IBM Systems - iSeries

Security
Intrusion detection

Version 5 Release 4

IBM Systems - iSeries

Security
Intrusion detection

Version 5 Release 4

Note

Before using this information and the product it supports, read the information in “Notices,” on page 15.

First Edition (February 2006)

This edition applies to version 5, release 4, modification 0 of i5/OS (product number 5722–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.

Contents

Intrusion detection

Intrusion detection involves gathering information about unauthorized access attempts and attacks coming in over the TCP/IP network. Security administrators can analyze the auditing records that intrusion detection provides to secure the iSeries™ network from these types of attacks.

Intrusion encompasses many undesirable activities such as information theft and denial of service attacks. The objective of an intrusion may be to acquire information that a person is not authorized to have (information theft). The objective may be to cause a business harm by rendering a network, system, or application unusable (denial of service), or it may 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, detecting access attempts, and attack behaviors.

You can create an intrusion detection policy that audits suspicious intrusion events that come in through the TCP/IP network. Examples of problems that the intrusion detection function looks for includes:

- Denial of service attacks
- Port scans
- Malformed packets
- Internet protocol (IP) fragments
- Restricted IP options and protocols
- Internet Control Message Protocol (ICMP) redirect messages
- Perpetual echo attacks on User Datagram Protocol (UDP) port 7 (the echo port)

You also can write an application to analyze the auditing data and report to the security administrator if TCP/IP intrusions are likely to be underway.

Important: The term intrusion detection is used two ways in the iSeries 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 an invalid user ID, 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 new intrusion detection function that uses policies to monitor suspicious traffic on the system.

What's new for V5R4

The entire intrusion detection topic is new in V5R4.

Intrusion detection

Using an intrusion detection policy, you can detect intrusions to the TCP/IP network and create auditing records.

You can perform the following intrusion detection functions to keep your system secure:

- Create an intrusion detection policy in the idspolicy.conf file to monitor specific types of unauthorized access attempts and attacks to the TCP/IP network.

- Audit suspicious intrusion activities.
- Analyze the auditing data and make recommendations to the security administrator if TCP/IP intrusions are likely to be underway.

How to see what's new or changed

Revision bars are not used in this topic because it is new.

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

Printable PDF

Use this to view and print a PDF of this information.

To view or download the PDF version of this document, select Intrusion detection (about 257 KB).

You can view or download these related topics:

- Plan and set up system security (3907 KB) which discusses techniques for detecting other types of intrusions.
- Quality of Service (QoS) (947 KB) which discusses how to use the QoS commands to activate an intrusion detection policy.

Saving PDF files

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

1. Right-click the PDF in your browser (right-click the link above).
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 (www.adobe.com/products/acrobat/readstep.html)  .

Concepts

This topic describes how the intrusion detection system works.

Intrusion detection uses the idspolicy.conf file that contains a set of policies for intrusion events. Each policy has an associated condition and action, but there might be more than one condition associated with the same action. The TCP/IP stack reports the most common potential intrusion events and audits them, so that you can write an application to analyze the data and report to the security administrator if intrusions are likely to be underway. The following diagram shows how the intrusion detection function works.

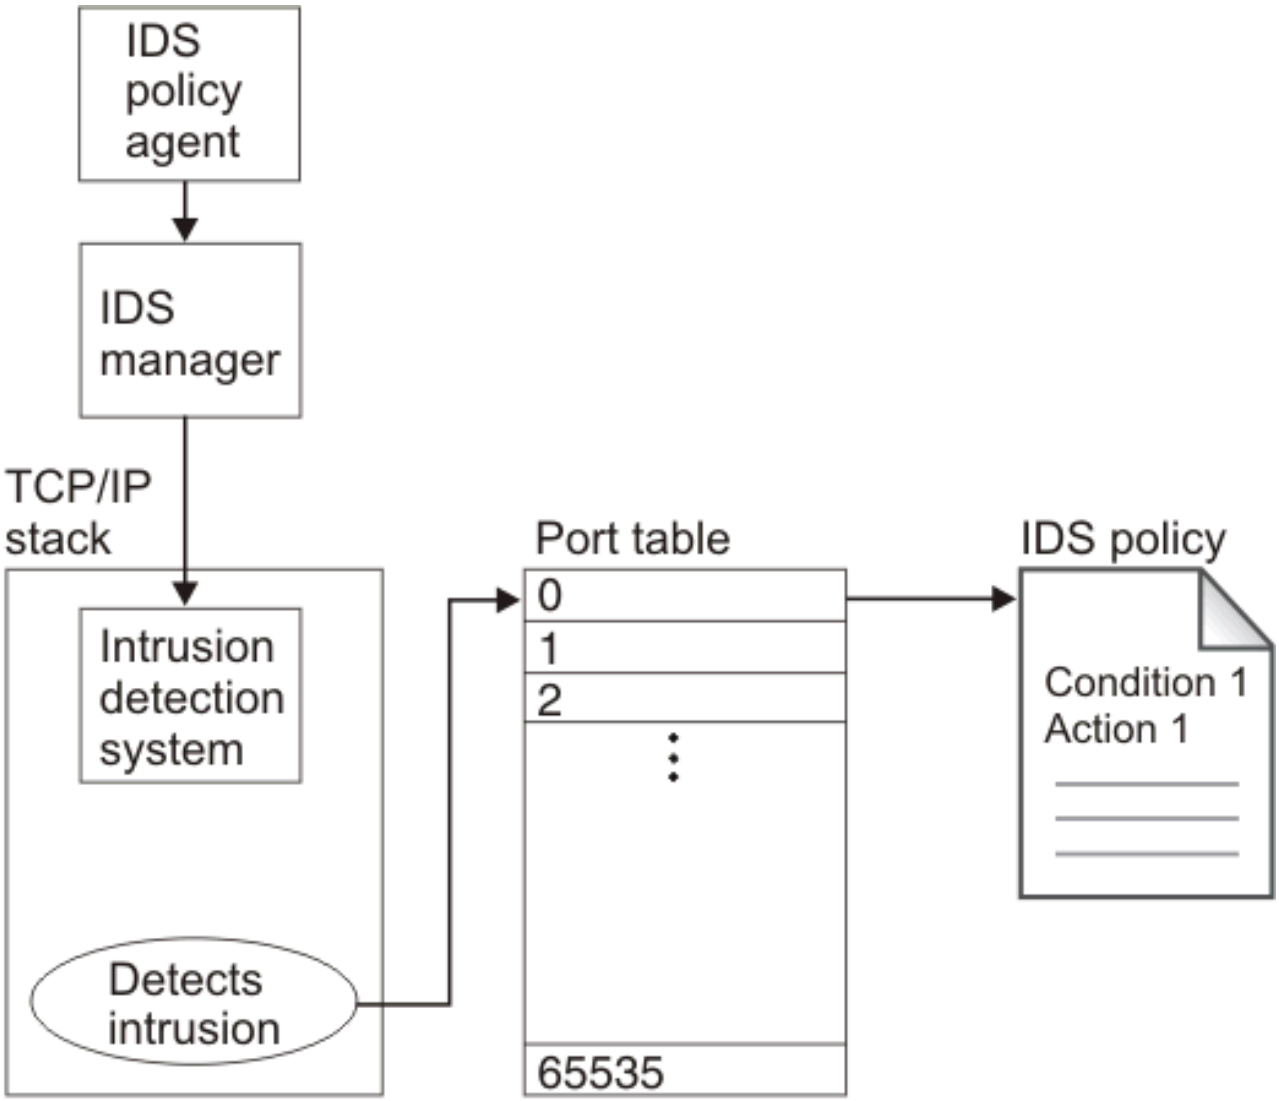

Intrusion detection system

RZAUB500-0

1. You edit the idspolicy.conf file to detect specific types of intrusions, and then you start the QoS server.
2. The QoS policy agent reads the intrusion detection policy in the idspolicy.conf file.
3. The QoS policy agent sends a message with machine instructions to the QoS manager.
4. The QoS manager interprets the machine instructions and sends them to the intrusion detection system inside the TCP/IP stack. The TCP/IP stack manages outbound traffic and inbound traffic in the network, and routes requests to other computers in the network.
5. The intrusion detection system creates the policies in the port table. The port table entries represent ports 0 through port 65 535. For example, port 0, which contains conditions that apply to all ports, points to intrusion condition 1 which points to action 1. Similarly, port 1 points to condition 2 which points to action 2. Port 1 also points to condition 3 which points to action 1, and so on.
6. When the TCP/IP stack detects an intrusion, it looks for matching conditions in the port table and executes the specific action, for example, creating an IM auditing record or system statistics.
7. The system creates an IM audit record which describes the type of intrusion event.

8. The system administrator analyzes the IM audit record to determine which security actions to take, such as closing off the port from where the intrusion originated.

Terminology

This topic defines intrusion detection terms.

denial-of-service (DOS) attack
In computer security, 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. Network service is interrupted for some period.

Internet Control Message Protocol (ICMP)
An Internet protocol that is used by a gateway to communicate with a source host, for example, to report an error in a datagram.

ICMP scan
An attack that tries to use ICMP to overload the system. This is typically a denial-of-service attack.

intrusion detection
A broad term encompassing the detection of many undesirable activities. 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 now 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, detecting information gathering, access attempts, and attack behaviors are vital parts of intrusion detection.

port scan
An attack that attempts to connect to unused ports looking for a way to break into the system.

Quality of Service (QoS)
Any operation that allows traffic priorities to be designated. Through QoS, different traffic throughout a network can be classified and administered.

traffic regulation (TR)
Used for intrusion detection policies that specify the data/connection rate thresholds.

User Datagram Protocol (UDP)
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.

Set up a new intrusion detection policy

Learn how to set up an intrusion detection policy for the first time.

An intrusion detection (IDS) policy consists of two parts:

- An IDS condition that identifies the conditions (such as the port, protocol, or IP address) that applies to the IDS policy.
- An IDS action that identifies the actions to take when a condition is met. Multiple conditions can point to the same action.

The IDS policy file, idspolicy.conf, is shipped with the i5/OS™ system and stored in the /QIBM/ProdData/OS400/QOS/idspolicy.conf directory. A sample IDS policy, which is commented out, is included in this shipped file.

Ensure that you have authority to the /QIBM/UserData/OS400/QOS/ETC/ directory and the idspolicy.conf file. Follow these steps to set up your intrusion detection policy for the first time:

1. Issue the following command to set IP QoS enablement to Yes: `CHGTCPA IPQOSENB(*YES)`
2. Issue the WRKSYSVAL command to set the auditing system values. Then you will see a list of system values.
   a. Type `2` to display the auditing options for the QAUDLVL system value.
   b. Add `*ATNEVT` to the list of auditing options.

      If there is no room in QAUDLVL to set *ATNEVT, be sure that *AUDLVL2 is set in QAUDLVL, as described below. Press PF3 to exit.
   c. Type `2` to display the auditing options for the QAUDLVL2 system value.
   d. Add `*ATNEVT` to the list of auditing options. Press PF3 to exit.
3. To configure the IDS policy file, copy the file from /QIBM/ProdData/OS400/QOS/idspolicy.conf to /QIBM/UserData/OS400/QOS/ETC/.
4. Edit the IDS policy file.
5. Start the QoS server using the following command: `strtcpsvr *qos`

   When you start the QoS server, it looks in the ETC directory for the idspolicy.conf file. If the idspolicy.conf file is not found, it is copied from the /QIBM/ProdData/OS400/QOS/ directory into the /QIBM/UserData/OS400/QOS/ETC/ directory.
6. Issue the Work with Active Jobs (WRKACTJOB) command to verify that the QoS server has started. You will see QTOQSRVR in the list of started servers.

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

Related reference

"Keywords in the IDS policy file" on page 7
Most of the keywords in the IDS policy file are supported in this release, but a few of them are not supported.

Change the intrusion detection policy file

Learn the steps for changing your intrusion detection policy file.

Follow these steps to edit your intrusion detection policy file:

1. Stop the QoS server using the following command: `endtcpsvr *qos`
2. Edit the IDS policy file in the /QIBM/UserData/OS400/QOS/ETC/ directory.
3. Start the QoS server using the following command: `strtcpsvr *qos`
4. Issue the Work with Active Jobs (WRKACTJOB) command to verify that the QoS server has started. You will see QTOQSRVR in the list of started servers.

Example: Traffic regulation policy

This sample 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 generates statistics and when user-specified thresholds are met, the system generates an audit record. Use the ibm-idsMaxEventMessage parameter in the IDS policy file to limit the number of records written to the audit journal.

This policy points to a single IDS traffic regulation (TR) condition and a single IDS action. The IDS condition selects the TCP protocol, local port 8000, and a local host IP address.

The IDS action specifies a TCP connection limit of 1000 for the listening server, a statistics interval of 10 minutes, and 10 percent of the TR connections. This example shows the local host IP addresses as a range of addresses from 9.10.11.000 through 9.10.11.255.

```
 ibm-idsConditionAuxClass     rule1     # IDS condition
 {
 ibm-idsConditionType         TR
 ibm-idsLocalPortRange        8000
 ibm-idsProtocolRange         6
 ibm-idsLocalHostIPAddress    2-9.10.11.000-24
 policyIdsActionName          idsact1
 }

 ibm_idsActionAuxClass        idsact1   # IDS action
 {
 ibm-idsActionType            TR
 ibm-idsStatInterval          10
 ibm-idsTRtcpTotalConnections 1000
 ibm-idsTRtcpPercentage       10
}
```

Example: Restricted IP protocol policy

This example is of an IDS attack-type policy that targets restricted IP protocols in the range of 200 to 205.

```
 ibm-idsConditionAuxClass     idscond4   # IDS condition
{
ibm-idsConditionType         ATTACK
ibm-idsAttackType            RESTRICTED_IP_OPTIONS
ibm-idsProtocolRange         200-205
ibm-policyIdsActionName      idsact2
}

 ibm-idsActionAuxClass        idsact2
 {
 ibm-idsActionType            ATTACK
 ibm-idsMaxEventMessage       5
 }
```

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.

This example uses the same IDS action, `idsact2`, as the "Example: Restricted IP protocol policy" example.

```
  ibm-idsConditionAuxClass    idscond5   # IDS condition
  {
  ibm-idsConditionType        ATTACK
  ibm-idsAttackType           PERPETUAL_ECHO
  ibm-idsLocalPortRange       7
  ibm-idsRemotePortRange      7
  ibm-policyIdsActionName     idsact2
  }
```

Example: Intrusion detection scan policy

This example is of a scan policy that uses stand-alone conditions and actions.

The TCP/IP stack detects port scans on a port-by-port basis. The stack itself cannot detect a global scan. When a port scan is suspected, it generates a SCAN_EVENT that calls the intrusion detection system. The intrusion detection system processes the scan event and calls the SCAN_GLOBAL code to generate statistics and monitor thresholds.

This IDS policy targets TCP ports 1 through 5000 for suspicious events.

```
ibm-idsConditionAuxClass     idscond10  # IDS condition
{
ibm-idsConditionType         SCAN_EVENT
ibm-policyIdsActionName      idsscan1
ibm-idsProtocolRange         6
ibm-idsLocalPortRange        1-5000
}
ibm-idsActionAuxClass        idsscan1   # IDS action
{
ibm-idsActionType            SCAN_GLOBAL
ibm-idsFSInterval            10
ibm-idsFSThreshold           10         # fast scanning threshold
ibm-idsSSInterval            100
ibm-idsSSThreshold           20         # slow scanning threshold
}
```

Keywords in the IDS policy file

Most of the keywords in the IDS policy file are supported in this release, but a few of them are not supported.

Supported keywords

The IDS policy contains the following supported keywords:

- ibm-policyIdsActionName
- ibm-idsICMPRedirect
- ibm-idsConditionAuxClass
- ibm-idsConditionType
- ibm-idsAttackType
- ibm-idsLocalPortRange
- ibm-idsRemotePortRange
- ibm-idsProtocolRange
- ibm-idsIPOptionRange
- ibm-idsLocalHostIPAddress
- ibm-idsRemoteHostIPAddress
- ibm-idsActionAuxClass
- ibm-idsActionType
- ibm-idsStatInterval
- ibm-idsMaxEventMessage
- ibm-idsTRtcpTotalConnections
- ibm-idsTRtcpPercentage
- ibm-idsTRtcpLimitScope
- ibm-idsTRudpQueueSize
- ibm-idsFSInterval
- ibm-idsFSThreshold
- ibm-idsSSInterval
- ibm-idsSSThreshold

The following keywords in the IDS policy file, while allowed, are ignored in this release.

ibm-idsMessageDest
Specifies to which queue the IDS-generated messages should go. (All messages result in audit records and are not sent to queues.)

ibm-idsNotification
Specifies whether the log file or the console gets notified. (All messages go to the audit journal only.)

ibm-idsLoggingLevel
Specifies a limit to the number of messages logged to a log file.

ibm-idsTypeActions
Specifies the type of action to take for a condition. (The only action taken is to create an audit record.)

ibm-idsSensitivity
Specifies the priority of the condition. (All conditions are treated as having equal priority.)

ibm-idsScanExclusion
Specifies an array of IP addresses and ports that should be exempt from statistical bookkeeping if a scan is detected. No IP addresses or ports are exempt from the statistics that are associated with a scan event.

Related tasks

"Set up a new intrusion detection policy" on page 4
Learn how to set up an intrusion detection policy for the first time.

Back up the intrusion detection policy file

You should back up your intrusion detection (IDS) policies to eliminate the need to re-create your policies in the event of a server outage or power loss.

Your IDS policies can be stored locally or exported to a directory server. You must back up the IDS policies in the following directories:

QIBM/UserData/OS400/QOS/ETC

QIBM/ProdData/OS400/QOS/

You must also back up your directory server publishing agent for the QoS server. The publishing agent contains the directory server name, the distinguished name (DN) for the QoS server, port used to access the directory server, and authentication information. In the event of a loss, your backups can save you the time and work it takes to re-create your policies from scratch.

Follow these steps to ensure that you can easily replace lost IDS policies:

1. **Use integrated file systems backup and recovery programs.**
   The *Backup and recovery* book provides instructions on conducting backups from integrated file systems.
2. **Print out the policies.**
   You can store the printouts wherever they are most likely to be secure and re-enter the information as necessary.
3. **Copy the information to a disk.**
   Copying has an advantage over printouts: rather than reentering manually, the information exists electronically. It provides you a straightforward method for transporting information from one online source to another.

Related information

Backup and Recovery PDF

Manage the intrusion detection policy file

You can configure an intrusion detection program to send e-mail to a system administrator to alert them to suspicious events and provide suggestions as to what action to take.

You also can write a program to analyze the statistics for certain patterns. 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 should take suspicious events into account as well as network problems that occur for other reasons such as hardware or configuration problems. For example, 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.

Audit intrusion detection activities

Learn how to audit intrusion detection activities. If the intrusion detection system (IDS) flags a suspicious event, it writes an IM audit record.

The audit record is written to the security audit journal whenever the QAUDCTL system value contains *AUDLVL and either the QAUDLVL or QAUDLVL2 system value contains *ATNEVT.

Note: To set *ATNEVT in QAUDLVL2, you must first set *AUDLVL2 in QAUDLVL.

To view the IM audit records, follow these steps:

1. Issue the following command from the command line to display all of the audit journals: `DSPJRN QAUDJRN`

   If you find an audit record of type IM, that means that IDS has flagged a suspicious event. If no IM audit records display, IDS has not detected any suspicious events. (To display only the IM audit records, issue the `DSPJRN QAUDJRN ENTTYP(IM)` command.)
2. Type `5` to view the contents of the IM audit record.
3. Report suspicious events to your systems administrator to take appropriate action, such as closing down the port or tracking down the spoofed IP address.

Now, you are ready to analyze the IM audit records. The audit record is the only way of alerting a system administrator that a suspicious event has taken place.

Note: Some fields in the IM record are in hexadecimal format. To view those hexadecimal fields, press F11.

Related reference

"Analyze the auditing data" on page 10
Learn how to analyze the auditing data for intrusion detection activities, and obtain reference information about the fields in the IM audit record.

Analyze the auditing data

Learn how to analyze the auditing data for intrusion detection activities, and obtain reference information about the fields in the IM audit record.

The following example shows an IM audit record entry with information about an intrusion event.

```
            Display Journal Entry

Object . . . . . .:          Library  . . . . . .:
Member . . . . . .:
Incomplete data  . .:   No      Minimized entry data: *NONE
Sequence . . . . . .:   5
Code . . . . . . . .:   T  - Audit trail entry
Type . . . . . . . .:   IM - Intrusion detection monitor

           Entry specific data
Column    *...+....1....+....2....+....3....+4....+....5.
00001    'P2005-06-06-15.01.32.6482729999 000009.10.11.0    '
00051    '                                 000009.10.11.255'
00101    '                         ,        ATTACK    RESTP'
00151    'ROT
```

The following table shows the layout of the IM audit record.

Table 1. Layout of the IM audit record

Field Type	Format	Description	Sample Entry
Entry type	Char(1)	Potential intrusion event detected.	P
Time of event	TIMESTAMP	Timestamp of when the event was detected.	2005-06-06-15.01.32.648272
Detection point identifier	Char(4)	Unique identifier for the processing location that detected the intrusion event. This field is for use by service personnel.	9999
Local address family	Char(1)	Local IP address family associated with the detected event.	This field is hidden and appears blank. Press F11 to display the information.
Local port number	Zoned(5,0)	Local port number associated with the detected event. (A value of 00000 represents an intrusion on any port because there is no port 0.)	00000
Local IP address	Char(46)	Local IP address associated with the detected event.	9.10.11.0
Remote address family	Char(1)	Remote address family associated with the detected event.	This field is hidden and appears blank. Press F11 to display the information.
Remote port number	Zoned(5,0)	Remote port number associated with the detected event.	00000
Remote IP address	Char(46)	Remote IP address associated with the detected event.	9.10.11.255

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

Field Type	Format	Description	Sample Entry
Probe type identifier	Char(6)	Identifies the type of probe used to detect the potential intrusion. Possible values include: **ATTACK** Attack action event **TR** Traffic regulation trace action event **SCANG** Scan global action event **SCANE** Scan event action event	ATTACK
Event correlator	Char(4)	Unique identifier for this specific intrusion event. You can use this identifier to correlate this audit record with other intrusion detection information.	This field is hidden and appears blank. Press F11 to display the information.
Event type	Char(8)	Identifies the type of potential intrusion that was detected. The possible values include: **MALFPKT** Malformed packet **FLOOD** Flood event **ICMPRED** Internet Control Message Protocol (ICMP) redirect **PERPECH** Perpetual echo **IPFRAG** IP fragment **RESTPROT** Restricted IP protocol (RESTP)	RESTP
Suspected packet	Char(1002)	This variable-length, binary field might contain up to the first 1000 bytes of the IP packet that is associated with the detected event. The first two bytes of this field contain the length of the suspected packet information.	This field is hidden and appears blank. Press F11 to display the information.

Related tasks

"Audit intrusion detection activities" on page 9
Learn how to audit intrusion detection activities. If the intrusion detection system (IDS) flags a suspicious event, it writes an IM audit record.

Scan events

The intrusion detection system detects scans to individual ports.

Through statistics gathering and auditing, the intrusion detection system determines whether the system has been the target of a global scan. When the TCP/IP stack detects an intrusion event is detected, the stack calls the intrusion detection function and generates statistics and audit records.

If an IDS scan policy does not exist in the IDS policy file, no action is taken. If an IDS scan policy exists, the intrusion detection system creates an audit record when it detects a scan event.

TCP port scans

You can classify TCP events as normal, possibly suspicious, or highly suspicious. In the IDS policy, you can define restricted ports that no one can use.

The intrusion detection system (IDS) scans and classifies the following types of TCP events. Typically, the TCP/IP stack discards the suspicious event.

Table 2. TCP scan events classified as suspicious

Scan Event	TCP/IP Connection State	Event Classification
Receive any packet	Unbound, not restricted	Possibly suspicious (possibly a failed application)
Receive a packet with the reset (RST) bit set in the TCP header. (In this situation, the host immediately terminates the connection, which results in a denial of service until that connection is reestablished.)	Half-open connection	Possibly suspicious (peer covering tracks)
Final timeout	Any connected state	Possibly suspicious (peer abandoned connection)
Receive unexpected flags	Any	Highly suspicious
Receive any packet from a restricted TCP/IP port	This TCP/IP port is RESERVED	Highly suspicious
Final timeout	Half-open connection	Highly suspicious (peer abandoned handshake)

User Datagram Protocol (UDP) port scans

You can classify UDP events as normal, possibly suspicious, or highly suspicious. In the IDS policy, you can define restricted ports that no one can use. Any datagram received for a restricted port is treated as a highly suspicious event. Datagrams received for unbound but unrestricted ports are treated as possibly suspicious events. Datagrams received for bound ports that are rejected by the QoS policy or FW filters are treated as possibly suspicious. All other datagrams received for bound ports are treated as normal events.

If an IDS scan policy does not exist in the IDS policy file, no action is taken. If an IDS scan policy exists, the intrusion detection system creates an audit record when it detects a scan event.

Table 3. UDP scan events

Scan Event	TCP/IP Connection State	Event Classification
QoS policy rejects packet	Bound	Normal
Receive any packet	Bound	Normal
FW filtering rejects packet	Bound	Possibly suspicious
Receive any packet	Unbound	Possibly suspicious (possibly failed application)
Receive any packet	This TCP/IP port is restricted	Highly suspicious

Internet Control Message Protocol (ICMP) port scans

You can use ICMP requests to map network topology. Any request sent to a subnet base or broadcast address is treated as a highly suspicious event. Echo (ping) requests and timestamp requests are very

common, so they are treated as normal events. The intrusion detection system audits ICMP redirect events.

Attack events

The intrusion detection system detects different types of attack events and writes an IM audit record in the QAUDJRN audit journal.

The intrusion detection system detects the following types of attack events:
- Malformed packets
- Denial of service floods
- ICMP redirect messages
- Perpetual echo on UDP ports
- IP fragments
- Restricted IP options and protocols
- Fragmented packets

The number of audit records that the system generates depends on the value of the maximum event message in the IDS policy.

Malformed packet events

A malformed packet is built in such a way as to cause a system to crash or hang when it is processed. When the IDS policy detects a malformed packet, it writes an audit record. The TCP/IP stack deletes the malformed packets.

Fragment restriction events

An invalid fragment overlays IP or transport headers in an attempt to bypass firewall checks. On the iSeries system, it is not possible to overlay an IP header. The TCP/IP stack checks to ensure that the first fragment of a fragmented datagram is a minimum of 576 bytes. The stack also checks that each fragment beyond the first one has an offset of greater than 256 bytes.

The IDS policy audits invalid IP fragments.

IP option restrictions

The IP options field in a datagram is a variable-length list of optional information. Some of the IP Options, such as Loose Source Route, can be used in network attacks. You can use the IDS policy to restrict which IP options that an inbound packet can contain. For example, you can specify whether an inbound packet with a restricted IP option be ignored or audited. You also can generate statistics on the number of inbound packets with restricted IP options.

IP protocol restrictions

The IP protocol field is an 8-bit field in the IP header. Undefined IP protocols are sometimes used to establish back door attacks on the network. You can use the IDS policy to restrict which IP protocols that an inbound packet can contain. The policy can specify whether an inbound packet with a restricted IP protocol be audited. You also can generate statistics on the number of inbound packets with restricted IP protocols.

SYN flood events

TCP SYN flood events create a large number of half-open sockets. These flood events fill up the socket connection backlog for a given application and deny valid connections from being accepted. A SYN flood

event spoofs the source IP address with the address of an unreachable system. The IDS policy flags SYN flood events and writes an audit record.

ICMP redirect events

You can use Internet Control Message Protocol (ICMP) redirect messages to override intended network routes. You can specify the IGNOREREDIRECT option in the IDS policy file to either ignore or process ICMP redirect messages.

Perpetual echo on UDP ports

You can use port 7, which is called the echo port, to test a UDP connection. (Both the source port and target port are set to port 7, which causes each port to echo back what it gets.) Whatever data is sent through UDP is echoed back. A perpetual echo is an attack on UDP port 7. The TCP/IP stack detects the event if the source port is equal to the target port. If there is an IDS policy for attack-type events, the system writes an audit record whenever it detects a perpetual echo attack on the UDP port.

Related information for intrusion detection

Listed here are the product manuals and IBM® Redbooks™ (in PDF format), Web sites, and information center topics that relate to the intrusion detection topic. You can view or print any of the PDFs.

Manuals

- iSeries Security Reference  (13 682 KB) .

Other information

- Plan and set up system security which discusses techniques for detecting other types of intrusions.
- Quality of service which discusses how to use the QoS commands to activate an intrusion detection policy.

Saving PDF files

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

1. Right-click the PDF in your browser (right-click the link above).
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 (www.adobe.com/products/acrobat/readstep.html)  .

Appendix. 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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, 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 information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, IBM License Agreement for Machine Code, 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.

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.

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_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This intrusion detection publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of i5/OS.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

e(logo)server
eServer
i5/OS
IBM

IBM (logo)
iSeries
Redbooks

Other company, product, and service names may be trademarks or service marks of others.

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.

Printed in USA