SECURITY

Email Security with Cisco IronPort

The definitive guide to deploying and maintaining secure email architectures with Cisco IronPort ESA

ciscopress.com

Chris Porter

Email Security with Cisco IronPort

Chris Porter

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

Note In addition to public and private listeners, ESA also supports something called a blackhole listener. You can only create a blackhole listener in the CLI—the option isn't even displayed in the GUI. A blackhole does exactly what you'd guess: It accepts and then just dumps messages. There aren't many real-world scenarios that call for a blackhole listener, but they can be useful for testing as a sink for other systems or for the ESA to deliver back to itself. Just be careful; as with a real blackhole, anything that goes in is lost forever.

At this point, you have a running ESA, and because of the defaults in the HAT and incoming mail policies, you have a fairly secure setup that can actually be used right away, provided that you gave accurate information to **listenerconfig** prompts. There are a few other things that you normally want to configure before you do that, however. These tasks are normally taken care of by the SSW, but you need to do them manually:

1. Make sure DNS servers are specified using **dnsconfig**. ESAs need one or more DNS servers to resolve names and IPs, or need access on port 53/UDP and port 53/TCP to the Internet so that they can act as their own caching DNS resolver. The ESA will be extremely unhappy if DNS is not working: Connections will be slow to accept as DNS queries time out; SenderBase queries will fail, causing the system to accept potential junk connections; and email accepted by the system will likely not be deliverable. **nslookup** (and the hidden command **dig**) are available to test name resolution.
2. Set the time zone, date, and time, or set up NTP servers. Having accurate time and date is critical for email. **settz** allows you to pick a time zone or use a GMT offset. **settime** allows you to manually set both the date and time, but a better option is to use Network Time Protocol (NTP) synchronized to a local or public time server that you specify in **ntpconfig**. If you don't have one, you can use time.ironport.com.
3. Specify delivery destinations using **smtproutes**. Without a destination specified for a given domain, the ESA defaults to using DNS MX lookups to resolve delivery hosts. For your local domains, this is almost certainly wrong, and you need to add entries to use **smtproutes** to override this.
4. Set up alerts. Alerts are email messages generated by the ESA for error, warning, or informational events. **alertconfig** is where you establish email addresses that should receive these alerts, and each entry allows you to specify which categories and severities (or all) should be sent to whom. It's recommended to have at least one email address that will receive all categories and all severities. From there, if you have different roles for individuals on your team or across groups, you can set up categories appropriately.

Remember to **commit** all of your changes. Forgetting to **commit** changes is the most common cause of head scratching during troubleshooting. Running **commit** isn't harmful if there are no changes to commit, so just get into the habit of committing often.

Commands in Depth

We looked at a few CLI commands in this chapter, and there will be references to CLI commands throughout the rest of this book. Suffice to say, you need to know how to use them. Now, we examine the CLI commands in depth, organized by their functional area and demonstrated through some real-world examples.

The focus here is on information and configuration that can't be done in the WUI or on commands that have expanded capabilities compared with the WUI.

Troubleshooting Example

To better illustrate the use of CLI commands, we'll follow a typical mail delivery troubleshooting procedure. In this case, imagine that we have an unknown problem whose symptom is that outbound mail isn't being delivered to an external company. Users have complained to the IT helpdesk, and the helpdesk investigation found that the users' email clients and the company's groupware servers show that the message has been delivered. Because the groupware servers deliver all outgoing mail to the ESAs, they've escalated the problem to the messaging technicians, and that's you. There's no more information available. We're going to exclusively use the CLI to find the root cause of the delivery problem.

Status and Performance Commands

The WUI provides a quick-glance overview of the status of an ESA, but for critical, real-time information on status and performance, the CLI commands are ideal. The best place to start is the **status** command, but using the *status detail* option gives us all the same information, with additional performance and queue statistics. Example 5-5 shows the output of the **status detail** command.

Example 5-5 *Output of the* **status detail** *Command*

```
esa02.cisco.com> status detail

Status as of:                  Wed Aug 18 16:10:12 2010 EDT
Up since:                      Sat Aug 07 12:24:19 2010 EDT (11d 3h 45m 53s)
Last counter reset:            Never
System status:                 Online
Oldest Message:                15 mins 43 secs
Feature - Bounce Verification: Perpetual
Feature - Virus Outbreak Filters: 723 days
Feature - IronPort Email Encryption: 723 days
Feature - IronPort Anti-Spam:  723 days
Feature - Incoming Mail Handling: Perpetual
Feature - IronPort Intelligent Multi-Scan: 79 days
Feature - RSA Email Data Loss Prevention: 79 days
Feature - Sophos Anti-Virus:   723 days
```

```
Counters:                              Reset          Uptime        Lifetime
  Receiving
    Messages Received                143,126           3,932         143,126
    Recipients Received              154,609           3,939         154,609
    Gen. Bounce Recipients               267              19             267
  Rejection
    Rejected Recipients              506,679          14,359         506,679
    Dropped Messages                      40               0              40
  Queue
    Soft Bounced Events                    3               0               3
  Completion
    Completed Recipients             153,425           3,917         153,425
      Hard Bounced Recipients            524              38             524
        DNS Hard Bounces                 254              19             254
        5XX Hard Bounces                 259              19             259
        Expired Hard Bounces               4               0               4
        Filter Hard Bounces                7               0               7
        Other Hard Bounces                 0               0               0
      Delivered Recipients           109,741           2,630         109,741
      Deleted Recipients              43,160           1,249          43,160
        Global Unsub. Hits                 0               0               0
    DomainKeys Signed Msgs             60,719           2,529          60,719
  Current IDs
    Message ID (MID)                                                   154987
    Injection Conn. ID (ICID)                                          697826
    Delivery Conn. ID (DCID)                                            74870

Rates (Events Per Hour):            1-Minute       5-Minutes      15-Minutes
  Receiving
    Messages Received                   416             536             528
    Recipients Received                 416             536             528
  Queue
    Soft Bounced Events                   0               0               0
  Completion
    Completed Recipients                  0              72              32
      Hard Bounced Recipients             0               0               0
      Delivered Recipients                0              36              16

Gauges:                                Current
  System
    RAM Utilization                       4%
    CPU Utilization
      Total                               2%
```

```
    MGA                                2%
    Anti-Spam                          5%
    Anti-Virus                         0%
    Reporting                          0%
    Quarantine                         0%
  Disk I/O Utilization                 0%
  Resource Conservation                 0
  Logging Disk Usage                  14%
  Logging Disk Available              42G
Connections
  Current Inbound Conn.                 2
  Current Outbound Conn.                5
Queue
  Active Recipients                   435
    Unattempted Recipients            432
    Attempted Recipients                3
  Messages In Work Queue                0
  Messages In Quarantine            1,143
  Destinations In Memory                8
  Kilobytes Used                   29,877
    Kilobytes In Quarantine        29,877
  Kilobytes Free                8,358,731
```

From this output, we can begin troubleshooting. You can see at the top that the system status is online, which means that the ESA is actively accepting connections and delivering mail. The next line shows the oldest message age of more than 15 minutes, which is not necessarily cause for concern, but it's worth looking into.

The next lines of output start with the word Feature and refers to the active feature keys on the system. All of these show either perpetual or a non-zero duration, and so expired keys are not the source of the delivery problem. If any keys are expired or near expiration, the ESA generates email alerts to notify administrators, and you or someone in your group pays diligent attention to those kinds of alerts.

The next section of output shows combined statistics for messages and recipients accepted and completed. Data on messages is shown, because it's more straightforward for users and administrators to think in terms of messages. More pertinent information is shown based on recipients, because multirecipient messages require that each recipient be filtered, processed, and delivered separately. It's possible that a message with two recipients ends up delivered to one of the recipients and not to the other or delivered at a much later time. You must be familiar with thinking in terms of recipients, because most of the ESA troubleshooting tools work on a recipient basis, and most of the

filtering and delivery rules do, too. The counters here show three columns: Reset, which is since the last counter reset (using the **resetcounters** command); Uptime, which is since the last reboot or shutdown; and Lifetime, which is, of course, since the ESA came from the factory.

The Gauges section is of most interest to us in our delivery troubleshooting. The first gauges are for system performance: CPU load, broken out by component, Disk I/O, and use of physical disk space on the system. Because all the gauges in this report show very low usage, this is not the source of our delivery problem, and we'll move on.

Below the Gauges, we see Connections, organized into inbound and outbound connections. These are live values at the time when the command is run. In this context, inbound refers to the direction of the SMTP connection, not the direction of email. Inbound means initiated by any external IP in to the ESA, whereas outbound means initiated by the ESA to any external IP. We can have inbound connections from our internal groupware server, over which outgoing email is transmitted. An outbound connection could be from the ESA to local groupware, delivering incoming messages from an Internet sender to a local recipient. Only successful SMTP connections are shown here. Connections that are still being initiated, or which are rejected either by ESA or the remote server, are not added to the count.

In this troubleshooting example, we have both inbound and outbound SMTP connections, so we can conclude that the system is actively receiving and delivering email. The number of connections is important, but needs to be taken into context with how busy the system is. ESA automatically opens as many connections as it needs to deliver mail efficiently, subject to limits set on the system and subject to its *Good Neighbor* algorithm, which avoids flooding destination servers with massive numbers of connections.

The Queue section starts with Active Recipients, which is a count of all recipients on all messages that have completed filtering and are waiting to be delivered. You can think of it as the total count of all recipients in a delivery queue, although ESA does not actually use a disk or memory queue to store them. Recipients are actually organized in a tree structure, grouped by recipient domain. Barring a problem with a remote server, the ESA should either be actively attempting to deliver them or waiting to make another delivery attempt based on the retry schedule configured in the appropriate bounce profile. If there is a problem with any remote server, recipients are counted here until they are delivered or the message exceeds its maximum time in queue and is bounced.

Active Recipients is broken out into Attempted Recipients and Unattempted Recipients and, in this example, output most are unattempted. Unattempted Recipients are those whose destination hosts have never been reached, so the SMTP RCPT TO command has never been attempted. This can mean that the host is down and not responding or is refusing or rejecting connections. It can also mean that the ESA has the maximum number of connections open, due to system settings, and is waiting to open more. ESA continues delivery attempts by periodically trying to open SMTP connections to these hosts, but recipients are not marked as attempted until SMTP is successful. Attempted Recipients refer to recipients that have been presented to the receiving host using the SMTP RCPT TO command, but received a temporary (4yz) error code in response to

that recipient. This response could be caused by a busy server, an unavailable mailbox, or rate limiting by the remote host. Recipients in the Attempted category are reattempted until either the maximum number of retries or maximum time in queue has been exceeded. Getting to the SMTP RCPT TO command presumes that ESA was able to make a successful SMTP connection at some point.

Back to this example, we have five outbound SMTP connections, but 432 recipients in queue have never been attempted because the ESA can't make connections to the various remote servers. Although this isn't necessarily a severe problem all the time, it could certainly be the root of the delivery problem. The combination of a low number of outbound connections combined with a high number of active recipients may mean that we're experiencing delivery delays, so it is worth investigating.

Before we move on to the next step, we must examine the last lines of output. The next, Messages in Work Queue, is a count of those messages that are awaiting processing by the security engines. All the processing for AntiSpam, AntiVirus, Content Filtering, Outbreak Filters, Encryption, and DLP is performed in the work queue. Messages that will be scanned by one or more of these engines go into the work queue where they are handed off to an available security engine running in a separate process. Because the security engines are the performance bottleneck in most ESA configurations, having a non-zero value in the **Messages In Work Queue** field can be the first indication that an ESA has reached its maximum throughput, at least for the current instant. Because this value is zero here, this is not the source of our delivery delays.

If you have a value in this field, you can drill further into work queue statistics. For status on the work queue and specific engines, or for a view of performance, see the commands in Table 5-1.

Table 5-1 *Work Queue and Filtering Engine Commands*

Command	Description
workqueue status	Displays the status of the work queue (operational or paused) and the number of messages in the work queue.
workqueue rate	Shows a continual running count of messages entering and leaving the work queue, updated by the interval it prompts you for.
antispamstatus	Status of any enabled anti-spam engines, including date and time of last engine and rule update.
antivirusstatus	Status of any enabled antivirus engines, including date and time of last engine and rule update.
vofstatus	Virus Outbreak Filters, shows last rule update and all current active outbreak rules.
encryptionstatus	Shows date and time of last encryption engine update.
sbstatus	Shows the status of the SenderBase lookups, including the queries used to find reputation scores, but also for the queries submitted as part of SenderBase Network Participation (SBNP).