



Testing and Troubleshooting

This chapter contains the following sections:

- [Debugging Mail Flow Using Test Messages: Trace](#), on page 1
- [Using the Listener to Test the Email Gateway](#), on page 7
- [Troubleshooting the Network](#), on page 10
- [Troubleshooting the Listener](#), on page 15
- [Troubleshooting Email Delivery From the Appliance](#), on page 17
- [Troubleshooting Performance](#), on page 19
- [Web Interface Appearance and Rendering Issues](#) , on page 20
- [Responding to Alerts](#) , on page 20
- [Troubleshooting Hardware Issues](#) , on page 20
- [Remotely Resetting Email Gateway Power](#) , on page 20
- [Working with Technical Support](#), on page 21

Debugging Mail Flow Using Test Messages: Trace

You can use **System Administration > Trace** page (the equivalent of the trace command in the CLI) to debug the flow of messages through the system by emulating sending a test message. The Trace page (and **trace** CLI command) emulates a message as being accepted by a listener and prints a summary of features that would have been “triggered” or affected by the current configuration of the system (*including uncommitted changes*). The test message is not actually sent. The Trace page (and **trace** CLI command) can be a powerful troubleshooting or debugging tool, especially if you have combined many of the advanced features available on the email gateway.



Note Trace is not effective for testing file reputation scanning.

The Trace page (and **trace** CLI command) prompts you for the input parameters listed in the following table.

Table 1: Input for the Trace page

Value	Description	Example
Source IP address	Type the IP address of the remote client to mimic the source of the remote domain. This can be an Internet Protocol version 4 (IPv4) or version 6 (IPv6) address. Note: The trace command prompts for an IP address and a fully-qualified domain name. It does <i>not</i> attempt to reverse the IP address to see if it matches the fully-qualified domain name. The trace command does not allow the fully-qualified domain name field to be blank, so it is impossible to test a scenario where the DNS does not reverse match properly.	203.45.98.109 2001:0db8:85a3::8a2e:0370:7334
Fully Qualified Domain Name of the Source IP	Type the fully-qualified remote domain name to mimic. If left null, a reverse DNS lookup will be performed on the source IP address.	smtp.example.com
Listener to Trace Behavior on	Choose from the list of listeners configured on the system to emulate sending the test message to.	InboundMail
Network Owner Organization ID	Type the unique identification number of the network owner, or allow the system to Lookup network owner ID associated with source IP address. You can view this information if you added network owners to sender groups via the GUI.	34
IP Reputation Score	Type the IP Reputation score you want to provide for the spoofed domain, or allow the system to look up the IP Reputation score associated with the source IP address. This can be helpful when testing policies that use IP Reputation scores. Note that manually entered IP Reputation scores are not passed to the Context Adaptive Scanning Engine (CASE). See Editing IP Reputation Filtering Score Thresholds for a Listener for more information.	-7.5
Envelope Sender	Type the Envelope Sender of the test message.	admin@example.net
Envelope Recipients	Type a list of recipients for the test message. Separate multiple entries with commas.	joe frank@example.com

Value	Description	Example
Message Body	Type the message body for the test message, including headers. Type a period on a separate line to end entering the message body. Note that “headers” are considered part of a message body (separated by a blank line), and omitting headers, or including poorly formatted ones can cause unexpected trace results.	To: 1@example.com From: ralph Subject: Test this is a test message .

After you have entered the values, click **Start Trace**. A summary of all features configured on the system affecting the message is printed.

You can upload message bodies from your local file system. (In the CLI, you can test with message bodies you have uploaded to the `/configuration` directory. See [FTP, SSH, and SCP Access](#) for more information on placing files for import onto the email gateway.)

After the summary is printed, you are prompted to view the resulting message and re-run the test message again. If you enter another test message, the Trace page and the trace command uses any previous values from the above table you entered.



Note The sections of configuration tested by the trace command listed in the following table are performed *in order*. This can be extremely helpful in understanding how the configuration of one feature affects another. For example, a recipient address transformed by the domain map feature will affect the address as it is evaluated by the RAT. A recipient that is affected by the RAT will affect the address as it is evaluated by alias table, and so on.

Table 2: Viewing Output When Performing a Trace

trace Command Section	Output
Host Access Table (HAT) and Mail Flow Policy Processing	<p>The Host Access Table settings for the listener you specified are processed. The system reports which entry in the HAT matched from the remote IP address and remote domain name you entered. You can see the default mail flow policies and sender groups and which one matched the given entries.</p> <p>If the email gateway was configured to reject the connection (either through a REJECT or TCPREFUSE access rule), the trace command exits at the point in the processing.</p> <p>For more information on setting HAT parameters, see Understanding Predefined Sender Groups and Mail Flow Policies.</p>
<p>Envelope Sender Address Processing</p> <p>These sections summarize how the email gateway configuration affects the Envelope Sender you supply. (That is, how the MAIL FROM command would be interpreted by the configuration of the email gateway.) The <code>trace</code> command prints “Processing MAIL FROM:” before this section.</p>	

trace Command Section	Output
Default Domain	<p>If you specified that a listener to change the default sender domain of messages it receives, any change to the Envelope Sender is printed in this section.</p> <p>For more information, see Configuring the Gateway to Receive Email.</p>
Masquerading	<p>If you specified that the Envelope Sender of a message should be transformed, the change is noted here. You enable masquerading for the Envelope Sender on private listeners using the listenerconfig -> edit -> masquerade -> config subcommands.</p> <p>For more information, see Configuring Routing and Delivery Features.</p>
<p>Envelope Recipient Processing</p> <p>These sections summarize how the email gateway affects the Envelope Recipients you supply. (That is, how the RCPT TO command would be interpreted by the configuration of the email gateway.) The trace command prints “ Processing Recipient List: ” before this section.</p>	
Default Domain	<p>If you specified that a listener to change the default sender domain of messages it receives, any changes to the Envelope Recipients are printed in this section.</p> <p>For more information, see Configuring the Gateway to Receive Email.</p>
Domain Map Translation	<p>The domain map feature transforms the recipient address to an alternate address. If you specified any domain map changes and a recipient address you specified matches, the transformation is printed in this section.</p> <p>For more information, see Configuring Routing and Delivery Features.</p>
Recipient Access Table (RAT)	<p>Each Envelope Recipient that matches an entry in the RAT is printed in this section, in addition to the policy and parameters. (For example, if a recipient was specified to bypass limits in the listener’s RAT.)</p> <p>For more information on specifying recipients you accept, see Configuring the Gateway to Receive Email.</p>
Alias Table	<p>Each Envelope Recipient that matches an entry in the alias tables configured on the email gateway (and the subsequent transformation to one or more recipient addresses) is printed in this section.</p> <p>For more information, see Configuring Routing and Delivery Features.</p>
<p>Pre-Queue Message Operations</p> <p>These sections summarize how the email gateway affects each message after the message contents have been received, but before the messages are enqueued on the work queue. This processing occurs before the final 250 ok command is returned to the remote MTA.</p> <p>The trace command prints “Message Processing : ” before this section.</p>	

trace Command Section	Output
Virtual Gateways	<p>The altsrchost command assigns messages to a specific interface, based on a match of the Envelope Sender's full address, domain, or name, or IP address. If an Envelope Sender matches entries from the altsrchost command, that information is printed in this section.</p> <p>Note that the virtual gateway address assigned at this point may be overridden by message filter processing below.</p> <p>For more information, see Configuring Routing and Delivery Features.</p>
Bounce Profiles	<p>Bounce profiles are applied at three different points in the processing. This is the first occurrence. If a listener has a bounce profile assigned to it, it is assigned at this point in the process. That information is printed in this section.</p> <p>For more information, see Configuring Routing and Delivery Features.</p>
<p>Work Queue Operations</p> <p>The following group of functions are performed on messages in the work queue. This occurs after the message has been accepted from the client, but before the message is enqueued for delivery on a destination queue. "Messages in Work Queue" is reported by the status and status detail commands.</p>	
Masquerading	<p>If you specified that the To:, From:, and CC: headers of messages should be masked (either from a static table entered from a listener or via an LDAP query), the change is noted here. You enable masquerading for the message headers on private listeners using the listenerconfig -> edit -> masquerade -> config subcommands.</p> <p>For more information, see Configuring Routing and Delivery Features.</p>
LDAP Routing	<p>If LDAP queries have been enabled on a listener, the results of LDAP acceptance, re-routing, masquerading, and group queries are printed in this section.</p> <p>For more information, see LDAP Queries.</p>
Message Filters Processing	<p>All messages filters that are enabled on the system are evaluated by the test message at this point. For each filter, the rule is evaluated, and if the end result is "true," each of the actions in that filter are then performed in sequence. A filter may contain other filters as an action, and the nesting of filters is unlimited. If a rule evaluates to "false" and a list of actions is associated with an else clause, those actions are evaluated instead. The results of the message filters, processed in order, are printed in this section.</p> <p>See Using Message Filters to Enforce Email Policies.</p>

trace Command Section	Output
<p>Mail Policy Processing</p> <p>The mail policy processing section displays the Anti-Spam, Anti-Virus, Outbreak Filters feature, and disclaimer stamping for all recipients you supplied. If multiple recipients match multiple policies in Email Security Manager, the following sections will be repeated for each matching policy. The string: “Message Going to” will define which recipients matched which policies.</p>	
<p>Anti-Spam</p>	<p>This section notes messages that are not flagged to be processed by anti-spam scanning. If messages are to be processed by anti-spam scanning for the listener, the message is processed and the verdict returned is printed. If the email gateway is configured to bounce or drop the messages based on the verdict, that information is printed and the trace command processing stops.</p> <p>Note: This step is skipped if anti-spam scanning is unavailable on the system. If anti-spam scanning is available but has not been enabled with a feature key, that information is also printed in this section.</p> <p>See Managing Spam and Graymail.</p>
<p>Anti-Virus</p>	<p>This section notes messages that are not flagged to be processed by anti-virus scanning. If messages are to be processed by anti-virus scanning for the listener, the message is processed and the verdict returned is printed. If the email gateway is configured to “clean” infected messages, that information is noted. If configured to bounce or drop the messages based on the verdict, that information is printed and the trace command processing stops.</p> <p>Note: This step is skipped if anti-virus scanning is unavailable on the system. If anti-virus scanning is available but has not been enabled with a feature key, that information is also printed in this section.</p> <p>See the Anti-Virus.</p>
<p>Content Filters Processing</p>	<p>All content filters that are enabled on the system are evaluated by the test message at this point. For each filter, the rule is evaluated, and if the end result is “true,” each of the actions in that filter are then performed in sequence. A filter may contain other filters as an action, and the nesting of filters is unlimited. The results of the content filters, processed in order, are printed in this section.</p> <p>See Content Filters.</p>
<p>Outbreak Filters Processing</p>	<p>This section notes that messages that contain attachments are to bypass the Outbreak Filters feature. If messages are to be processed by Outbreak Filters for the recipient, the message is processed and the evaluation. If the email gateway is configured to quarantine, bounce, or drop the messages based on the verdict, that information is printed and the processing stops.</p> <p>See Outbreak Filters.</p>

trace Command Section	Output
Footer Stamping	This section notes whether a footer text resource was appended to the message. The name of the text resource is displayed. See Message Disclaimer Stamping in Text Resources .
<p>Delivery Operations</p> <p>The following sections note operations that occur when a message is delivered. The trace command prints “Message Enqueued for Delivery” before this section.</p>	
Global Unsubscribe per Domain and per User	<p>If any recipients you specified as input for the trace command match recipients, recipient domains, or IP addresses listed in the in the Global Unsubscribe feature, any unsubscribed recipient addresses are printed in this section.</p> <p>See Configuring Routing and Delivery Features.</p>
Sender Domain Reputation	<p>Displays the processing results of messages that are blocked by the reputation verdict of the sender domain.</p> <p>The following is a sample example of the processing results:</p> <ul style="list-style-type: none"> • Reputation: Neutral • Domain Age: 33 years 9 months 6 days • Category: N/A • Action: Message Dropped
<p>Final Result</p> <p>When all processing has been printed, you are prompted with the final result. In the CLI, Answer y to the question, “Would you like to see the resulting message?” to view the resulting message.</p>	


Using the Listener to Test the Email Gateway

“Sinkhole” listeners allow you to test your message generation systems and to also get a rough measure of receiving performance. Two types of sinkhole listeners are *queueing* and *non-queueing*.

- The queueing listener saves the message to the queue, but then immediately deletes it. Use a queueing listener when you are interested in measuring the performance of the entire injection portion of your message generation system.
- The non-queueing listener accepts a message, and then immediately deletes it without saving it. Use the non-queueing listener when you want to troubleshoot the connection from your message generation system to the email gateway.

For example, in the following figure, you could create a sinkhole listener “C” to mirror the private listener labeled “B.” A non-queueing version tests the performance path of the system from the groupware client to the groupware server to the email gateway. A queueing version tests that same path *and* the email gateway’s ability to enqueue messages and prepare them for delivery via SMTP.

Figure 1: Sinkhole Listener for an Enterprise Gateway



In the following example, the `listenerconfig` command is used to create a sinkhole queueing listener named `Sinkhole_1` on the Management interface. This Host Access Table (HAT) for the listener is then edited to accept connections from the following hosts:

- `yoursystem.example.com`
- `10.1.2.29`
- `badmail.tst`
- `.tst`



Note The final entry, `.tst`, configures the listener so that any host in the `.tst` domain can send email to the listener named `Sinkhole_1`.

Example

```
mail3.example.com> listenerconfig

Currently configured listeners:

1. InboundMail (on PublicNet, 192.168.2.1) SMTP Port 25 Public
2. OutboundMail (on PrivateNet, 192.168.1.1) SMTP Port 25 Private

Choose the operation you want to perform:

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

[ ]> new

Please select the type of listener you want to create.

1. Private
2. Public
```



```
3. Sinkhole

[2]> 3

Do you want messages to be queued onto disk? [N]> y

Please create a name for this listener (Ex: "OutboundMail"):

[]> Sinkhole_1

Please choose an IP interface for this Listener.

1. Management (192.168.42.42/24: mail3.example.com)
2. PrivateNet (192.168.1.1/24: mail3.example.com)
3. PublicNet (192.168.2.1/24: mail3.example.com)

[1]> 1

Choose a protocol.

1. SMTP
2. QMQP

[1]> 1

Please enter the IP port for this listener.

[25]> 25

Please specify the systems allowed to relay email through the IronPort C60.
Hostnames such as "example.com" are allowed.
Partial hostnames such as ".example.com" are allowed.
IP addresses, IP address ranges, and partial IP addresses are allowed.
Separate multiple entries with commas.

[]> yoursystem.example.com, 10.1.2.29, badmail.tst, .tst

Do you want to enable rate limiting per host? (Rate limiting defines
the maximum number of recipients per hour you are willing to receive from a remote
domain.) [N]> n

Default Policy Parameters
=====

Maximum Message Size: 100M

Maximum Number Of Connections From A Single IP: 600

Maximum Number Of Messages Per Connection: 10,000

Maximum Number Of Recipients Per Message: 100,000

Maximum Number Of Recipients Per Hour: Disabled
```

```

Use SenderBase for Flow Control: No

Spam Detection Enabled: No

Virus Detection Enabled: Yes

Allow TLS Connections: No

Allow SMTP Authentication: No

Require TLS To Offer SMTP authentication: No

Would you like to change the default host access policy? [N]> n

Listener Sinkhole_1 created.

Defaults have been set for a Sinkhole Queuing listener.

Use the listenerconfig->EDIT command to customize the listener.

Currently configured listeners:

1. Sinkhole_1 (on Management, 192.168.42.42) SMTP Port 25 Sinkhole Queuing
2. InboundMail (on PublicNet, 192.1681.1) SMTP Port 25 Public
3. OutboundMail (on PrivateNet, 192.168.1.1) SMTP Port 25 Private

Choose the operation you want to perform:

- NEW - Create a new listener.
- EDIT - Modify a listener.
- DELETE - Remove a listener.
- SETUP - Change global settings.

[ ]>

```



Note Remember to issue the commit command for these changes to take effect.

After you have configured a sinkhole queuing listener and modified the HAT to accept connections from your injection system, use your injection system to begin sending email to the email gateway. Use the `status`, `status detail`, and `rate` commands to monitor system performance. You can also monitor the system via the Graphical User Interface (GUI). For more information, see:

- [Monitoring Using the CLI](#)
- [Other Tasks in the GUI](#)

Troubleshooting the Network

If you suspect that the email gateway has network connectivity issues, first confirm that the email gateway is working properly.

Testing the Network Connectivity of the Email Gateway

Procedure

- Step 1** Connect to the system and log in as the administrator. After successfully logging in, the following messages are displayed:

```
Last login: day month date hh:mm:ss from IP address
```

```
Copyright (c) 2001-2003, IronPort Systems, Inc.
```

```
AsyncOS x.x for Cisco
```

```
Welcome to the Cisco Messaging Gateway Appliance(tm)
```

- Step 2** Use the `status` or `status detail` commands.

```
mail3.example.com> status
```

or

```
mail3.example.com> status detail
```

The `status` command returns a subset of the monitored information about email operations. The statistics returned are grouped into two categories: counters and gauges. For complete monitoring information about email operations including rates, use the `status detail` command. Counters provide a running total of various events in the system. For each counter, you can view the total number of events that have occurred since the counter was reset, since the last system reboot, and over the system's lifetime. (For more information, see [Monitoring Using the CLI](#).)

- Step 3** Use the `mailconfig` command to send mail to a known working address.

The `mailconfig` command generates a human-readable file including all configuration settings available to the email gateway. Attempt to send the file from the email gateway to a known working email address to confirm that the email gateway is able to send email over the network.

```
mail3.example.com> mailconfig
```

```
Please enter the email address to which you want to send the  
configuration file.
```

```
Separate multiple addresses with commas.
```

```
[> user@example.com
```

```
Do you want to include passphrases? Please be aware that a configuration without  
passphrases will fail when reloaded with loadconfig. [N]> y
```

```
The configuration file has been sent to user@example.com.
```

```
mail3.example.com>
```

Troubleshooting

After you have confirmed that the email gateway is active on the network, use the following commands to pinpoint any network problems.

- You can use the `netstat` command to display network connections (both incoming and outgoing), routing tables, and a number of network interface statistics, including the following information:
 - List of active sockets
 - State of network interfaces
 - Contents of routing tables
 - Size of the listen queues
 - Packet traffic information
- You can use the `diagnostic -> network -> flush` command to flush all network related caches.
- You can use the `diagnostic -> network -> arpshow` command to show the system ARP cache.
- You can use the `packetcapture` command to intercept and display TCP/IP and other packets being transmitted or received over a network to which the computer is attached.

To use `packetcapture`, set the network interface and the filter. The filter uses the same format the UNIX `tcpdump` command. Use `start` to begin the packet capture and `stop` to end it. After stopping the capture, you need to use SCP or FTP to download the files from the `/pub/captures` directory. For more information, see [Running a Packet Capture, on page 25](#).

- Use the `ping` command to a known working host to confirm that the email gateway has an active connection on the network and is able to reach specific segments of your network.

The `ping` command allows you to test connectivity to a network host from the email gateway.

```
mail3.example.com> ping
```

```
Which interface do you want to send the pings from?
```

1. Auto
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)

```
[1]> 1
```

```
Please enter the host you wish to ping.
```

```
[> anotherhost.example.com
```

```
Press Ctrl-C to stop.
```

```

PING anotherhost.example.com (x.x.x.x): 56 data bytes
64 bytes from 10.19.0.31: icmp_seq=9 ttl=64 time=0.133 ms
64 bytes from 10.19.0.31: icmp_seq=10 ttl=64 time=0.115 ms
^C
--- anotherhost.example.com ping statistics ---
11 packets transmitted, 11 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.115/0.242/1.421/0.373 ms

```



Note You must use Control-C to end the ping command.

- Use the `traceroute` command to test connectivity to a network host from the email gateway and debug routing issues with network hops.

```

mail3.example.com> traceroute

Which interface do you want to trace from?

1. Auto
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)

[1]> 1

Please enter the host to which you want to trace the route.

[]> 10.1.1.1

Press Ctrl-C to stop.

traceroute to 10.1.1.1 (10.1.1.1), 64 hops max, 44 byte packets
 1 gateway (192.168.0.1) 0.202 ms 0.173 ms 0.161 ms
 2 hostname (10.1.1.1) 0.298 ms 0.302 ms 0.291 ms

mail3.example.com>

```

- Use the `diagnostic -> network -> smtping` command to test a remote SMTP server.
- Use the `nslookup` command to check the DNS functionality.

The `nslookup` command can confirm that the email gateway is able to reach and resolve hostnames and IP addresses from a working DNS (domain name service) server.

```

mail3.example.com> nslookup

Please enter the host or IP to resolve.

```

```
[ ]> example.com

Choose the query type:

1. A
2. CNAME
3. MX
4. NS
5. PTR
6. SOA
7. TXT

[1]>

A=192.0.34.166 TTL=2d
```

Table 3: Checking DNS Functionality: Query Types

A	the host's Internet address
CNAME	the canonical name for an alias
MX	the mail exchanger
NS	the name server for the named zone
PTR	the hostname if the query is an Internet address, otherwise the pointer to other information
SOA	the domain's "start-of-authority" information
TXT	the text information

- Use the `tophosts` command via the CLI or the GUI, and sort by Active Recipients.

The `tophosts` command returns a list of the top 20 recipient hosts in queue. This command can help you determine if network connectivity problems are isolated to a single host or group of hosts to which you are attempting to send email. (For more information, see "Determining the Make-up of the Mail Queue".)

```
mail3.example.com> tophosts

Sort results by:

1. Active Recipients
2. Connections Out
3. Delivered Recipients
4. Soft Bounced Events
5. Hard Bounced Recipients
```

```
[1]> 1

Status as of: Mon Nov 18 22:22:23 2003

ActiveConn.Deliv.SoftHard

# Recipient HostRecipOutRecip.BouncedBounced

1 aol.com36510255218

2 hotmail.com29071982813

3 yahoo.com13461231119

4 excite.com9838494

5 msn.com8427633 29

^C
```

- “Drill-down” to use the `hoststatus` command on the top domains listed from the `tophosts` command results.

The `hoststatus` command returns monitoring information about email operations relating to a specific recipient host. DNS information stored in the AsyncOS cache and the last error returned from the recipient host are also given. Data returned is cumulative since the last `resetcounters` command. (For more information, see [Monitoring the Status of a Mail Host](#).)

Using the `hoststatus` command on the top domains can isolate the performance issues with DNS resolution to the either the email gateway or the internet. For example, if the `hoststatus` command for the top active recipient host shows many pending outbound connections, then try to determine if that particular host is down or unreachable, or if the email gateway cannot connect to all or the majority of hosts.

- Check firewall permissions.

The email gateway may need all of the following ports to be opened in order to function properly: ports 20, 21, 22, 23, 25, 53, 80, 123, 443, and 628. (See [Firewall Information](#).)

- Send email from the email gateway on your network to `dnscheck@ironport.com`

Send an email from within your network to `dnscheck@ironport.com` to perform basic DNS checks on your system. An auto-responder email will respond with the results and details of the following four tests:

DNS PTR Record - Does the IP address of the Envelope From match the PTR record for the domain?

DNS A Record - Does the PTR record for the domain match the IP address of the Envelope From?

HELO match - Does the domain listed in the `SMTP HELO` command match the DNS hostname in the Envelope From?

Mail server accepting delayed bounce messages - Does the domain listed in the `SMTP HELO` command have MX records that resolve IP addresses for that domain?

Troubleshooting the Listener

If you suspect problems with injecting email, use the following strategies:

- Confirm the IP address that you are injecting from, and then use the `listenerconfig` command to check for allowed hosts.

Is the IP address allowed to connect to the listener you have created? Use the `listenerconfig` command to examine the Host Access Table (HAT) for the listener. Use these commands to print the HAT for a listener:

```
listenerconfig -> edit -> listener_number -> hostaccess -> print
```

The HAT can be configured to refuse connections by IP address, block of IP addresses, hostname, or domains. For more information, see “Specifying Hosts that are Allowed to Connect”.

You can also use the `limits` subcommand to check the maximum number of connections allowed for a listener:

```
listenerconfig -> edit -> listener_number -> limits
```

- On the machine that you are injecting from, use Telnet or FTP to manually connect to the email gateway. For example:

```
injection_machine% telnet appliance_name
```

You can also use the `telnet` command within the email gateway itself to connect from the listener to the actual email gateway:

```
mail3.example.com> telnet
```

```
Please select which interface you want to telnet from.
```

1. Auto
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)

```
[1]> 3
```

```
Enter the remote hostname or IP.
```

```
[> 193.168.1.1
```

```
Enter the remote port.
```

```
[25]> 25
```

```
Trying 193.168.1.1...
```

```
Connected to 193.168.1.1.
```

```
Escape character is '^'.
```

If you cannot connect from one interface to another, you may have issues with the way in which the email gateway's Management and Data1 and Data2 interfaces are connected to your network. See [FTP, SSH, and SCP Access](#) for more information. You can telnet to port 25 of the listener and enter SMTP commands manually (if you are familiar with the protocol).

- Examine the IronPort text mail logs and injection debug logs to check for receiving errors.

Injection debug logs record the SMTP conversation between the email gateway and a specified host connecting to the system. Injection debug logs are useful for troubleshooting communication problems between the email gateway and a client initiating a connection from the Internet. The log records all bytes transmitted between the two systems and classifies them as “Sent to” the connecting host or “Received from” the connecting host.

For more information, see [Using Text Mail Logs](#) and [Using Injection Debug Logs](#).

Troubleshooting Email Delivery From the Appliance

If you suspect problems with delivering email from the email gateway, try the following strategies:

- Determine if the problem is domain-specific.

Use the `tophosts` command to get immediate information about the email queue and determine if a particular recipient domain has delivery problems.

Are there problem domains returned when you sort by “Active Recipients?”

When you sort by Connections Out, does any one domain reach the maximum connections specified for a listener? The default maximum number of connections for a listener is 600. The default maximum system-wide number of connections is 10,000 (set by the `deliveryconfig` command). You can examine the maximum number of connections for a listener using the command:

```
listenerconfig -> edit -> listener_number -> limits
```

Are the connections for a listener further limited by the `destconfig` command (either by system maximum or by Virtual Gateway addresses)? Use this command to examine the `destconfig` connection limits:

```
destconfig -> list
```

- Use the `hoststatus` command.

“Drill-down” using the `hoststatus` command on the top domains listed from the results listed by the `tophosts` command.

Is the host available and accepting connections?

Are there problems with one specific MX record mail server for the given host?

The `hoststatus` command reports the last “5XX” status code and description returned by the host if there is a 5XX error (Permanent Negative Completion reply) for the specified host. If the last outgoing TLS connection to the host failed, the `hoststatus` command displays the reason why it failed.

- Configure and/or examine the domain debug, bounce, and text mail logs to check if the recipient host is available.

Domain debug logs record the client and server communication during an SMTP conversation between the email gateway and a specified recipient host. This log file type can be used to debug issues with specific recipient hosts.

For more information, see [Using Domain Debug Logs](#).

Bounce logs record all information pertaining to each bounced recipient.

For more information, see [Using Bounce Logs](#).

Text mail logs contain details of email receiving, email delivery and bounces. These logs are a useful source of information to understand delivery of specific messages and to analyze system performance.

For more information, see [Using Text Mail Logs](#).

- Use the `telnet` command to connect from the email gateway to the problem domain:

```
mail3.example.com> telnet

Please select which interface you want to telnet from.

1. Auto
2. Management (192.168.42.42/24: mail3.example.com)
3. PrivateNet (192.168.1.1/24: mail3.example.com)
4. PublicNet (192.168.2.1/24: mail3.example.com)

[1]> 1

Enter the remote hostname or IP.

[]> problemdomain.net

Enter the remote port.

[25]> 25
```

- You can use the `tlsverify` command to establish an outbound TLS connection on demand and debug any TLS connection issues concerning a destination domain. To create the connection, specify the domain to verify against and the destination host. AsyncOS checks the TLS connection based on the Required (Verify) TLS setting.

```
mail3.example.com> tlsverify

Enter the TLS domain to verify against:

[]> example.com

Enter the destination host to connect to. Append the port (example.com:26) if you are
not connecting on port 25:

[example.com]> mx.example.com:25

Connecting to 1.1.1.1 on port 25.

Connected to 1.1.1.1 from interface 10.10.10.10.

Checking TLS connection.

TLS connection established: protocol TLSv1, cipher RC4-SHA.

Verifying peer certificate.

Verifying certificate common name mx.example.com.

TLS certificate match mx.example.com
```

```
TLS certificate verified.  
TLS connection to 1.1.1.1 succeeded.  
TLS successfully connected to mx.example.com.  
TLS verification completed.
```

Troubleshooting Performance

If you suspect that there are performance problems with the email gateway, utilize the following strategies:

- Use the `rate` and `hostrate` commands to check the current system activity.

The `rate` command returns real-time monitoring information about email operations. For more information, see [Displaying Real-time Activity](#).

The `hostrate` command returns real-time monitoring information for a specific host.

- Use the `status` command to cross-check the historical rates to check for degradation.
- Use the `status detail` command to check the RAM utilization.

You can use the `status detail` command to quickly see the system's RAM, CPU, and Disk I/O utilization.



Note RAM utilization should always be less than 45%. If RAM utilization exceeds 45%, then, the email gateway will enter “resource conservation mode;” it initiates a “back-off” algorithm to prevent over-subscription of resources and sends out the following email alert:

```
This system (hostname: hostname) has entered a 'resource conservation' mode in order  
to  
prevent the rapid depletion of critical system resources.
```

```
RAM utilization for this system has exceeded the resource conservation threshold of  
45%.  
The allowed injection rate for this system will be gradually decreased as RAM  
utilization approaches 60%.
```

This situation occurs only with an aggressive injection with poor deliverability facilities. If you encounter RAM utilization exceeding 45%, check the number of messages in the queue and see if a particular domain is down or unavailable for delivery (via the `hoststatus` or `hostrate` commands). Also check the status of the system and ensure that delivery is not suspended. If after stopping the injection you continue to experience a high RAM utilization, contact Cisco Customer Support.

- Is the problem specific to one domain?

Use the `tophosts` command to get immediate information about the email queue and determine if a particular recipient domain has delivery problems.

Check the size of the queue. You can delete, bounce, suspend, or redirect messages in the email queue to manage its size, or to deal with recipients to a specific, problematic domain. For more information, see [Managing the Email Queue](#). Use these commands:

- `deleterecipients`
- `bouncerecipients`
- `redirectrecipients`
- `suspenddel / resumedel`
- `suspendlistener / resumelister`

Use the `tophosts` command to check the number of soft and hard bounces. Sort by “Soft Bounced Events” (option 4) or “Hard Bounced Recipients” (option 5). If the performance for a particular domain is problematic, use the commands above to manage the delivery to that domain.

Web Interface Appearance and Rendering Issues

See [Overriding Internet Explorer Compatibility Mode](#).

Responding to Alerts

- [Troubleshooting Alerts That Miscellaneous Disk Usage is Approaching the Quota](#), on page 20

Troubleshooting Alerts That Miscellaneous Disk Usage is Approaching the Quota

Problem

You receive an alert that Miscellaneous disk usage is approaching its quota.

Solution

You can either increase the quota or delete files. See [Managing Disk Space for the Miscellaneous Quota](#).

Troubleshooting Hardware Issues

Lights on the front and/or rear panels of your hardware email gateway indicate health and status of your email gateway. For descriptions of these indicators, see the hardware guides, such as the *Cisco x90s Series Content Security Appliances Installation and Maintenance Guide*, available from <http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-installation-guides-list.html>.

Specifications for your email gateway, such as temperature ranges, are also available in these documents.

Remotely Resetting Email Gateway Power

If the email gateway requires a hard reset, you can reboot the email gateway chassis remotely using a third-party Intelligent Platform Management Interface (IPMI) tool.

Restrictions

- Remote power cycling is available only on certain hardware.
For specifics, see [Enabling Remote Power Cycling](#).
- If you want to be able to use this feature, you must enable it in advance, before you need to use it.
For details, see [Enabling Remote Power Cycling](#).
- Only the following IPMI commands are supported:
 - **status, on, off, cycle, reset, diag, soft**
 - Issuing unsupported commands will produce an “insufficient privileges” error.

Before You Begin

- Obtain and set up a utility that can manage devices using IPMI version 2.0.
- Understand how to use the supported IPMI commands. See the documentation for your IPMI tool.

Procedure

-
- Step 1** Use IPMI to issue a supported power-cycling command to the IP address assigned to the Remote Power Cycle port, which you configured earlier, along with the required credentials.

For example, from a UNIX-type machine with IPMI support, you might issue the command:

```
ipmitool -I lan -H 192.0.2.1 -U remoteresetuser -P password chassis power reset
```

where **192.0.2.1** is the IP address assigned to the Remote Power Cycle port and **remoteresetuser** and **password** are the credentials that you entered while enabling this feature.

- Step 2** Wait at least eleven minutes for the email gateway to reboot.
-

Working with Technical Support

- [Technical Support for Virtual Email Gateway, on page 21](#)
- [Opening or Updating a Support Case From the Email Gateway, on page 22](#)
- [Enabling Remote Access for Cisco Technical Support Personnel, on page 22](#)
- [Running a Packet Capture, on page 25](#)

Technical Support for Virtual Email Gateway

Requirements for getting technical support for your virtual email gateway are described in the Cisco Content Security Virtual Appliance Installation Guide available from

<http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-installation-guides-list.html>

Opening or Updating a Support Case From the Email Gateway

Do not contact Cisco IronPort Customer Support for help with Cloud Email Security. See the *Cisco IronPort Cloud Email Security / Hybrid Email Security Overview Guide* for information on getting support for Cloud and Hybrid Email Security.

Before You Begin

- If your issue is urgent, do not use this method. Instead, contact support using one of the other methods listed in [Cisco Customer Support](#).
Use the following procedure only for issues such as a request for information or a problem for which you have a workaround, but would like an alternate solution.
- Consider other options for getting help:
 - [Knowledge Base](#)
 - [Cisco Support Community](#)
- To access Cisco technical support directly from the email gateway, your Cisco.com user ID must be associated with your service agreement contract for this email gateway. To view a list of service contracts that are currently associated with your Cisco.com profile, visit the Cisco.com Profile Manager at https://rpfa.cloudapps.cisco.com/rpfa/profile/profile_management.do. If you do not have a Cisco.com user ID, register to get one. See [Registering for a Cisco Account](#).
Be sure to save your Cisco.com user ID and support contract ID in a safe location.
- When you open a support case using this procedure, the email gateway configuration file is sent to Cisco Customer Support. If you do not want to send the email gateway configuration, you can contact Customer Support using a different method.
- In cluster configurations, support requests and their saved values are machine-specific.
- The email gateway must be connected to the internet and able to send email.
- If you are sending information about an existing case, make sure you have the case number.

Procedure

-
- Step 1** Sign in to the email gateway.
 - Step 2** Choose **Help and Support > Contact Technical Support**.
 - Step 3** Complete the form.
 - Step 4** Click **Send**.

Note CCO User IDs and the last-entered Contract ID are saved on the email gateway for future use.

Enabling Remote Access for Cisco Technical Support Personnel

Only Cisco Customer Assistance can access your email gateway using these methods.

- [Enabling Remote Access to Email Gateways With an Internet Connection](#) , on page 23
- [Enabling Remote Access to Email Gateways Without a Direct Internet Connection](#) , on page 23
- [Disabling Remote Access](#) , on page 24

- [Disabling a Tech Support Tunnel](#) , on page 24
- [Checking the Status of the Support Connection](#) , on page 24

Enabling Remote Access to Email Gateways With an Internet Connection

Support accesses the email gateway through an SSH tunnel that this procedure creates between the email gateway and the `upgrades.ironport.com` or server.

Before You Begin

Identify a port that can be reached from the internet. The default is port 25, which will work in most environments because the system also requires general access over that port in order to send email messages. Connections over this port are allowed in most firewall configurations.

Procedure

- Step 1** Log in to the email gateway.
- Step 2** From the top right side of the GUI window, choose **Help and Support > Remote Access**.
- Step 3** Click **Enable**.
- Step 4** Enter the following information:

Option	Description
Seed String	The seed string is used to generate a secure shared secret to be used by Cisco Customer Support to access this email gateway.
Secure Tunnel	Select the check box to use a secure tunnel for the remote access connection. Enter a port for the connection. The default is port 25 , which will work in most environments.

- Step 5** Click **Submit**.

What to do next

When remote access for support personnel is no longer required, see [Disabling a Tech Support Tunnel](#) , on page 24.

Enabling Remote Access to Email Gateways Without a Direct Internet Connection

For appliances without a direct internet connection, access is made through a second email gateway that is connected to the internet.

Before You Begin

- The email gateway must be able to connect on port 22 to a second email gateway that is connected to the internet.
- On the email gateway with the internet connection, follow the procedure in [Enabling Remote Access to Email Gateways With an Internet Connection](#) , on page 23 to create a support tunnel to that email gateway.

Procedure

- Step 1** From the command-line interface of the email gateway requiring support, enter the **techsupport** command.
- Step 2** Enter **sshaccess**.
- Step 3** Follow the prompts.
-

What to do next

When remote access for support personnel is no longer required, see the following:

- [Disabling Remote Access](#) , on page 24
- [Disabling a Tech Support Tunnel](#) , on page 24

Disabling a Tech Support Tunnel

An enabled techsupport tunnel remains connected to upgrades.ironport.com for 7 days. After that time, established connections will not be disconnected but will be unable to re-attach to the tunnel once disconnected.

To disable the tunnel manually:

Procedure

- Step 1** Log in to the email gateway.
- Step 2** From the top right side of the GUI window, choose **Help and Support > Remote Access**.
- Step 3** Click **Disable**.
-

Disabling Remote Access

A remote access account that you create using the techsupport command remains active until you deactivate it.

Procedure

- Step 1** From the command-line interface, enter the techsupport command.
- Step 2** Enter sshaccess .
- Step 3** Enter disable .
-

Checking the Status of the Support Connection

Procedure

- Step 1** From the command-line interface, enter the techsupport command.

Step 2 Enter status .

Running a Packet Capture

Packet Capture allows support personnel to see the TCP/IP data and other packets going into and out of the email gateway. This allows Support to debug the network setup and to discover what network traffic is reaching the email gateway or leaving the email gateway.

Procedure

Step 1 Choose **Help and Support > Packet Capture**.

Step 2 Specify packet capture settings:

- a) In the **Packet Capture Settings** section, click **Edit Settings**.
- b) (Optional) Enter duration, limits, and filters for the packet capture.

Your Support representative may give you guidance on these settings.

If you enter a capture duration without specifying a unit of time, AsyncOS uses seconds by default.

In the **Filters** section:

- Custom filters can use any syntax supported by the UNIX tcpdump command, such as host 10.10.10.10 && port 80 .
- The client IP is the IP address of the machine connecting to the email gateway, such as a mail client sending messages through the email gateway.
- The server IP is the IP address of the machine to which the email gateway is connecting, such as an Exchange server to which the email gateway is delivering messages.
- You can use the client and server IP addresses to track traffic between a specific client and a specific server, with the email gateway in the middle.

- c) Click **Submit**.

Step 3 Click **Start Capture**.

- Only one capture may be running at a time.
- When a packet capture is running, the Packet Capture page shows the status of the capture in progress by showing the current statistics, such as file size and time elapsed.
- The GUI only displays packet captures started in the GUI, not from the CLI. Similarly, the CLI only displays the status of a current packet capture run started in the CLI.
- The packet capture file is split into ten parts. If the file reaches the maximum size limit before the packet capture ends, the oldest part of the file is deleted (the data is discarded) and a new part starts with the current packet capture data. Only 1/10 of the packet capture file is discarded at a time.
- A running capture started in the GUI is preserved between sessions. (A running capture started in the CLI stops when the session ends.)

Step 4 Allow the capture to run for the specified duration, or, if you have let the capture run indefinitely, manually stop the capture by clicking **Stop Capture**.

Step 5 Access the packet capture file:

- Click the file in the **Manage Packet Capture Files** list and click **Download File**.
- Use FTP or SCP to access the file in the captures subdirectory on the email gateway.

What to do next

Make the file available to Support:

- If you allow remote access to your email gateway, technicians can access the packet capture files using FTP or SCP. See [Enabling Remote Access for Cisco Technical Support Personnel](#), on page 22.
- Email the file to Support.