IP Reputation Filtering

This chapter contains the following sections:

Overview of Sender IP Reputation Filtering

Sender IP reputation filtering is the first layer of spam protection, allowing you to control the messages that come through the email gateway based on senders' trustworthiness as determined by the Sender IP Reputation Service.

The email gateway can accept messages from known or highly reputable senders — such as customers and partners — and deliver them directly to the end user without any content scanning. Messages from unknown or less reputable senders can be subjected to content scanning, such as anti-spam and anti-virus scanning, and you can also throttle the number of messages you are willing to accept from each sender. Email senders with the worst reputation can have their connections rejected or their messages bounced based on your preferences.

Note File reputation filtering is a separate service. For information, see File Reputation Filtering and File Analysis

IP Reputation Service

The IP Reputation Service, using global data from the Talos Affiliate network, assigns a IP Reputation Score (IPRS) to email senders based on complaint rates, message volume statistics, and data from public blocked lists and open proxy lists. The IP Reputation Score helps to differentiate legitimate senders from spam sources. You can determine the threshold for blocking messages from senders with low reputation scores.

The Talos Security Network website (https://talosintelligence.com) provides a global overview of the latest email and web-based threats, displays current email traffic volume by country, and allows you to look up reputation scores based on IP address, URI or Domain.

Note The IP Reputation Service is only available with a current anti-spam feature key.

Related Topics

- IP Reputation Score , on page 2
- How Sender IP Reputation Filters Work , on page 3
- Recommended Settings for Different Sender IP Reputation Filtering Approaches , on page 3
- Outbreak Filters
- Using Email Security Monitor

IP Reputation Score

The IP Reputation Score is a numeric value assigned to an IP address based on information from the IP Reputation Service. The IP Reputation Service aggregates data from over 25 public blocked lists and open proxy lists, and combines this data with global data from Talos to assign a score from -10.0 to +10.0, as follows:

Score	Meaning
-10.0	Most likely to be a source of spam
0	Neutral, or not enough information to make a recommendation
+10.0	Most likely to be a trustworthy sender

The lower (more negative) the score, the more likely that a message is spam. A score of -10.0 means that this message is "guaranteed" to be spam, while a score of 10.0 means that the message is "guaranteed" to be legitimate.

Using the IP Reputation Score, you configure the email gateway to apply mail flow policies to senders based on their trustworthiness. (You can also create message filters to specify "thresholds" for IP Reputation Scores to further act upon messages processed by the system. For more information, refer to "IP Reputation Rule" and "Bypass Anti-Spam System Action.")

Figure 1: The IP Reputation Service

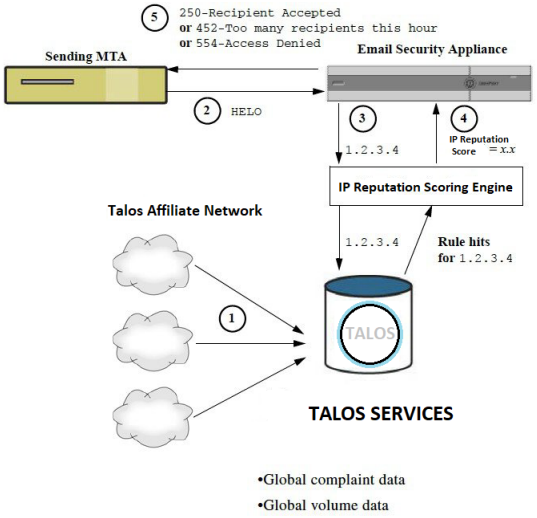

1. Talos affiliates send real-time, global data
2. Sending MTA opens connection with the email gateway
3. Email gateway checks global data for the connecting IP address
4. IP Reputation Service calculates the probability that this message is spam and assigns an IP Reputations Score
5. Cisco returns the response based on the IP Reputation Score

How Sender IP Reputation Filters Work

Sender IP Reputation filter technology aims to shunt as much mail as possible from the remaining security services processing that is available on the email gateway. (See Understanding the Email Pipeline.)

When sender reputation filtering is enabled, mail from known bad senders is simply refused. Known good mail from global 2000 companies is automatically routed around the spam filters, reducing the chance of false positives. Unknown, or "grey" email is routed to the anti-spam scanning engine. Using this approach, Sender IP Reputation filters can reduce the load on the content filters by as much as 50%.

Figure 2: Sender IP Reputation Filtering Example

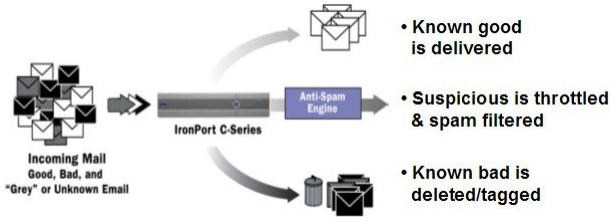

Recommended Settings for Different Sender IP Reputation Filtering Approaches

Depending on the objectives of your enterprise, you can implement a conservative, moderate, or aggressive approach.

Approach	Characteristics	Allowed _List	Blocked_ List	Suspectlist	Unknownlist
		Sender IP Reputation Score range:			
Conservative	Near zero false positives, better performance	7 to 10	-10 to -4	-4 to -2	-2 to 7
Moderate (Installation default)	Very few false positives, high performance	Sender IPReputation Scores are not used.	-10 to -3	-3 to -1	-1 to +10
Aggressive	Some false positives, maximum performance. This option shunts the most mail away from Anti-Spam processing.	4 to 10	-10 to -2	-2 to -1	-1 to 4
All approaches		**Mail Flow Policy:**			
		Trusted	Blocked	Throttled	Accepted

Editing IP Reputation Filtering Score Thresholds for a Listener

Use this procedure if you want to change the default IP Reputation Service score thresholds or add a sender group for reputation filtering.

Note Other settings related to IP Reputation Score thresholds, and Mail Flow Policy settings, are described in Defining Which Hosts Are Allowed to Connect Using the Host Access Table

Before You Begin

- If your email gateway is set to receive mail from a local MX/MTA, identify upstream hosts that may mask the sender's IP address. See Determining Sender IP Address In Deployments with Incoming Relays for more information.
- Understand IP Reputation Scores. See Defining Sender Groups by IP Reputation Score.
- Choose a filtering approach for your organization and note the recommended settings for that approach. See Recommended Settings for Different Sender IP Reputation Filtering Approaches , on page 3.

Procedure

Step 1 Select **Mail Policies** > **HAT Overview**.

Step 2 Select the public listener from the **Sender Groups (Listener)** menu.

Step 3 Click the link for a sender group.

For example, click the “SUSPECTLIST” link.

Step 4 Click **Edit Settings**.

Step 5 Enter the range of IP Reputation Scores for this sender group.

For example, for “ALLOWED_LIST,” enter the range 7.0 to 10.

Step 6 Click **Submit**.

Step 7 Repeat as needed for each sender group for this listener.

Step 8 Commit changes.

What to do next

Related Topics

- Testing IP Reputation Filtering Using the IP Reputation Scores, on page 5
- Defining Which Hosts Are Allowed to Connect Using the Host Access Table
- How to Configure the Email Gateway to Scan Messages for Spam

Testing IP Reputation Filtering Using the IP Reputation Scores

Unless you regularly receive a large portion of spam, or you have set up “dummy” accounts to specifically receive spam for your organization, it may be difficult to immediately test the IP Reputation policies you have implemented. However, if you add entries for reputation filtering with IP Reputation Scores into a listener’s HAT as indicated in the following table, you will notice that a smaller percentage of inbound mail will be “unclassified.”

Test the policies using the `trace` command with an arbitrary IP Reputation scores. See Debugging Mail Flow Using Test Messages: Trace. The `trace` command is available in the CLI as well as the GUI.

Table 1: Suggested Mail Flow Policies for Implementing the IP Reputation Scores

Policy Name	Primary Behavior (Access Rule)	Parameters	Value
`$BLOCKED`	`REJECT`	`None`	
`$THROTTLED`	`ACCEPT`	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: Use Spam Detection: Use TLS: Maximum recipients / hour: Use SenderBase:	10 20 1 MB 10 ON OFF **20** *(recommended)* ON

Policy Name	Primary Behavior (Access Rule)	Parameters	Value
`$ACCEPTED` (Public Listener)	`ACCEPT`	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: Use Spam Detection: Use TLS: Use SenderBase:	1,000 1,000 100 MB 1,000 ON OFF ON
`$TRUSTED`	`ACCEPT`	Maximum messages / session: Maximum recipients / message: Maximum message size: Maximum concurrent connections: Use Spam Detection: Use TLS: Maximum recipients / hour: Use SenderBase:	1,000 1,000 100 MB 1,000 OFF OFF -1 *(disabled)* OFF

Note In the $THROTTLED policy, the maximum recipients per hour from the remote host is set to 20 recipients per hour, by default. Note that this setting controls the maximum throttling available. You can increase the number of recipients to receive per hour if this parameter is too aggressive. For more information on Default Host Access policies, see Understanding Predefined Sender Groups and Mail Flow Policies.

Entering Low IP Reputation Scores in the Message Subject

Although Cisco recommends throttling, an alternate way to use the IP Reputation Service is to modify the subject line of suspected spam messages. To do this, use the message filter shown in the following table. This filter uses the reputation filter rule and the strip-header and insert-header filter actions to replace the subject line of messages having a IP Reputation Score lower than -2.0 with a subject line that includes the actual IP Reputation Score represented as: **{Spam** *IP Reputation Score***}**. Replace *listener_name* in this example with the name of your public listener. (The period on its own line is included so that you can cut and paste this text directly into the command line interface of the filters command.)

Table: Message Filter to Modify Subject Header with IP Reputation**: Example 1**

```
iprs_filter:

if ((recv-inj == "listener_name
" AND subject != "\\{Spam -?[0-9.]+\\}"))
```

```
{

        insert-header("X-IPRS", "$REPUTATION");

        if (reputation <= -2.0)

{

        strip-header("Subject");

        insert-header("Subject", "$Subject \\{Spam $REPUTATION\\}");

       }

 }

.
```

Related Topic

- Using Message Filters to Enforce Email Policies