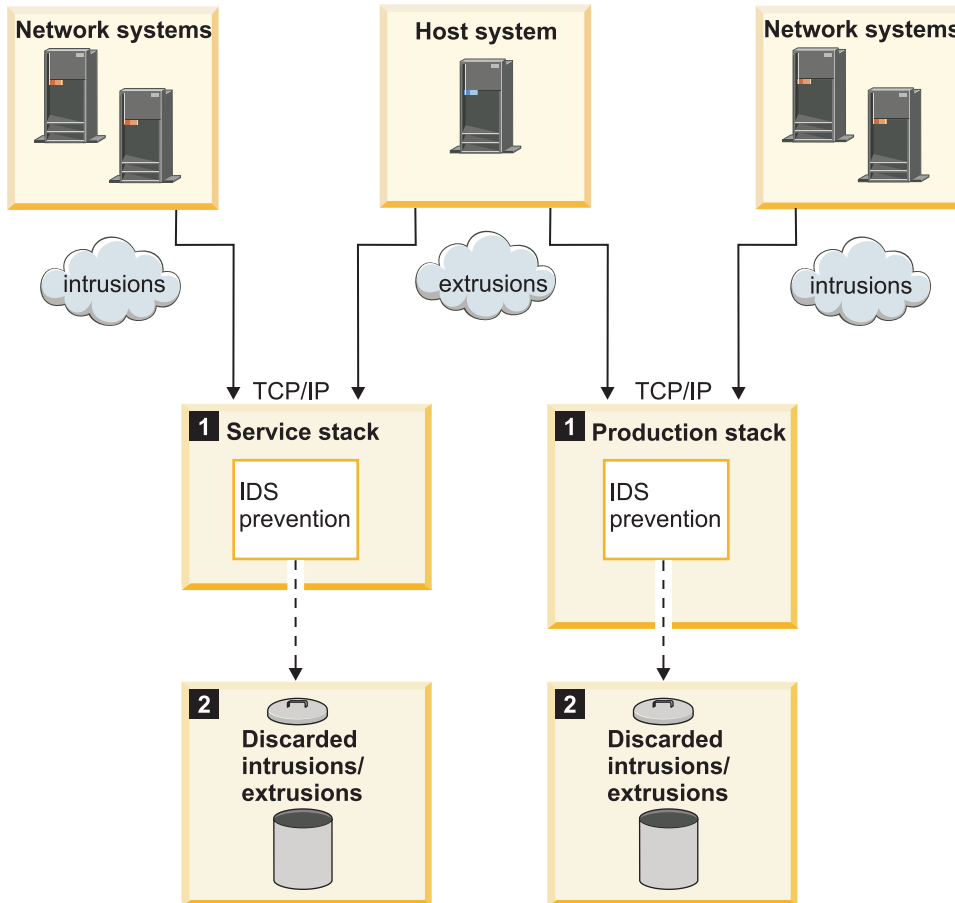
You can analyze the intrusion events to determine which security actions to take. For example, you could end the interface from which the intrusion originated, or use techniques such as variable dynamic throttling to limit or prevent the intrusions from occurring.

Intrusion detection and prevention

You can use the intrusion detection system to prevent intrusions and extrusions from occurring.

Intrusion prevention is a system that attempts to deny potentially malicious activity. The denial mechanisms could involve filtering packets, variable dynamic throttling, or using Quality of Service (QoS) to vary connection rates and burst limits.

The following graphic shows how IDS detects and prevents intrusions and extrusions from occurring.



RZAUB502-0

1. The TCP/IP service and production stacks detect intrusions from systems in the network and extrusions from the host system.
2. If you have variable dynamic throttling enabled, IDS limits or discards the intrusion or extrusion.

You can configure variable dynamic throttling for each IDS policy. Throttling detects all types of intrusions and extrusions. Variable dynamic throttling is a prevention method that automatically starts if certain intrusion event thresholds are met. Throttling stays active until thresholds are no longer exceeded for an interval of time. You can choose to throttle network traffic from all or specific ports and IP addresses. You also can specify the slow and fast scan thresholds, the maximum event message thresholds, or use the default values for those thresholds in your IDS policies. Throttling is activated once a threshold for that policy has been exceeded, and stays active for either a user-defined or system-defined time interval. If the threshold is exceeded at any time during the interval, the throttling is increased immediately and the time interval is reset. Throttling could eventually lead to denying all packets from a given interface. This process continues until the number of offending packets no longer exceeds the thresholds for an entire time interval. When the number of packets drop below the thresholds, throttling is deactivated and normal packet flow resumes.

You also can specify in the ICMP tab of the IDS Properties page whether or not to allow Internet Control Message Protocol (ICMP) redirect messages. *ICMP* is a protocol that is used to send error or informational messages. ICMP is used by some utilities, such as **traceroute**, and the **ping** tool to determine if a host is

reachable. Examples of ICMP messages include: echo replies, echo requests, redirect, destination unreachable, and time exceeded.

Related concepts

Variable dynamic throttling

You can specify *variable dynamic throttling* in each intrusion detection (IDS) policy. If an enabled IDS policy has throttling specified, throttling occurs after a suspected intrusion or extrusion has occurred and certain thresholds have been reached. Variable dynamic throttling starts discarding packets when a threshold has been exceeded within a given statistics or scan interval.

Real-time intrusion and extrusion detection notification

IDS is a notification system. You can configure IDS to send real-time intrusion notifications as messages to a message queue and as e-mail. That way, you can alert systems administrators about specific types of intrusions and extrusions so that they can take appropriate actions.

Use the Notification tab on the IDS Properties page to set up e-mail and message notification. You can send e-mail to a maximum of three e-mail addresses, to a message queue, or to both places. You also can enable or disable e-mail notification for individual intrusion detection policies.

IDS notification generates e-mail using the following format:

- The **Sender** line specifies `qsys@system_name`, which is the name of the system where the intrusion was detected.
- The **Subject** line summarizes the type of intrusion or extrusion that was detected on that system.
- The body of the e-mail describes the intrusion or extrusion in detail.

If you (the system administrator) determine that an attack is underway, you can take the appropriate steps to prevent further attacks.

Related tasks

Setting up e-mail and message notification

IDS is a notification system. You can optionally configure IDS to send real-time intrusion notifications to a message queue and to specific e-mail addresses. That way, you can alert systems administrators about specific types of intrusions and extrusions so that they can take actions to stop further intrusions from occurring. You can enable or disable IDS e-mail and message notification per policy.

Intrusion and extrusion types

The intrusion detection system detects many types of intrusion and extrusion events.

- Address poisoning
- Fraggle attack
- Internet Control Message Protocol (ICMP) redirect message
- Internet Protocol (IP) fragment
- Malformed packet
- Outbound raw
- Perpetual echo on User Datagram Protocol (UDP) ports
- Ping-of-death attack
- Restricted IP option
- Restricted IP protocol (intrusions only)
- Slow and fast scans
- Smurf attack
- SYN floods
- TCP ACK storm
- Traffic regulation conditions

Attack events

Attack policies monitor for various types of attacks against the system. Your system can be attacked or used as the source of an attack. When IDS detects an attack, it writes an intrusion event to the audit record.

For example, an intruder might attempt to cause a system to crash or hang, tie up system resources and deny services, slip through a firewall, or gain back-door entry to a system. The intrusion detection system detects the following types of attack events:

Address poisoning

A hacking technique that redirects data to a different system (for snooping packets) or to nonexistent addresses. Address poisoning also is called Address Resolution Protocol (ARP) spoofing in IPv4 and Neighbor Discovery spoofing in IPv6. IDS is notified whenever the ARP cache or Neighbor Discovery cache changes.

Fraggle attack

A type of denial-of-service attack in which User Datagram Protocol (UDP) echo requests are sent to a broadcast or multicast address, with the source address spoofed as the victim's address. The target port of a fraggle attack is echo port 7. The attacker's purpose is to overload a system with a high volume of traffic, as each host on the network replies to each broadcast or multicast packet that the attacker sends. Spoofing the source address makes the receiver of the multiple responses the victim of the denial-of-service (DoS) attack. (A *denial-of-service attack* is an assault on a network that brings down one or more hosts on a network such that the host is unable to perform its functions properly.)

IDS is notified when a UDP echo request is received to determine if the destination address is an IP broadcast or multicast address. If the destination address is a broadcast or multicast address, IDS signals the attack.

ICMP redirect message

An out-of-band message that is designed to inform a host of a more optimal route through a network, but possibly used maliciously for attacks that redirect traffic to a specific system. In this type of an attack, the hacker, posing as a router, sends an Internet Control Message Protocol (ICMP) redirect message to a host, which indicates that all future traffic must be directed to a specific system as the more optimal route for the destination. You can set up IDS to notify you when these ICMP redirect messages occur or to ignore them.

IP fragment

An Internet Protocol (IP) datagram that contains only a portion of the user data from a larger IP datagram. In an attack, the IP fragment might be less than 576 bytes in length, or have an offset of less than 256 bytes. When the IP fragments are too small, the intent might be a malicious attempt to slip through a firewall, but it could just be a normal case of packet retransmission. IDS detects IP fragments that are suspicious.

Malformed packet

A packet that does not conform to TCP/IP standards for size, destination, or flags in the TCP header. The intent might be to crash or to hang a system. IDS also checks for restricted IP protocols and options in a

malformed packet attack. The TCP/IP stack notifies IDS of these malformed packets and usually discards them.

Outbound raw attack

An outbound packet that uses a nonstandard protocol. Outbound packets are a type of extrusion. Outbound restricted IP protocols are covered under outbound raw attacks. Standard protocols include TCP, UDP, ICMP, ICMPv6, Internet Group Management Protocol (IGMP), or Open Shortest Path First (OSPF).

Perpetual echo

A denial-of-service attack on the UDP echo port 7. If the source port and target port are set to port 7, the request is echoed back and forth. An attacker sends a UDP echo request to an IP broadcast or multicast address and provides a spoofed source address for all the targets to echo back responses. The spoofed source address, which is not the hacker's address, becomes the victim of a potentially large amount of network traffic. A perpetual echo can be an intrusion or extrusion.

Ping-of-death

An attack that involves sending a ping packet that is larger than the maximum IP packet size of 65 536 bytes, which can overload a system.

Restricted IP option

An IP option, such as Loose Source and Record Route (LSRR), that is used to map a network's topology and discover private IP addresses. A hacker might try to use restricted IP options to get through firewalls. You can use the IDS policy to restrict which IP options an inbound or outbound packet can contain. A restricted IP option can be an intrusion or extrusion.

Restricted IP protocol

An unrecognized protocol that can be used to establish an attack on a network. An IP protocol other than ICMPv6, ICMP, IGMP, TCP, or UDP is an unrecognized protocol. A hacker might program directly to a raw socket without going through the TCP/IP programming interface. IDS is notified of the potential intrusion by classifying it as a restricted protocol attack. If there is no corresponding IDS policy for restricted protocols, the notification goes unrecorded. Non-mainstream outbound protocols are covered under outbound raw attacks.

Open Shortest Path First (OSPF) is an interior gateway protocol that is used to send information to routers regarding the shortest path to each node in a network. Unlike the other well-known protocols that IDS is not notified about, IDS is notified about inbound packets that contain the OSPF protocol with a "restricted protocol" attack. If networks in the system are using OSPF, consider excluding OSPF from the range of protocols to restrict. OSPF might display in the audit journal rather frequently if it is included in the restricted protocol range in the policy. If you receive an intrusion notification about the OSPF protocol, review the information to determine whether the system is using OSPF for legitimate purposes.

Smurf

A denial-of-service attack in which a spoofed source address is flooded with echo replies. The replies are caused when many ping (ICMP echo) requests using the spoofed source address are sent to one or more broadcast or multicast addresses.

SYN floods

A type of denial-of-service attack in which an attacker sends a large number of TCP connection requests to a target computer, without answering the target computer's acknowledgment requests. The target computer becomes overloaded and denies service to legitimate users.

TCP ACK storm

A denial-of-service attack on a server where a hacker or cracker secretly inserts data into a client/server session in an attempt to disrupt the session. If the hacker uses the correct sequence number on the inserted data, the server sends the client an ACK packet containing a sequence number that it is not expecting. The real client then tries to resynchronize with the server by sending an ACK packet with the sequence number that it is expecting. This ACK packet contains a sequence number that the server is not expecting. The server then sends the last ACK packet that it sent, and so on. The resulting acknowledgements (ACKs) bounce back and forth and a TCP ACK storm ensues after the hacker has hijacked multiple client/server sessions.

Related tasks

Creating an attack policy

To protect your system from various types of attacks, such as smurf attacks or SYN floods, you can create one or more intrusion detection attack policies from the **All Policies** view or the **Attack Policies** view.

Extrusion events

An *extrusion* is an attack, traffic regulation, or scan event that originates from the local host system against a remote system. For example, a trusted insider might use a company machine as the origin of a denial-of-service attack. An extrusion also is called an *outbound intrusion*.

IDS detects the following types of outbound attacks:

- Outbound attacks, such as fraggle, flood, UDP echo requests, or smurf attacks. These attacks might show up as broadcast or multicast attempts to the subnet to which a host is connected. These attacks show up as XATTAC in the intrusion monitor record.
- Outbound raw packets that use a nonstandard protocol. Standard protocols include TCP, UDP, ICMP, ICMPv6, IGMP, and OSPF.
- IPv6 routing headers.
- Outbound scans to nonlistening or closed ports. These attacks show up as XSCAN in the intrusion monitor record.
- Outbound traffic regulation events for UDP. These attacks show up as XTRUDP in the intrusion monitor record.
- Outbound traffic regulation events for TCP. These attacks show up as XTRTCP in the intrusion monitor record.

Scan events

A *scan* is an attack that attempts to connect to unused ports looking for a way to break into the system. A scan also can be a connection request from a spoofed IP address. After the open ports are discovered, the hacker tries to discover the weaknesses and gain access to the system.

IDS detects both inbound and outbound scan events.

A *port scan* is used by administrators to check the security of a network, and by hackers or crackers to find open ports and vulnerabilities in the system.

A scan policy can monitor both slow and fast scans. Fast scans might indicate quick attempts at gathering information or attempts to deny service. Slow scans might indicate that a perpetrator is seeking information about which ports to probe or what operating system is running.

If IDS is active before the system IPL, the service stack detects intrusions and extrusions, even if no IDS policies exist. If an IDS scan policy exists, IDS creates an audit record when it detects a scan event, if the slow or fast scan thresholds are exceeded.

Sometimes a high rate of scans indicates that a user is trying to connect to a service that is down, rather than a genuine attack on the system. For example, if the Telnet or TCP/IP server is down, that might look like a scan and IDS would detect it.

Related tasks

Creating a scan policy

To protect your system from unauthorized scans for open ports, you can create an intrusion detection scan policy from the **All Policies** view or the **Scan Policies** view.

Related reference

Example: Intrusion detection scan policy

This example shows an intrusion detection scan policy that monitors for both slow scans and fast scans on all IP addresses and ports 1-5000.

Example: Variable dynamic throttling for scan events

This is an example of how to set variable dynamic throttling for a scan policy. If your system is being attacked, you can set up throttling to limit or deny intrusions.

Traffic regulation events

Traffic regulation policies monitor the established TCP connections on all or specific IP addresses and ports.

A traffic regulation policy might look for an inordinate number of connections to a certain range of addresses, ports, or applications, or a denial-of-service attack on a system. A traffic regulation policy also can catch User Datagram Protocol (UDP) errors.

Sometimes a high rate of network traffic indicates that many legitimate users or applications are accessing the system at the same time, rather than a hacker trying to tie up the network. If you determine that normal network traffic is generating traffic regulation events, you can adjust the traffic regulation policy accordingly.

UDP is an Internet Protocol that provides unreliable, connectionless datagram service. It enables an application program on one machine or process to send a datagram to an application program on another machine or process. IDS detects the following types of UDP traffic regulation events:

- Socket errors.
- Not connected to the sender.
- Not enough room for the datagram (buffer overflow).

Related tasks

Creating a traffic regulation policy

To monitor for extensive network traffic that could potentially overload your system, you can create one or more intrusion detection traffic regulation policies from the **All Policies** view or the **Traffic Regulation Policies** view. You can monitor for TCP or UDP connections.

Related reference

Example: Traffic regulation policy

This example traffic regulation policy traces suspicious traffic across the network, such as an unusually high rate of TCP connections.

Example: Variable dynamic throttling for traffic regulation events

This is an example of how to set variable dynamic throttling for a traffic regulation policy to limit or deny intrusions.

Variable dynamic throttling

You can specify *variable dynamic throttling* in each intrusion detection (IDS) policy. If an enabled IDS policy has throttling specified, throttling occurs after a suspected intrusion or extrusion has occurred and certain thresholds have been reached. Variable dynamic throttling starts discarding packets when a threshold has been exceeded within a given statistics or scan interval.

Consider using variable dynamic throttling when a specific intrusion detection policy is generating a lot of intrusion events to prevent overloading the system with what might be a denial-of-service (DoS) attack.

If you specify variable dynamic throttling for an IDS policy, throttling applies to all IP addresses in that policy. To avoid blocking valid connections to your system, you should not specify variable dynamic throttling before you see a number of intrusion events from a given IP address. You should only specify variable dynamic throttling on a single suspicious address or a limited range of IP addresses.

After you see a number of intrusion events from a given IP address, you can compose an IDS policy that specifies variable dynamic throttling for that IP address. In the policy, you specify that throttling should be activated when the thresholds that you specified in the policy have been exceeded. This minimizes the possibility that throttling would block valid connections and maximize its effect on potential intruders.

After the **Maximum number of events to log** threshold is reached within the interval, variable dynamic throttling begins. If scans continue to be a problem, throttling continues through the next interval.

IDS throttling is both variable and dynamic. Throttling is dynamic in that it goes into effect as soon as a threshold is exceeded. Throttling is variable in that the rate of dropping packets increases as thresholds continue to be exceeded in successive intervals.

For example, if you throttle at 100%, which is the default value, all of the packets that conform to the policy port and IP address ranges are allowed through until a threshold is exceeded twice in a row. In all cases, when a threshold has been exceeded during a throttled interval, the throttling rate is automatically decremented by 10%. If you throttle at 100%, a second throttled interval allows only 9 out of 10 packets through. If you throttle at 50%, 1 out of every 2 packets within the interval is discarded. If you throttle at 0%, all packets are discarded for the throttled interval.

If you specified throttling in your IDS policy, it starts automatically when a threshold is exceeded and gets decremented by 10% for each successive throttled interval. You can use throttling with both intrusions and extrusions.

Related concepts

[Intrusion detection and prevention](#)

You can use the intrusion detection system to prevent intrusions and extrusions from occurring.

Related reference

[Example: Variable dynamic throttling for scan events](#)

This is an example of how to set variable dynamic throttling for a scan policy. If your system is being attacked, you can set up throttling to limit or deny intrusions.

[Example: Variable dynamic throttling for traffic regulation events](#)

This is an example of how to set variable dynamic throttling for a traffic regulation policy to limit or deny intrusions.

Using the Intrusion Detection System GUI

You can use the Intrusion Detection System GUI to configure and manage intrusion detection policies, and display the intrusion events that have been logged in the audit journal.

Using the IDS GUI in IBM Navigator for i

To use the Intrusion Detection System GUI in IBM Navigator for i, perform these steps:

1. Expand **IBM i Management > Security**.
2. Click **Intrusion Detection** to display the **Intrusion Detection Management** page.

After you have started IDS, you can perform the following tasks:

- Start IDS.
- Stop IDS.
- Intrusion detection system setup.
- Manage intrusion detection policies.
- Display intrusion detection events.

IDS does not need to be started to display existing policies or the intrusion events, but IDS must be started to pick up the new policies and to monitor the system for new intrusions and extrusions.

Setting up intrusion detection system properties

In IDS Properties, you can set up e-mail and message notification for IDS, and determine whether to allow Internet Control Message Protocol (ICMP) redirect messages.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to be able to display or change the IDS properties.

To open the IDS Properties page in IBM Navigator for i, expand **IBM i Management > Security > All Tasks > Intrusion detection**. Click **Intrusion Detection System properties**.

Related information

[Special authority](#)

Setting up e-mail and message notification

IDS is a notification system. You can optionally configure IDS to send real-time intrusion notifications to a message queue and to specific e-mail addresses. That way, you can alert systems administrators about specific types of intrusions and extrusions so that they can take actions to stop further intrusions from occurring. You can enable or disable IDS e-mail and message notification per policy.

Prerequisites:

- You must have *ALLOBJ and *IOSYSCFG authority to be able to display or change the IDS properties.
- To use IDS e-mail notification, System i SMTP must be configured and running. For information on how to use SMTP, see [E-mail](#).

To set up e-mail and message notification for IDS, perform these steps:

1. In IBM Navigator for i, expand **IBM i Management > Security > All Tasks > Intrusion detection**.
2. Click **Intrusion Detection System properties**.
3. In the IDS Properties page, select the **Notification** tab.
4. To send intrusion messages to a message queue, select the **Send message notifications** check box and specify the name of the message queue and library. (If the check box remains cleared, IDS does not send notifications to a message queue.)
5. To send intrusion messages to an e-mail address, select the **E-mail address** check box and enter the e-mail address. You can send intrusion messages to up to three e-mail addresses. (If the check box remains cleared, IDS does not send notifications to an e-mail address.)
6. To allow Internet Control Message Protocol (ICMP) redirect messages, click the **ICMP** tab and select the check box. (If the check box remains cleared, IDS does not notify you of ICMP redirect messages.)

ICMP redirect messages are used to inform a host of a more optimal route to a destination. However, a hacker could send an ICMP redirect message to a host to have future traffic directed to the hacker's system.

Intrusion detection events are sent to the specified message queue and e-mail addresses. The IDS Properties settings apply to all of the intrusion detection policies.

Tip: You can configure each intrusion detection policy to send e-mail and message notifications when an intrusion event is detected. To do this, select the **Notification** tab on the IDS Policy Properties page for the specific policy.

Related concepts

[Real-time intrusion and extrusion detection notification](#)

IDS is a notification system. You can configure IDS to send real-time intrusion notifications as messages to a message queue and as e-mail. That way, you can alert systems administrators about specific types of intrusions and extrusions so that they can take appropriate actions.

Starting the intrusion detection system

You need to start the intrusion detection system (IDS) before you can monitor your system for intrusions and extrusions.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG special authority to be able to start IDS.

You no longer need to set the *ATNEVT option in the QAUDLVL system value to enable auditing for intrusions, because that step is done for you automatically when you start IDS.

Starting IDS in IBM Navigator for i

To start IDS in IBM Navigator for i, perform these steps:

1. Expand **IBM i Management > Security**.
2. Click **Intrusion Detection** to display the **Intrusion Detection Management** page.
3. Click **Start** on the **Intrusion Detection Management** page.

The **Intrusion detection notification status** is refreshed to show that IDS is started, and auditing for intrusions is enabled.

Stopping the intrusion detection system

When you stop the intrusion detection system (IDS), it no longer monitors your system for intrusions and extrusions.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG special authority to be able to stop IDS.

Related information

[Remove TCP/IP Table \(RMVTCPTBL\)](#)

[End TCP/IP Server \(ENDTCPSVR\)](#)

Stopping IDS in IBM Navigator for i

To stop IDS in IBM Navigator for i, perform these steps:

1. Expand **IBM i Management > Security**.
2. Click **Intrusion Detection** to display the **Intrusion Detection Management** page.
3. Click **Stop** on the **Intrusion Detection Management** page.

The **Intrusion detection notification status** is refreshed to show that IDS is stopped.

Important: If the IDS GUI is unavailable (perhaps because a user has accidentally shut down the Web server port), use the Remove TCP/IP Table (**RMVTCPTBL *IDS**) CL command to stop IDS manually.

Creating intrusion detection policies

You can create a set of default intrusion detection policies that will monitor for all types of intrusions or extrusions for the entire system. You can also create specific attack, scan, and traffic regulation policies.

You can create intrusion detection policies from the Intrusion Detection Policies page. If you are in the All Policies view, you can create attack, scan, or traffic regulation policies. If you are in the Attack Policies, Scan Policies, or Traffic Regulation Policies view, you can create that specific type of policy.

Related tasks

[Changing an intrusion detection policy](#)

You can change all of the properties of a user-created intrusion detection policy. However, you cannot change many of the properties of a default policy.

[Deleting an intrusion detection policy](#)

You can delete an intrusion detection policy that you no longer need to use.

[Enabling an intrusion detection policy](#)

IDS monitors intrusions for only enabled intrusion detection policies.

[Disabling an intrusion detection policy](#)

IDS monitors intrusions for only enabled intrusion detection policies. You can temporarily disable an intrusion detection policy to prevent IDS from using it to monitor intrusions.

Creating a set of default intrusion detection policies

Create a set of default intrusion detection policies that you can use to monitor for all intrusions and extrusions across all IP addresses and ports on your system.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to work with intrusion detection policies.

The default intrusion detection policies include attack, scan, and traffic regulation policies. To create a set of default intrusion detection policies, perform these steps:

1. In IBM Navigator for i, expand **IBM i Management > Security > All Tasks > Intrusion Detection**.
2. Click **Manage IDS policies**.
3. In the **Intrusion Detection Policies** page, select **New** from the **Actions** menu. The **New intrusion detection policy** wizard is displayed.
4. In the **Select Policy to Create** page, select **Create a set of default intrusion detection policies**. (This function is disabled if the default policies already exist.)
5. Follow the instructions in the wizard to create the policies.
6. Click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Now your system is ready to catch suspicious events coming in through the TCP/IP network.

Many of the properties settings in the default IDS policies are read-only, while all of the properties settings in the user-created IDS policies are editable. The default IDS policies provide intrusion detection coverage for the entire system. If you want more specific policies that cover a specific range of IP addresses or ports, for example, you can create a policy based on a default policy and change those settings. Then you can configure the new policy to take precedence over the default policy. The user-created IDS policy monitors for a subset of intrusions and the system-supplied IDS policy monitors for the rest of the intrusions.

Creating an attack policy

To protect your system from various types of attacks, such as smurf attacks or SYN floods, you can create one or more intrusion detection attack policies from the **All Policies** view or the **Attack Policies** view.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to work with intrusion detection policies.

Related tasks

[Creating an intrusion detection policy based on another policy](#)

If you want to create an IDS policy with many of the same characteristics (such as IP addresses, ports, and notification method), you can create an IDS policy that is based on another policy.

Related reference

Attack events

Attack policies monitor for various types of attacks against the system. Your system can be attacked or used as the source of an attack. When IDS detects an attack, it writes an intrusion event to the audit record.

Creating attack policies in IBM Navigator for i

To create one or more attack policies in IBM Navigator for i, perform these steps:

1. Expand **IBM i Management > Security > All Tasks > Intrusion Detection**.
2. Click **Manage IDS policies**.
3. In the **Intrusion Detection Policies** page, display either the **All Policies** or **Attack Policies** view.
4. Select **New** from the **Actions** menu.
5. Decide whether to create policies for all attack types or one attack type. Then follow the instructions in the wizard to create the attack policies.
6. After you finish creating the policies, click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Creating a scan policy

To protect your system from unauthorized scans for open ports, you can create an intrusion detection scan policy from the **All Policies** view or the **Scan Policies** view.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to work with intrusion detection policies.

Related tasks

Creating an intrusion detection policy based on another policy

If you want to create an IDS policy with many of the same characteristics (such as IP addresses, ports, and notification method), you can create an IDS policy that is based on another policy.

Related reference

Scan events

A *scan* is an attack that attempts to connect to unused ports looking for a way to break into the system. A scan also can be a connection request from a spoofed IP address. After the open ports are discovered, the hacker tries to discover the weaknesses and gain access to the system.

Creating a scan policy in IBM Navigator for i

To create a scan policy in IBM Navigator for i, perform these steps:

1. Expand **IBM i Management > Security > All Tasks > Intrusion Detection**.
2. Click **Manage IDS policies**.
3. In the **Intrusion Detection Policies** page, display either the **All Policies** or **Scan Policies** view.
4. Select **New** from the **Actions** menu.
5. Follow the instructions in the wizard to create the scan policy.
6. After you finish creating the policy, click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Creating a traffic regulation policy

To monitor for extensive network traffic that could potentially overload your system, you can create one or more intrusion detection traffic regulation policies from the **All Policies** view or the **Traffic Regulation Policies** view. You can monitor for TCP or UDP connections.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to work with intrusion detection policies.

Related concepts

Traffic regulation events

Traffic regulation policies monitor the established TCP connections on all or specific IP addresses and ports.

Related tasks

[Creating an intrusion detection policy based on another policy](#)

If you want to create an IDS policy with many of the same characteristics (such as IP addresses, ports, and notification method), you can create an IDS policy that is based on another policy.

Creating traffic regulation policies in IBM Navigator for i

To create traffic regulation policies in IBM Navigator for i, perform these steps:

1. Expand **IBM i Management > Security > All Tasks > Intrusion Detection**.
2. Click **Manage IDS policies**.
3. In the **Intrusion Detection Policies** page, display either the **All Policies** or **Traffic Regulation Policies** view.
4. Follow the instructions in the wizard to create the traffic regulation policies.
5. After you finish creating the policies, click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Creating an intrusion detection policy based on another policy

If you want to create an IDS policy with many of the same characteristics (such as IP addresses, ports, and notification method), you can create an IDS policy that is based on another policy.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to work with intrusion detection policies.

Related tasks

[Creating an attack policy](#)

To protect your system from various types of attacks, such as smurf attacks or SYN floods, you can create one or more intrusion detection attack policies from the **All Policies** view or the **Attack Policies** view.

[Creating a scan policy](#)

To protect your system from unauthorized scans for open ports, you can create an intrusion detection scan policy from the **All Policies** view or the **Scan Policies** view.

[Creating a traffic regulation policy](#)

To monitor for extensive network traffic that could potentially overload your system, you can create one or more intrusion detection traffic regulation policies from the **All Policies** view or the **Traffic Regulation Policies** view. You can monitor for TCP or UDP connections.

Creating a policy based on another in IBM Navigator for i

To create an intrusion detection policy based on another policy in IBM Navigator for i, perform these steps:

1. Expand **IBM i Management > Security > All Tasks > Intrusion Detection**.
2. Click **Manage IDS policies**.
3. Select an intrusion detection policy and select **New Based On** from the **Actions** menu.
4. On the **General** tab, enter the new policy name and change any of the other settings on the **Properties** tabs.
5. Click **OK** on the **Properties** page to create the intrusion detection policy. The new intrusion detection policy appears in the list of policies.
6. Click **OK** to apply the changes.

Managing intrusion detection policies

You can create, enable, disable, delete, or change a policy; create a policy based on another policy; or change the priority of an intrusion detection policy.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to work with intrusion detection policies.

From the **Intrusion Detection Policies** page, you can perform any of the following actions:

- To see all types of intrusion detection policies, select the **All Policies** view.
- To see one type of intrusion detection policy, select the **Attack Policies** view, **Scan Policies** view, or **Traffic Regulation Policies** view.
- In the Attack, Scan, or Traffic Regulation Policy views, use the **Move Up** and **Move Down** actions to change the priority of an intrusion detection policy. The policies are listed in order of priority with the highest priority policy at the top of the list.
- To create an intrusion detection policy, select **New** from the **Actions** menu.
- To create an intrusion detection policy based on another policy, select the policy and then select **New Based On** from the **Actions** menu.
- To disable an intrusion detection policy, select the policy and then select **Disable** from the **Actions** menu.
- To enable an intrusion detection policy, select the policy and then select **Enable** from the **Actions** menu.
- To delete an intrusion detection policy, select the policy and then select **Delete** from the **Actions** menu.
- To display the properties of an intrusion detection policy, select the policy and then select **Properties** from the **Actions** menu.
- In the **All Policies** view, you also can perform additional actions to further tailor the list of policies viewed, such as sorting and filtering the policies in the table.

Changing an intrusion detection policy

You can change all of the properties of a user-created intrusion detection policy. However, you cannot change many of the properties of a default policy.

Prerequisite: You must have *ALLOBJ and *IOSYSCFG authority to be able to change the properties for an intrusion detection policy.

To change an intrusion detection policy, perform these steps:

1. In IBM Navigator for i, expand **IBM i Management > Security > All Tasks > Intrusion Detection**.
2. Click **Manage IDS policies**.
3. In the **Intrusion Detection Policies** page, select a policy from the list, and select **Properties** from the **Actions** menu.
4. Make any of the following changes to the intrusion detection policy:
 - Use the **General** tab to change the description of the policy.
 - Use the **Local IP Addresses** tab to select which local IP addresses to monitor. You can monitor either IPv4 or IPv6 addresses.
 - Use the **Local Ports** tab to select which local ports to monitor.
 - Use the **Remote IP Addresses** tab to select which remote IP addresses to monitor. You can monitor either IPv4 or IPv6 addresses.
 - Use the **Remote Ports** tab to select which remote ports to monitor.
 - Use the **Notification** tab to change how this policy handles notifications, and whether to send an e-mail to the addresses that are defined in IDS Properties.

- Use the **Advanced** tab to control packet throttling. This setting is useful if you are getting too many notifications for a specific intrusion event.
- For a scan policy, use the **Scan Thresholds** tab to change the slow and fast-scan thresholds.
- For a traffic regulation policy, use the **TCP Thresholds** tab to specify when to send an intrusion notification based on the defined connection thresholds.

Related tasks

[Creating intrusion detection policies](#)

You can create a set of default intrusion detection policies that will monitor for all types of intrusions or extrusions for the entire system. You can also create specific attack, scan, and traffic regulation policies.

Changing the priority of an intrusion detection policy

In the Attack, Scan, or Traffic Regulation Policies views, you can change the priority of an intrusion detection policy. However, you cannot change the priority of an intrusion detection policy in the **All Policies** view.

The order that the policies are listed determines the priority of the policies. The policies are processed in this order when intrusion events occur for the IP addresses and ports defined in the policies.

To change the priority of an intrusion detection policy, perform these steps:

1. In the **Intrusion detection policies** table, select the **Attack Policies**, **Scan Policies**, or **Traffic Regulation Policies** view.
2. Select one or more intrusion detection policies:
 - Click **Move Up** from the **Actions** menu to increase the priority.
 - Click **Move Down** from the **Actions** menu to decrease the priority.

Deleting an intrusion detection policy

You can delete an intrusion detection policy that you no longer need to use.

To delete intrusion detection policies, perform these steps:

1. From the **Intrusion Detection Policies** page, select one or more policies to delete, and select **Delete** from the **Actions** menu.

You can delete policies from the All Policies, Attack Policies, Scan Policies, and Traffic Regulation Policies views.

2. Verify the policies to be deleted and click **OK**.
3. Click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Related tasks

[Creating intrusion detection policies](#)

You can create a set of default intrusion detection policies that will monitor for all types of intrusions or extrusions for the entire system. You can also create specific attack, scan, and traffic regulation policies.

Enabling an intrusion detection policy

IDS monitors intrusions for only enabled intrusion detection policies.

Related tasks

[Creating intrusion detection policies](#)

You can create a set of default intrusion detection policies that will monitor for all types of intrusions or extrusions for the entire system. You can also create specific attack, scan, and traffic regulation policies.

Enabling a policy from the Intrusion Detection Policies page

You can enable intrusion detection policies in two ways. To enable an intrusion detection policy from the **Intrusion Detection Policies** page, perform these steps:

1. In the **Intrusion Detection Policies** page, select one or more policies, and select **Enable** from the **Actions** menu. The selected policies are enabled.
2. Click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Enabling a policy from the IDS Policy Properties page

To enable an intrusion detection policy from the **IDS Policy Properties** page, perform these steps:

1. In the **Intrusion Detection Policies** page, select a policy, and select **Properties** from the **Actions** menu.
2. In the **General** tab, check **Policy enabled** and click **OK**. The selected policy is enabled.
3. Click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Disabling an intrusion detection policy

IDS monitors intrusions for only enabled intrusion detection policies. You can temporarily disable an intrusion detection policy to prevent IDS from using it to monitor intrusions.

Related tasks

[Creating intrusion detection policies](#)

You can create a set of default intrusion detection policies that will monitor for all types of intrusions or extrusions for the entire system. You can also create specific attack, scan, and traffic regulation policies.

Disabling a policy from the Intrusion Detection Policies page

You can disable intrusion detection policies in two ways. To disable an intrusion detection policy from the **Intrusion Detection Policies** page, perform these steps:

1. In the **Intrusion Detection Policies** page, select one or more policies.
2. Select **Disable** from the **Actions** menu. The selected policies are disabled.
3. Click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Disabling a policy from the IDS Policy Properties page

To disable an intrusion detection policy from the **IDS Policy Properties** page, perform these steps:

1. In the **Intrusion Detection Policies** page, select a policy, and select **Properties** from the **Actions** menu.
2. On the **General** tab, clear the **Policy enabled** check box and click **OK**. The selected policy is disabled.
3. Click **OK** on the **Intrusion Detection Policies** page to apply the changes.

Backing up the intrusion detection policy file

Back up your intrusion detection policies so that you can restore them if the system has to be scratch installed or if you want to move those policy definitions to another system.

Your intrusion detection policies can be stored locally or exported to a directory server. Back up the `idspolicy.conf` file in the `/QIBM/UserData/OS400/Q0S/ETC` directory.

To ensure that you can easily replace lost IDS policies, follow these steps:

1. Ensure that you have a backup and recovery strategy in place.
2. Decide whether to back up the IDS policies as part of a full-system backup or with other integrated file system files.

Consider maintaining two sets of IDS policies, one set for normal working hours and another set for night hours. For example, the traffic regulation policy for normal working hours would allow a large number of connections, but the policy for night hours would allow just a few connections. Store one set of policies in the `ETC` directory and the other set in some other directory. Then you can write a CL program that swaps the set of policies at the end of each day, and restarts IDS so that those policies take effect.

Related information

[Backing up the integrated file system](#)

Writing intrusion detection programs

You can write a program to analyze the auditing data and statistics to see if a pattern of intrusion or extrusion events emerges, or use journaling APIs to create an audit trail.

For example, the statistics might reveal that suspicious events are occurring during off-hours. The statistics might show that there were attempted attacks on the system. The statistics also might show that the network was misconfigured or not working correctly.

An intrusion detection program might take suspicious events into account as well as network problems that occur for other reasons, such as hardware or configuration problems. For example, Internet Control Message Protocol (ICMP) redirect messages might indicate that a router is not fully configured yet. Sometimes routers are slow to figure out which router in a network is the best route to a destination.

You also can use the QTOQIDSC API to start, stop, or recycle IDS, or to retrieve the IDS status.

Related information

[Control Intrusion Detection and Prevention \(QTOQIDSC, QtoqIDSCControl\) API](#)
[Journal and Commit APIs](#)

Displaying intrusion detection events

Use the Intrusion Detection System GUI to display a list of potential intrusion events as well as detailed information about each event.

To display intrusion detection events, perform these steps:

1. In IBM Navigator for i, expand **IBM i Management > Security**.
2. Click **Intrusion Detection** to display the **Intrusion Detection Management** page.
3. Click **Display intrusion detection events** to display the **Intrusion Detection Events** page.
4. By default, the **Intrusion Detection Events** page lists events that have occurred in the previous 24 hours. Perform any of the following tasks:
 - To refresh the intrusion detection events immediately, select **Refresh** from the **Actions** menu.
 - To display event details, select the event and select **Details** from the **Actions** menu. You also can find these event details in the intrusion monitor audit record.
 - To filter intrusion events, select **Include** from the **Actions** menu. For example, you can display all of the IDS events that have occurred on the system, or include only the events that have occurred in the past five hours.

Tips:

- If you get an intrusion detection event (or IM audit record) of type unknown with an IP address of 0.0.0.0 and any port for the port number, you can ignore it. This type of audit record occurs on system IPL when you specify IDS active.
- If you cannot retrieve the intrusion events using the IDS GUI, use the following CL command to display the intrusion monitor (IM) audit records on the system:

```
DSPJRN JRN(QSYS/QAUDJRN) RCVRNG(*CURCHAIN) ENTTYP(IM)
```

You also can copy the IM records to a file so that you can display all the IM records with their fields. This allows you to see if the intrusions are related by IP address, type, time of arrival, and so on. Use the following CL commands:

```
CPYAUDJRNE IM  
RUNQRY *NONE QAUDITIM
```

Filtering intrusion detection events

You can specify the inclusion criteria for intrusion detection events in the IDS GUI. For example, you could include all events, all events beginning from a specific date and time, or all events within a date and time interval.

By default, the Intrusion Detection Events page displays the events logged within the last 24 hours. To filter the intrusion detection events, perform these steps:

1. Display the **Intrusion Detection Events** page.
2. Select **Include** from the **Actions** menu, which will display the **IDS Events - Include** page.
3. Specify one of the following inclusion criteria for the events and click **OK**.
 - All intrusion detection events
 - All intrusion detection events beginning from a specific date and time
 - All intrusion detection events within a specific date and time interval

The **Intrusion Detection Events** page is immediately refreshed to show the events that meet the inclusion criteria. The **Include** field next to the **Edit** button displays the inclusion criteria that is being used.

Tip: You also can change the inclusion criteria for the intrusion events by clicking **Edit**.

Intrusion monitor audit record entries

The Intrusion Detection System (IDS) GUI displays the intrusion detection events that are generated from the intrusion monitor (IM) audit records in an easy-to-read format. However, you might want to examine the IM audit records while reviewing other audit records.

The following example shows an IM audit record entry with information about an intrusion event for an TCP ACK storm attack. Because the information in lines 151 through 201 are in hexadecimal, they might display as odd characters. Press F11 to display this information in readable form.

```
Display Journal Entry
Object . . . . . : Library . . . . . :
Member . . . . . :
Incomplete data . . . : No      Minimized entry data: *NONE
Sequence . . . . . : 1201
Code . . . . . : T - Audit trail entry
Type . . . . . : IM - Intrusion detection monitor

Entry specific data
00001 'P2007-06-08-13.38.06.8811471114 004499.5.6.170 '
00051 '                                020019.10.108.13'
00101 '6                                ATTACK0023ACKST'
```

The following example shows an IM audit record entry with information about an intrusion event for a SCAN attack.

```
Display Journal Entry
Column
Object . . . . . : Library . . . . . :
Member . . . . . :
Incomplete data . . . : No      Minimized entry data: *NONE
Sequence . . . . . : 209
Code . . . . . : T - Audit trail entry
Type . . . . . : IM - Intrusion detection monitor

Entry specific data
Column
00001 'P2007-05-25-16.03.28.8169131107 003899.5.138.154 '
00051 '                                250799.5.138.154'
00101 '                                SCANE 0024'
00151 ' 6 22 100000P ,
```

The following table shows the layout of the IM audit record. Use the information in this table to analyze and interpret the IM audit records.

Table 1. Layout of the IM audit record

Field Type	Format	Description
Entry type	Char(1)	The type of entry. P Potential intrusion event detected.
Time of Event	TIMESTAMP	Timestamp of when the event was detected in SAA timestamp format.
Detection Point Identifier	Char(4)	This is a unique identifier for the processing location that detected the intrusion event. This field is intended for use by service personnel.
Local Address Family	Char(1)	Local IP address family associated with the detected event.
Local Port Number	Zoned(5,0)	Local port number associated with the detected event.
Local IP Address	Char(46)	Local IP address associated with the detected event.
Remote Address Family	Char(1)	Remote address family associated with the detected event.
Remote Port Number	Zoned(5,0)	Remote port number associated with the detected event.
Remote IP Address	Char(46)	Remote IP address associated with the detected event.
Probe Type Identifier	Char(6)	The type of condition that is used to detect the potential intrusion or extrusion. Possible values include: ATTACK Attack action detected event TR-TCP Traffic Regulation action detected event over TCP TR-UDP Traffic Regulation action detected event over UDP SCANE Scan event action detected event SCANG Scan global action detected event XATTAC Possible extrusion attack XTRTCP Outbound TR detected event (TCP) XTRUDP Outbound TR detected event (UDP) XSCAN Outbound scan event detected
Event Correlator	Char(4)	Unique identifier for this specific intrusion event. This identifier can be used to correlate this audit record with other intrusion detection information.

Table 1. Layout of the IM audit record (continued)

Field Type	Format	Description
Event Type	Char(8)	<p>Identifies the type of potential intrusion that was detected. The possible values are:</p> <p>ACKSTORM TCP ACK storm</p> <p>ADRPOISN Address poisoning</p> <p>FLOOD Flood event</p> <p>FRAGGLE Fraggle attack</p> <p>ICMPRED ICMP (Internet Control Message Protocol) redirect</p> <p>IPFRAG IP fragment</p> <p>MALFPKT Malformed packet</p> <p>OUTRAW Outbound Raw</p> <p>PERPECH Perpetual echo</p> <p>PNGDEATH Ping of death</p> <p>RESTOPT Restricted IP options</p> <p>RESTPROT Restricted IP protocol</p> <p>SMURF Smurf attack</p>
Protocol	Char(3)	Protocol number.
Condition	Char(4)	Condition number from IDS policy file.
Throttling	Char(1)	<ul style="list-style-type: none"> • 0 = not active • 1 = active
Discarded Packets	Zoned(5,0)	Number of discarded packets when throttled.
Target TCP/IP Stack	Char(1)	<p>P Production Stack</p> <p>S Service Stack</p>
Reserved	Char(6)	Reserved for future use.
Suspected Packet	Char(1002)	This is a variable length field which may contain up to the first 1000 bytes of the IP packet associated with the detected event. This field contains binary data and should be treated as if it has a CCSID of 65535. The first 2 bytes contain the length of the suspected packet information.

Examples: Intrusion detection

Use the examples in this section to create various types of intrusion detection policies.

Example: Traffic regulation policy

This example traffic regulation policy traces suspicious traffic across the network, such as an unusually high rate of TCP connections.

Traffic regulation events correlate to completed handshakes for connections. The intrusion detection system tracks the TCP traffic over the IP addresses and ports that are specified in the IDS policy. When user-specified thresholds are met, IDS generates an intrusion event.

This intrusion detection policy specifies a TCP connection limit of 1000, a TCP connection percentage of 100%, a statistics interval of 60 minutes, and a maximum number of 5 event messages. When IDS detects the 1001st TCP connection to port 8000 at local addresses 9.10.11.000 through 9.10.11.255, it sends the intrusion notification to the specified e-mail addresses and logs the notification to the audit journal. Use the **Intrusion Detection Events** page to display the logged events. IDS can send a maximum of five intrusion notifications within each 60-minute interval.

The number of audit records that the system generates depends on the value of the Maximum event messages in the intrusion detection policy file.

Setting	Value
Policy name	TR_policy
Policy type	Traffic regulation (TCP)
Threshold for the total number of TCP connections	1000
TCP connection percentage	100
Local IP addresses	9.10.11.000-9.10.11.255
Local ports	8000
Remote IP addresses	All IP addresses
Remote ports	All ports
Statistics interval	60 minutes
Maximum event messages	5
Send e-mail notification ¹	Yes

¹ IDS sends e-mail notification only if you have enabled this support in the IDS Properties page, which is where the e-mail addresses are specified.

Related concepts

[Traffic regulation events](#)

Traffic regulation policies monitor the established TCP connections on all or specific IP addresses and ports.

Example: Restricted IP options policy

This example is of an IDS attack policy that targets restricted IP options for a single local IPv6 address, a range of remote IPv6 addresses, and all ports.

There are 256 possible IP options, with only a small number currently in common use. Checking for restricted IP options is performed on all inbound and outbound packets, even those forwarded to another

system. You can use the IDS policy to provide notification of a packet with a restricted IP option, as well as to discard the packet.

A hacker might try to use restricted IP options, such as Loose Source and Record Route (LSRR), to get through firewalls. LSRR is used to map a network's topology and discover private IP addresses.

<i>Table 3. Restricted IP options example</i>	
Setting	Value
Policy name	Restricted_IP_option_policy
Policy type	Attack
Attack type	Restricted IP options
Local IP addresses	2001:0db8:3c4d:0015:0000:0000:abcd:ef12
Local ports	All ports
Remote IP addresses	2002:9436:7a00:0000:0000:0000:0000:0000-2002:9436:7aff:ffff:ffff:ffff:ffff:ffff
Remote ports	All ports
Statistics interval	5 minutes
Maximum event messages	5
Send e-mail notification	Yes

Example: Perpetual echo policy

This example is of an IDS attack-type policy that targets perpetual echoes on local port 7 and remote port 7.

UDP port 7 is the echo port. In an attack, if the header specifies the source and target ports as port 7, the UDP datagram echoes back and forth between the local port 7 and the remote UDP port 7.

When a perpetual echo occurs on port 7, IDS sends an intrusion notification to the **Intrusion Detection Events** page and to the audit journal, but it does not send an e-mail notification.

Each event that is detected is logged. Ensure that IDS does not overload the system if it is logging large numbers of events. If IDS is logging too many events, you can reduce the number of events being logged by using any of the following methods:

- Using variable dynamic throttling.
- Changing the IDS policy to monitor fewer IP addresses.
- Limiting the maximum number of messages.

<i>Table 4. Perpetual echo policy example</i>	
Setting	Value
Policy name	Echoes_policy
Policy type	Attack
Attack type	Perpetual echo
Local IP addresses	All IP addresses
Local ports	7
Remote IP addresses	All IP addresses
Remote ports	7

<i>Table 4. Perpetual echo policy example (continued)</i>	
Setting	Value
Send messages for each intrusion	Yes
Send e-mail notification	No

Example: E-mail notification

In this example, IDS detected an intrusion on the local system and sent an e-mail notification to the systems administrator.

The following is an example of an e-mail notification received for a restricted IP options attack:

```
To: Sysadmin
Subject: A possible intrusion, suspicious inbound activity, was detected on sys1234.
```

The following information was gathered about the event:

```
Time of Event: date time
Extrusion Type: ATTACK
Attack Type: RESTOPT
Local IP Address: 224.0.0.1
Local Port: 0
Remote IP Address: 9.5.211.4
Remote Port: 0
Protocol: 2
Throttling Active: *NO
Discarded Packet Count: 0
Condition ID: 11
Stack: P
Event Correlator: 0001
Detection Point ID: 1001
Suspected Packet:
X'<long hexadecimal string>'
```

Recovery . . . : For more information on actions you can take to block and impede future suspicious inbound activity, see the Intrusion detection topic in the Security category in the IBM i Information Center.

Example: Intrusion detection scan policy

This example shows an intrusion detection scan policy that monitors for both slow scans and fast scans on all IP addresses and ports 1-5000.

A high number of *fast scans* might indicate quick attempts at gathering information or attempts to deny service. A high number of *slow scans* might indicate that a perpetrator is seeking information about which ports to probe or what operating system is running. Sometimes a high rate of scans indicates that a user is trying to connect to a system that is down, rather than a genuine attack on the system.

This IDS scan policy targets local and remote ports 1 through 5000 for suspicious events. An intrusion notification is logged if the number of slow scans within a 100-minute interval exceeds 64, or if the number of fast scans within a 1-minute interval exceeds 20. IDS can send up to five intrusion notifications during each scan interval.

<i>Table 5. Scan policy example</i>	
Setting	Value
Policy name	Fast_scan
Policy type	Scan
Slow scan interval	100 minutes
Slow scan threshold	64
Fast scan interval	1 minute
Fast scan threshold	20

<i>Table 5. Scan policy example (continued)</i>	
Setting	Value
Local IP addresses	All IP addresses
Local ports	1-5000
Remote IP addresses	All IP addresses
Remote ports	1-5000
Maximum event messages	5
Send e-mail notification	Yes

Related reference

Scan events

A *scan* is an attack that attempts to connect to unused ports looking for a way to break into the system. A scan also can be a connection request from a spoofed IP address. After the open ports are discovered, the hacker tries to discover the weaknesses and gain access to the system.

Example: Variable dynamic throttling for scan events

This is an example of how to set variable dynamic throttling for a scan policy. If your system is being attacked, you can set up throttling to limit or deny intrusions.

Throttling allows you to discard packets when the intrusion threshold has been exceeded within a scan interval. Throttling automatically starts when the intrusion threshold is exceeded. The throttling rate automatically decrements by 10% for each successive throttled interval. This means that 10% more packets are discarded in each successive throttled interval. You can use throttling with both intrusions and extrusions.

In this example, the IDS scan policy signals a scan event if the following conditions are met:

- A connection attempt is made to nonlistening ports 26 to 136 from remote IP addresses in the range of 9.0.0.0 to 9.255.255.255.
- If fast scans occur at a rate of five for a 1-minute interval, or if slow scans occur at a rate of 10 for an interval of 120 minutes.

Set throttling on the **Advanced** tab in **IDS Policy Properties**. If throttling is active and taking place at the rate of 50 percent, the first packet in the scan interval is discarded, and the second packet is allowed through. Throttling begins once the fast scan or slow scan threshold is exceeded. A threshold violation occurs when the number of scans received during a user-defined fast scan interval exceeds the fast scan threshold, or when the number of slow scans received during a user-defined slow scan interval exceeds the slow scan threshold.

If thresholds are not exceeded during a throttled interval, throttling will be active for only that interval. In this example, if the slow scan threshold is exceeded, throttling will be in effect for at least 120 minutes. If the threshold is exceeded during a throttled interval, the throttle rate is decremented by 10% to a minimum of 0%, at which time, all packets are discarded for that interval. Throttling is deactivated only when thresholds are not exceeded during a time interval.

A throttling value of 100% allows all packets through, while a throttling value of 0% stops all packets from coming through. If you want to totally shut down the source of an attack, you would set throttling to 0%.

<i>Table 6. Variable dynamic throttling for scan events</i>	
Setting	Value
Policy name	Scan_policy2
Policy type	Scan

Table 6. Variable dynamic throttling for scan events (continued)

Setting	Value
Fast scan interval	1 minute
Fast scan threshold	5
Slow scan interval	120 minutes
Slow scan threshold	10
Local IP addresses	All IP addresses
Local ports	All ports
Remote IP addresses	9.0.0.0 to 9.255.255.255
Remote ports	26 to 136
Maximum event messages	5
Throttling	50%

Related concepts

Variable dynamic throttling

You can specify *variable dynamic throttling* in each intrusion detection (IDS) policy. If an enabled IDS policy has throttling specified, throttling occurs after a suspected intrusion or extrusion has occurred and certain thresholds have been reached. Variable dynamic throttling starts discarding packets when a threshold has been exceeded within a given statistics or scan interval.

Related reference

Scan events

A *scan* is an attack that attempts to connect to unused ports looking for a way to break into the system. A scan also can be a connection request from a spoofed IP address. After the open ports are discovered, the hacker tries to discover the weaknesses and gain access to the system.

Example: Variable dynamic throttling for traffic regulation events

This is an example of how to set variable dynamic throttling for a traffic regulation policy to limit or deny intrusions.

Suppose you have created the following traffic regulation policy with throttling set to 50 percent. An intrusion event is generated when the number of established TCP connections exceeds 1000 connections, or exceeds 10% of the total number of connections to the system during a 10-minute interval. The maximum number of event messages during each statistics interval is 1. At this point, throttling begins. Input from all IP addresses coming in on port 80 is cut back to just 50% for 10 minutes (the statistics interval). During this time period, IDS keeps statistics for the given protocol, range of IP addresses, and ports. After the statistics interval ends, IDS evaluates whether to continue to throttle during the next 10-minute interval, based on the statistics gathered during the throttled interval.

Table 7. Variable dynamic throttling for traffic regulation events

Setting	Value
Policy name	TR_policy2
Policy type	Traffic regulation (TCP)
Threshold for the total number of TCP connections	1000
TCP connection percentage	10
Local IP addresses	All IP addresses
Local ports	80

Table 7. Variable dynamic throttling for traffic regulation events (continued)

Setting	Value
Remote IP addresses	All IP addresses
Remote ports	All ports
Statistics interval	10 minutes
Maximum event messages	1
Throttling	50%

Related concepts

Traffic regulation events

Traffic regulation policies monitor the established TCP connections on all or specific IP addresses and ports.

Variable dynamic throttling

You can specify *variable dynamic throttling* in each intrusion detection (IDS) policy. If an enabled IDS policy has throttling specified, throttling occurs after a suspected intrusion or extrusion has occurred and certain thresholds have been reached. Variable dynamic throttling starts discarding packets when a threshold has been exceeded within a given statistics or scan interval.

Related information for Intrusion detection

Product manuals, IBM Redbooks® publications, Web sites, and other information center topic collections contain information that relates to the Intrusion detection topic collection. You can view or print any of the PDF files.

Other information

- The [Plan and set up system security](#) topic, which discusses techniques for detecting other types of intrusions
- The [Security Reference](#) topic, which provides reference information about intrusion monitor journal entries
- The [Quality of service](#) topic, which discusses how to use the QoS commands to activate an intrusion detection policy in V5R4

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department YBWA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Programming interface information

This Planning and setting up system security publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Other product and service names might be trademarks of IBM or other companies.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

