



Implementing Antispam technologies in Forefront Protection 2010 for Exchange Server

Introduction	2
Critical Components of Forefront Protection 2010 for Exchange Server Antispam Defense	2
Source Analysis Layer.....	3
IP Allow and IP Block Lists	4
Forefront DNS Block List	5
Sender ID Framework	7
Source Analysis Layer summary.....	8
Source Analysis Layer Benefits.....	9
Protocol Analysis Layer	9
SMTP Tarptitting	10
Limiting Submission Rate	11
Safelists Enforcements.....	13
Sender Filtering.....	14
Recipient Filtering	16
Backscatter Filtering.....	18
Protocol Analysis Layer Summary.....	22
Protocol Analysis Layer Benefits.....	22
Content Analysis Layer.....	22
Content Filter Engine Scanning	25
FPE Content Filter Analysis Layer Summary:	26
FPE Content Filter benefits:	26

Introduction

The purpose of this whitepaper is to demonstrate the business value FPE provides for protecting Exchange Server organizations, the technological superiority of antispam technologies FPE enables for its customers, and the critical important FPE adds with its advanced features for shielding Exchange organizations against unsolicited and malicious e-mail while increasing the deliverability of legitimate correspondences. This paper summarizes the new e-mail hygiene protection features available in the Forefront Protection for Exchange Server 2010 (FPE 2010) release, illustrates functionality available for protecting Exchange Server deployments at the FPE 2010 RTM time, and outlines integration with Exchange server features contributing to the Forefront antispam solution.

Implementing FPE 2010 for Exchange Server as the primary defense system against spam adds significant business and technical value, enabling organizations to function more efficiently by removing security threats associated with unwanted and malicious content entering organizations via e-mail. It also results in increasing the end users satisfaction and overall productivity. In the end this evidently translates into monetary savings allowing FPE-protected Exchange organizations to save on Capital Expenditures associated with deployment of new hardware layers and Operational Expenditures associated with messaging hygiene management and helpdesk operations. FPE for Exchange enables integrated, technology-agnostic, multi-faceted, layered antispam solution that is straightforward to implement, cost-effective to operate, and inexpensive to support. And most importantly, FPE is natively integrated with Exchange Server 2010 technologies, backwards-compatible with Exchange 2007 server, supports end-to-end antispam framework across the family of Microsoft messaging products, and delivers solution designed for multi-directional malware scanning and protection.

Most modern antispam products are composed of several technologies tightly coupled together for detecting spam; however, FPE for Exchange is the only solution that enables a multi-layered, multi-directional, integrated, and distributed model for antispam filtering.

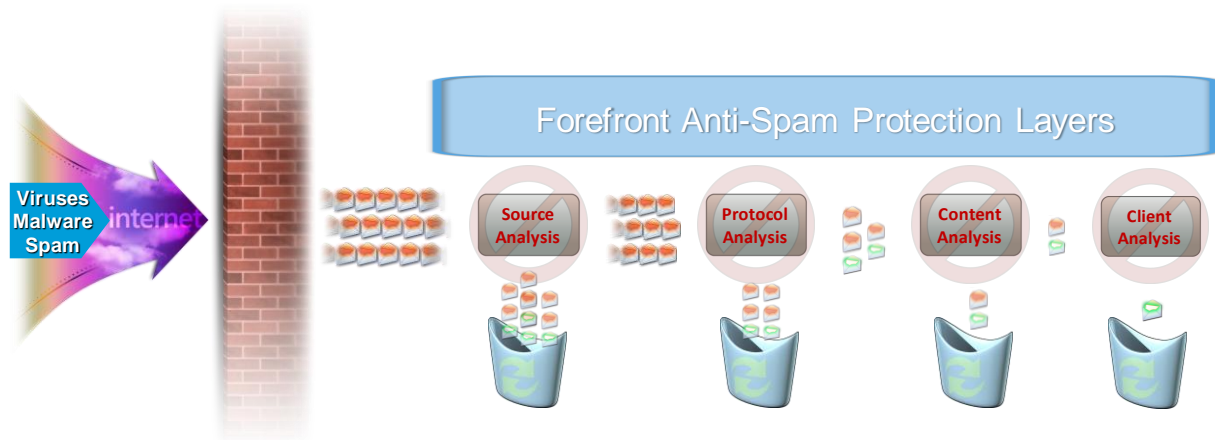
This whitepaper will concentrate on elaborating the Forefront technologies that contribute to multi-layered, multi-directional, and integrated approach in combating spam. Distributed model which is referenced as Hybrid Model in the Forefront antispam framework has been covered in a separate whitepaper which entirely concentrates on outlining the integration between Forefront on-premises and Forefront Online and the benefits acquired from implementing such architecture in Forefront Protection 2010 for Exchange Server deployments.

Critical Components of Forefront Protection 2010 for Exchange Server Antispam Defense

The Forefront Protection 2010 for Exchange Server (FPE) antispam solution is composed of four major spam detection layers:

1. Source Analysis

2. Protocol Analysis
3. Content Analysis
4. Client Analysis



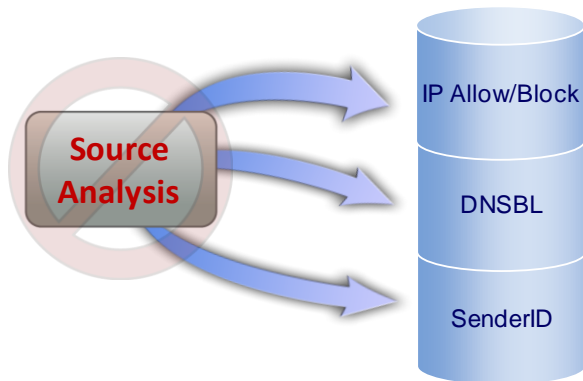
These layers are tightly integrated into FPE pipeline but are loosely coupled and can be deployed independently of each other. However, enabling all of them together allows for the most efficient spam scanning and produces the best results in protecting organizations from spam and malicious content. While Client Analysis, enforced on the Outlook client end (for example, blocking certain encoding lists or top-level domains), is an integral part of the end-to-end antis spam protection, it will not be covered in this whitepaper.

Source Analysis Layer

Ideally, to reduce the impact on an organization's hardware resources, network bandwidth, and the end users, unsolicited e-mail should be stopped at the organization's network perimeter. Forefront Protection 2010 for Exchange Server enables multiple technologies to evaluate the connecting party's identity, reputation, and the fidelity of the SMTP transaction. Some of these technologies, such as the Forefront DNS Block List are the result of collaborative work among multiple internal Microsoft teams and external vendors to provide simple but effective solutions for Forefront customers to effectively combat spam.

The Source Analysis Layer consists of 3 major filtering technologies:

1. IP Allow/IP Block Lists
2. DNSBL (DNS Block List also known as RBL)
3. SenderID Framework

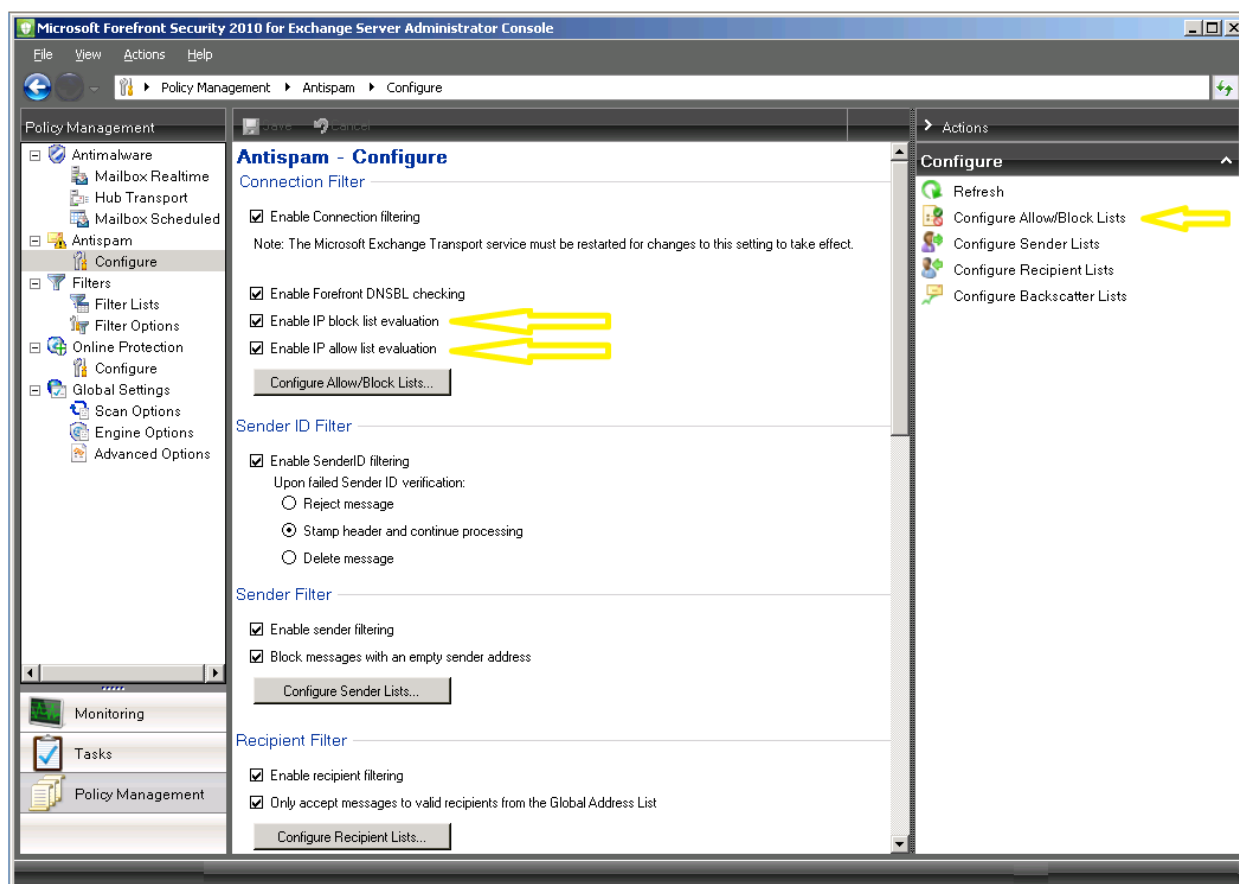


It is important to understand that while being considered as Source, or Connection Level, filtering these features execute at different points during an SMTP transaction. However, all of them are based on evaluating the reputation of the connecting IP addresses.

IP Allow and IP Block Lists.

IP Allow and IP Block lists are the first agents to execute in the Source Analysis Layer. While IP Connection Filtering is considered somewhat rudimentary technology because it requires manually maintaining a list of remote MTA hosts making connections to FPE-protected Exchange servers, it is still a very effective technology that allows Forefront administrators to quickly provide remedies if they are under attack from a particular IP Address. By entering a list of IP Addresses into the IP Block List, TCP connection requests originating from these IPs to SMTP port 25 are denied at the time of connection. The IP Block List agent is triggered on the connection event when the remote host tries to establish a connection with the FPE-protected Exchange Server.

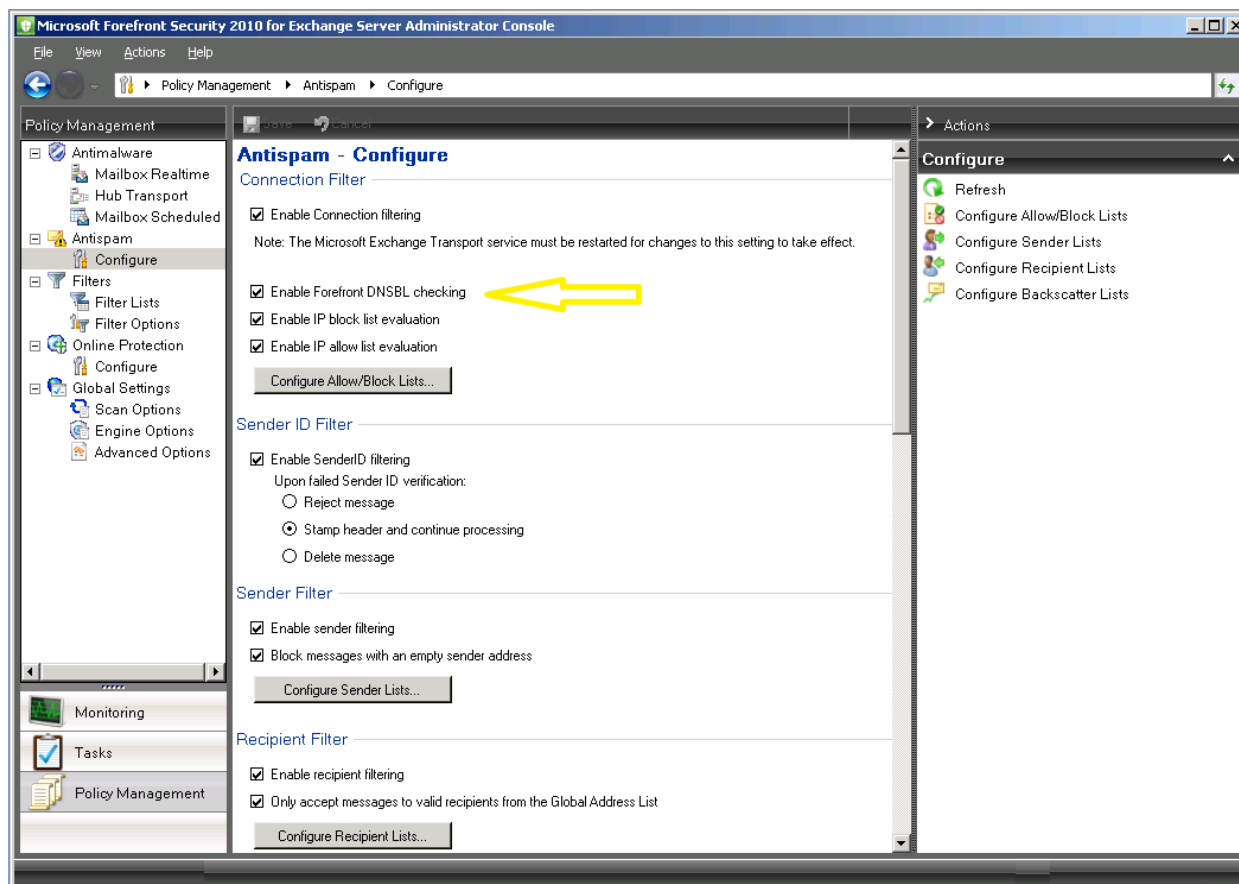
Similarly, entering a list of IPs onto the IP Allow List will allow all SMTP transactions originating from those IPs to bypass Forefront antispam filtering layers and increase the deliverability of legitimate mail. Both IP Allow and IP Block lists configurable in the FPE administrative console as shown below:



When adding an IP address to either the Allow or Block list, the Forefront administrator can select an expiration date when the selected action should be terminated. It is also possible to import the list from another FPE-protected server or export the list to another FPE-protected Exchange server. This functionality is very important to help administrators in scenarios when Forefront protects Exchange organizations that are not managed via the Forefront Protection Manager.

Forefront DNS Block List

Forefront DNS block list (DNSBL) is a feature that provides dynamic blocking of incoming connection requests based on the connecting IP address, which is checked against a Microsoft-maintained aggregated DNS blocklist that includes information collected from various third-party block lists. If the IP Address is on the list, SMTP transactions from this IP will be blocked. While the technology itself is hardly new, it is one of the most effective means to counter spam attacks. To protect the list from unauthorized queries, the Forefront DNSBL agent sends the query in an encrypted format. The management of the Forefront DNSBL is easy – all that needs to be done to enable the list is to select the check box next to the **Enable Forefront DNSBL checking** option in the FPE administrator console as shown below:



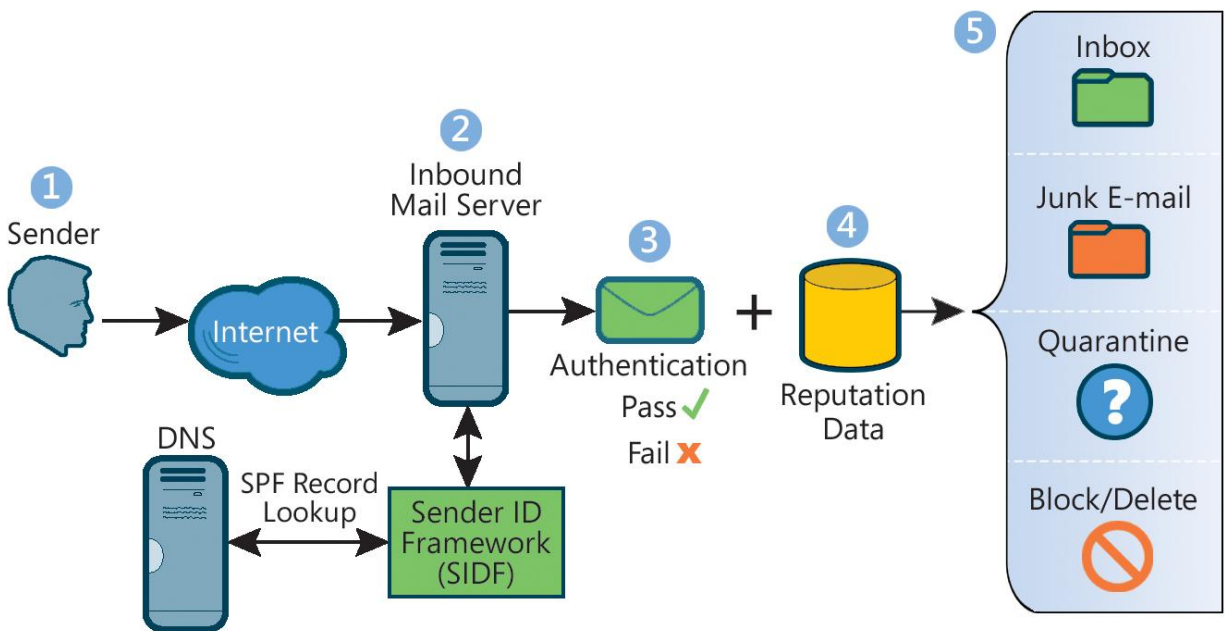
There are no additional actions needed from the administrator to enable the feature, everything happens under the Forefront hood. So what exactly happens under it? When an incoming connection is checked, the DNS filter consults the DNSBL to see if the IP address of the connecting server is on the list. To prevent unauthorized access to the DNSBL data, it is hashed on the DNSBL server end. The same algorithm has been applied on the FPE by the DNSBL agent to hash queries before sending them to the DNSBL provider service for verification. Once they reach the DNSBL service provider, these queries are evaluated for their fidelity. If the hash does not compute correctly, the DNSBL provider will not service the query. If the hash computes properly, the DNSBL query will be appropriately serviced. If the connecting IP address has been listed on one of the vendor's block lists, the returned query will contain the exact feed where it is listed and the proper instructions how to implement corrective actions. For example, if the connection request was blocked by IP address being on the Exchange Hosted Services blocklist, the Forefront DNSBL agent will issue the following response: "550 5.7.1. Mail Submission Rejected by {*blocklist_provider_name*}. Mail From IP Banned. To request removal from the list please forward this message to delist.forefront@messaging.microsoft.com".

If the IP was blocklisted mistakenly or maliciously, forwarding the message to the specified alias will successfully start the process of delisting IP from the blocklist.

Sender ID Framework

Sender ID Framework is the cornerstone of the e-mail authentication technologies integrated by Forefront Server Protection. SenderID Framework provides a way for FPE-protected Exchange servers to authenticate an incoming e-mail message and validate its source. SenderID Filter verifies that the incoming message originates from the IP that can legitimately send on behalf of the domain referenced in the sender's address. Successfully implemented SenderID Framework not only ensures brand name protection for businesses and legitimate senders but enables an additional layer of defense against spoofed messages that contain malicious payloads, especially during zero-day attacks.

Zero-day attacks exploit the lag in time between the discovery of the security vulnerability and antivirus vendors supplying new antivirus definitions. In many cases these exploits are being carried via messages with spoofed sender domains. In such scenarios SenderID Filtering is the most effective and in many cases the last and only line of defense in such scenarios. SenderID Framework is relatively easy to implement. The image below, taken from [Sender ID SPF Records Wizard](#) maintained by Microsoft, outlines the validation logic executed in the framework. The Wizard itself enables an administrator to create correct SPF records in a just a few clicks.

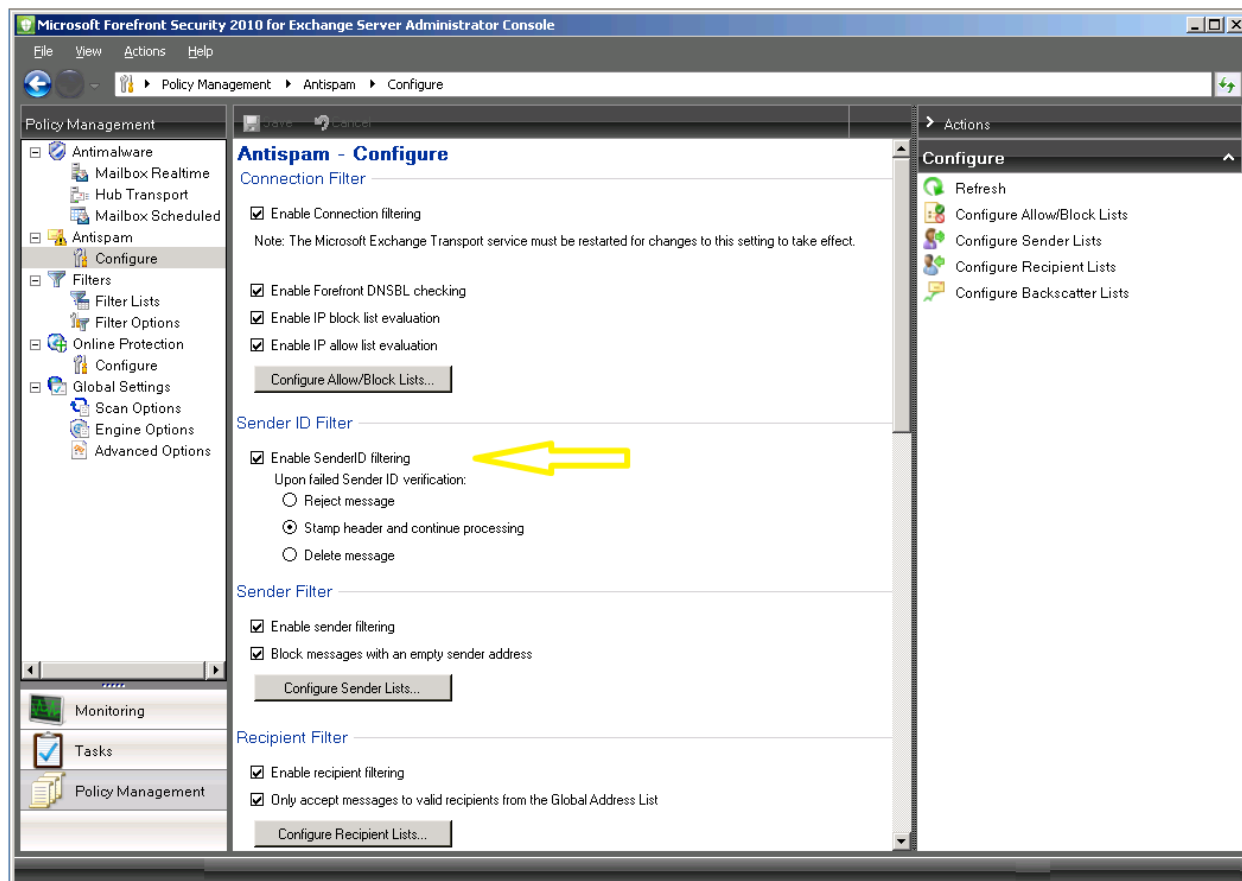


Having SenderID properly configured and functioning will greatly increase deliverability of legitimate mail and help to achieve strong protection against spoofing attacks. FPE exposes SenderID protection via the administrative console, which contains the following options (as shown below):

Enable SenderID Filtering:

- Reject message
- Delete message

- Stamp and continue processing



It is important to understand that the SenderID Filter exposed by FPE will Reject or Delete the message only if it is a clear (or identified) spoof. Both the Reject and Delete actions correspond to the Hard Fail verdict of SenderID Framework. If the verdict is Soft Fail or any other, the filter will stamp an appropriate header on the message and let it through without rejecting or deleting. To achieve the best results using SenderID Framework, FPE-protected Exchange organizations are encouraged to create SPF records with the following format: “-all”. Having “-all” in the SPF records allows SenderID to make reliable ruling about the integrity of SMTP transaction and trustworthiness of connecting IP sending on behalf of specified domain.

Source Analysis Layer summary

- The Source Analysis layer is based on:
 - IP address evaluations
 - Domain/IP resolutions and originating traffic patterns validation
 - Reputation discovery via feeds from partners and third-party providers
- Reputation and e-mail authentication are major contributing factors in ensuring the fidelity of connecting IPs “driver’s license” and protecting domain names/brands

- Key technology for circumventing distributed spam attacks and limiting overall impact to FPE-protected Exchange deployments during the connection event.

Source Analysis Layer Benefits

Source Analysis filtering is one of the most efficient layers in the entire antispam solution and provides the following benefits:

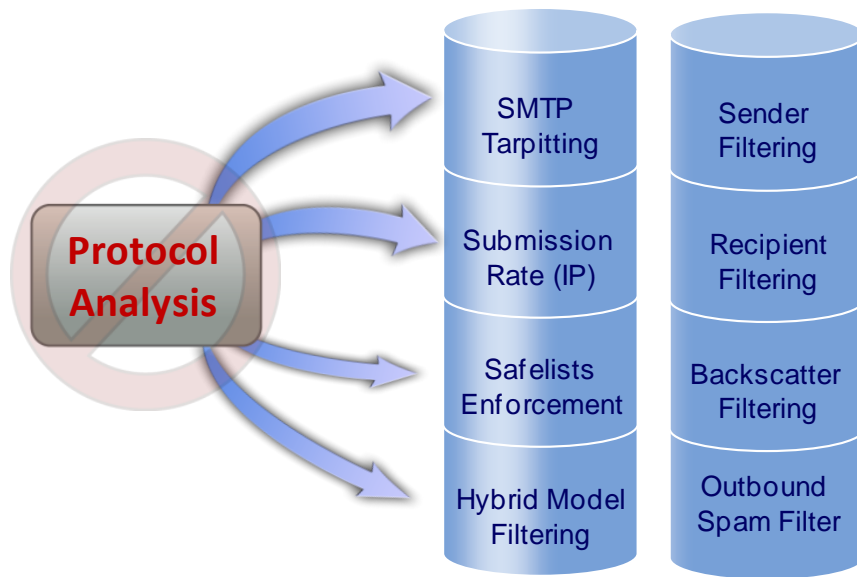
1. Unwanted or malicious e-mail is immediately rejected at the connection time
2. Early rejection means reduction in the carbon footprint of spam and no bandwidth/resource consumption as unwanted content won't get moved through the network layers
3. Safelisted IPs receive preferential treatment which translates into increased deliverability of legitimate mail
4. Circumvention of distributed spam attacks.

Protocol Analysis Layer

The Protocol Analysis Layer combines protection available from both the core components of the transport pipeline and technologies exposed by Forefront Server Protection. It is important to understand that the best results in antispam defense can be achieved by amalgamation of these technologies into a single solution.

The Protocol Analysis Layer includes the following protection technologies:

1. SMTP Tarpitting (via SMTP stack)
2. Message submission rate (per IP)
3. Sender Filtering
4. Recipient Filtering
5. Backscatter Filtering
6. Outbound Spam Filtering
7. Safelists Enforcement
8. Hybrid Model Filtering



SMTP Tarpitting

SMTP Tarpitting is one of the most useful features to protect Forefront organizations from Directory Harvesting Attacks (DHAs). If the Recipients Validation feature is enabled, the SMTP protocol stack will perform recipient verification against Active Directory and if the recipient does not exist it will reject the mail transaction for the specified recipient as shown below:

Client: MAIL FROM: sender@contoso.com

Server: 250 2.1.0 sender@domain.com....Sender OK

Client: RCPT TO: alex@fabrikam.com

Server: 250 2.1.5 alex@fabrikam.com ← This is how acceptance takes place

Client: RCPT TO: bogusrecipient@fabrikam.com

Server: 550 5.5.1 User Unknown ← This is how rejection takes place

Malicious users can take advantage of this very useful feature, execute a dictionary attack, and perform directory enumeration in a very short amount of time by parsing the server responses. To counter the attack at the SMTP protocol level, the receiving server can slow down its responses back to the client (this is called tarpitting) making the Directory Harvesting Attack impossible. By default, the tarpit interval on the FPE-protected Exchange servers is 5 seconds, meaning after every incorrectly issued command (in our example to a non-existing recipient) the server will issue a response after a five second wait. This interval is modifiable via *Set-ReceiveConnector* cmdlet using *TarpitInterval* parameter with the desired number of seconds. For example, to change the default tarpit interval to 10 seconds:

```
[PS] C:\ Set-ReceiveConnector -Identity <ReceiveConnectorIdParameter>
```

```
- TarpitInterval 00:00:10
```

Setting the tarpit interval to 00:00:00 effectively disables the tarpit feature.

If an attacker pipelines a number of RCPT TO: commands to non-existent recipients, server response to every command will get delayed by default for five seconds making enumeration of directory via dictionary attack a non-viable option for the attackers. To compliment the tarpitting feature, it is advisable to enforce strict compliance of connecting hosts with governing SMTP RFCs. Commercially implemented MTAs rarely violate SMTP protocol, so it is beneficial to limit a number of errors on a single connection from a particular IP address. The default number is five. To change the default to two:

```
[PS] C:\ Set-ReceiveConnector -Identity <ReceiveConnectorIdParameter>
```

```
- MaxProtocolErrors 2
```

Limiting Submission Rate

By default Exchange servers in the Edge role are set to accept 600 messages from a single IP per minute. However, in some cases, for example if under spam attack, it is practical to limit the submission rate to a lower number. Limiting the number of mail submissions will help to implement an immediate response to an unfolding spam attack. To implement the Submission Rate limit, the following Exchange cmdlet should be used:

```
[PS] C:\ Set-ReceiveConnector
```

Accommodating parameters:

```
- Identity <ReceiveConnectorIdParameter>
```

```
- MessageRateLimit <integer>
```

To obtain the value for the Identity parameter a Forefront administrator will need to run the following cmdlet:

```
Get-ReceiveConnector | fl
```

The output of the cmdlet will have the ID parameter of the receive connector and it is the value the Forefront administrator needs. The output below shows the needed Identity value:

```

Select Machine: ffnv-spam01 | Scope: FFNY-SPAM01D.testdomain.local
[PS] C:\Documents and Settings\AutoAdmin>Get-ReceiveConnector f1

AuthMechanism           : Tls, Integrated, BasicAuth, BasicAuth
                          RequireTls, ExchangeServer
Banner                  :
BinaryMimeEnabled       : True
Bindings                : <0000:0000:0000:0000:0000:0000:0.0.0.
                          0:25, 0.0.0.0:25>
ChunkingEnabled         : True
DefaultDomain           :
DeliveryStatusNotificat : True
ionEnabled              :
EightBitMimeEnabled     : True
DomainSecureEnabled     : False
EnhancedStatusCodesEnabled : True
LongAddressesEnabled    : False
OxarEnabled             : False
Rqdn                   : FFNY-SPAM01
Comment                :
Enabled                 : True
ConnectionTimeout       : 00:10:00
ConnectionInactivityTimeout : 00:05:00
MessageRateLimit        : unlimited
MaxInboundConnection    : 5000
MaxInboundConnectionPerSource : unlimited
MaxInboundConnectionPercentagePerSource : 100
MaxHeaderSize           : 64KB
MaxHopCount             : 30
MaxLocalHopCount        : 8
MaxLogonFailures        : 3
MaxMessageSize           : 200MB
MaxProtocolErrors       : 5
MaxRecipientsPerMessage : 5000
PermissionGroups         : AnonymousUsers, ExchangeUsers, ExchangeServers, ExchangeLegacyServers, Custom
PipeLiningEnabled       : True
ProtocolLoggingLevel    : None
RemoteIPRanges          : <0000:0000:0000:0000:0000:0000:0.0.0.
                          0-ffff:ffff:ffff:ffff:ffff:ffff:255.2
                          55.255.255, 0.0.0.0-255.255.255.255>
RequireEHLODomain       : False
RequireTls              : False
EnableAuthGSSAPI        : False
Server                  : FFNY-SPAM01
SizeEnabled             : EnabledWithoutValue
TarpitInterval          : 00:00:05
AdminDisplayName        :
ExchangeVersion         : 0.1 (0.0.535.0)
Name                    : Default FFNY-SPAM01
DistinguishedName       : CN=Default FFNY-SPAM01,CN=SMTP Receive
                          Connectors,CN=Protocols,CN=FFNY-SPA
                          M01,CN=Servers,CN=Exchange Administra
                          tive Group (PYDIBOHF23SPDLT),CN=Admin
                          istrative Groups,CN=First Organizatio
                          n,CN=Microsoft Exchange,CN=Services,C
                          N=Configuration,DC=FFNY-SPAM01D,DC=te
                          stdomain,DC=local
Identity                : FFNY-SPAM01\Default FFNY-SPAM01
Guid                    : 036727e6-7697-41a0-aa39-ce02508d1992
ObjectCategory          : FFNY-SPAM01D.testdomain.local/Configu
                          ration/Schema/ms-Exch-Smtp-Receive-Co
                          nconnector
ObjectClass              : {top, msExchSmtpReceiveConnector}
WhenChanged             : 1/21/2009 2:50:33 PM
WhenCreated             : 1/21/2009 2:40:05 PM
OriginatingServer       : ffnv-spam01.FFNY-SPAM01D.testdomain.l

```

The Identity parameter in this particular case has the following value:

“FFNY-SPAM01\Default FFNY-SPAM01”

After obtaining corresponding values to the cmdlet parameters, running the following task will force the SMTP stack to reject messages sent in excess of ten from a single client IP address:

```
[PS] C:\ Set-ReceiveConnector -Identity “FFNY-SPAM01\Default FFNY-SPAM01”
```

```
-MessageRateLimit 10
```

If a client IP tries to initiate additional SMTP transactions after submitting ten messages, a FPE-protected MTA will refuse such transaction requests. Although not exposed via the FPE administrator console, this feature contributes to Forefront end-to-end antispam framework in an important way by throttling down message volume per sender (remote MTA).

Other important Exchange options that contribute to incoming network traffic shaping are:

MaxInboundConnectionsPerSource: This option will throttle the number of simultaneous connections originated from the same IP address assignment (default is 100). To modify the default:

```
[PS] C:\ Set-ReceiveConnector -Identity <ReceiveConnectorIdParameter>
```

```
- MaxInboundConnectionsPerSource "Number"
```

MaxInboundConnection: This option defines the total number of connections accepted at the same time (default is 5000). To modify the default:

```
[PS] C:\ Set-ReceiveConnector -Identity <ReceiveConnectorIdParameter>
```

```
- MaxInboundConnection "Number"
```

MaxInboundConnectionPercentagePerSource: This option defines how many connections a single IP can establish with the receive connector (default value is 2%). To modify the default:

```
[PS] C:\ Set-ReceiveConnector -Identity <ReceiveConnectorIdParameter>
```

```
- MaxConnectionPercentagePerSource "Number"
```

MaxRecipientsPerMessage: This option defines the maximum number of recipients allowed per message (default number is 200). To modify the default:

```
[PS] C:\ Set-ReceiveConnector -Identity <ReceiveConnectorIdParameter>
```

```
- MaxRecipientsPerMessage: Number
```

MaxMessageSize: This option specifies the maximum size of a message accepted into FPE-protected Exchange organization (default size is 10MB). To modify the default:

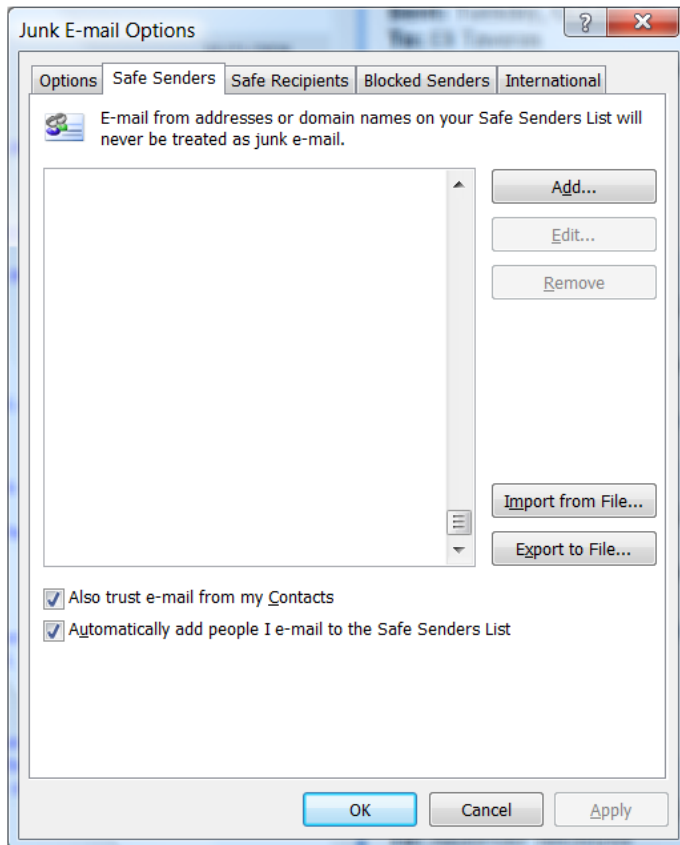
```
[PS] C:\ Set-ReceiveConnector -Identity <ReceiveConnectorIdParameter>
```

```
- MaxMessageSize: Number (for example :5MB)
```

These options, when combined together with Forefront filtering features, will help to secure strong protection against malicious attempts exploiting various attack vectors against FPE-protected Exchange organizations.

Safelists Enforcements

Safe lists proved to be extremely valuable additions to the FPE antispam framework. In some regard, these lists are the way for the end user to define the antispam behavior of FPE-protected Exchange servers. These lists have are exposed by the outlook client under the Junk E-Mail Options as shown below:

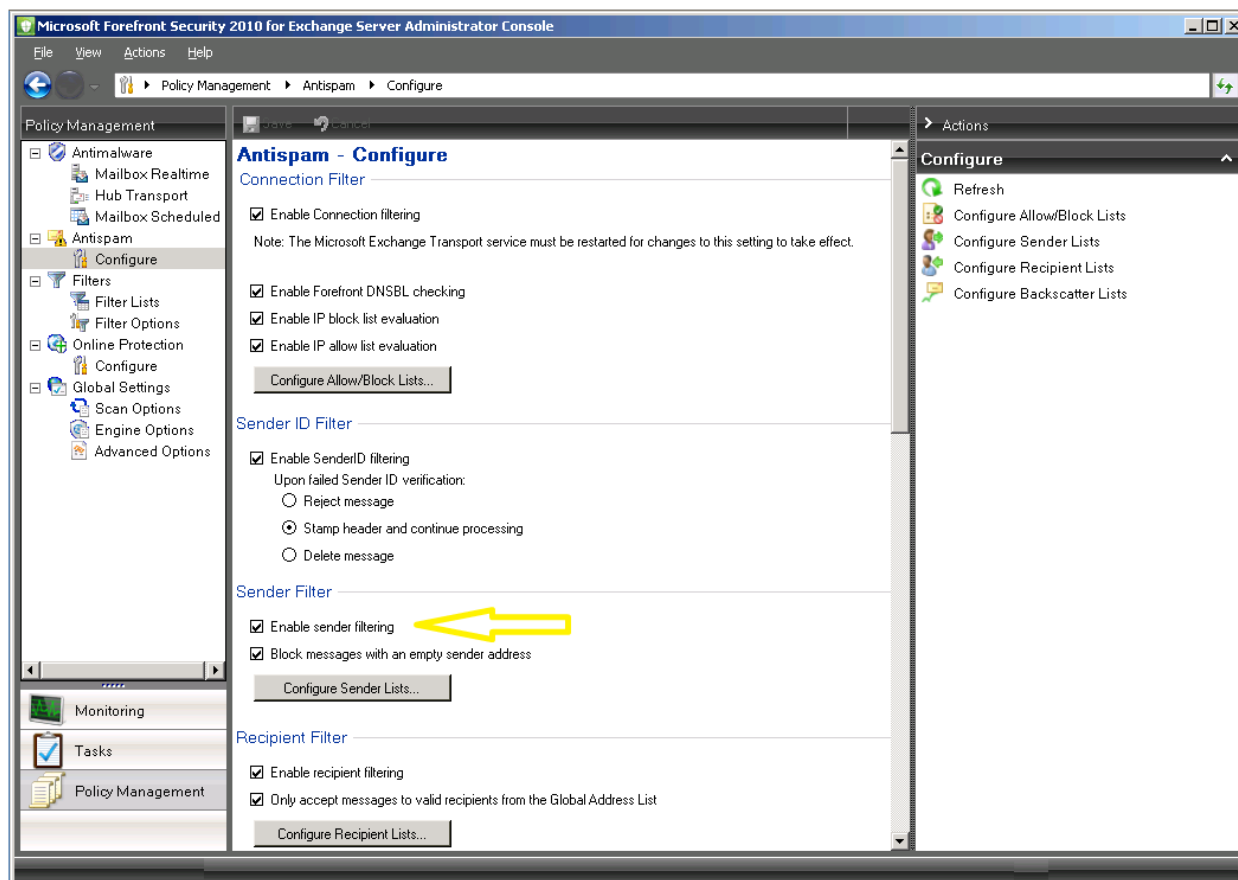


There are three main lists – Safe Senders, Safe Recipients, and Blocked Senders lists. Enforcement of Safe Senders, Safe Recipients, and Blocked Senders happens inside the Forefront antispam framework (enforcement of Blocked Senders in Exchange 2007 server happens on the client side). All the entries from the above Lists have been propagated by the Mailbox Assistant to Active Directory and then from there by EdgeSync to the network perimeter to all FPE antispam agents. During the antispam scanning of incoming mail, FPE antispam agents will aggregate per-recipient Safe Senders, Safe Recipients, and Blocked Senders information and act accordingly. Blocked Senders functionality happens inside the Sender Filter and is not available in Exchange 2007 server. If the recipient of a message item included the sender in Safe Senders List, the first FPE agent will remove the message from the antispam processing, stamp it with SCL:-1 value, and hand it off for antivirus scanning. This greatly improves the accuracy of spam processing, expedites e-mail delivery, and allows e-mail to arrive in the end users inbox richly rendered.

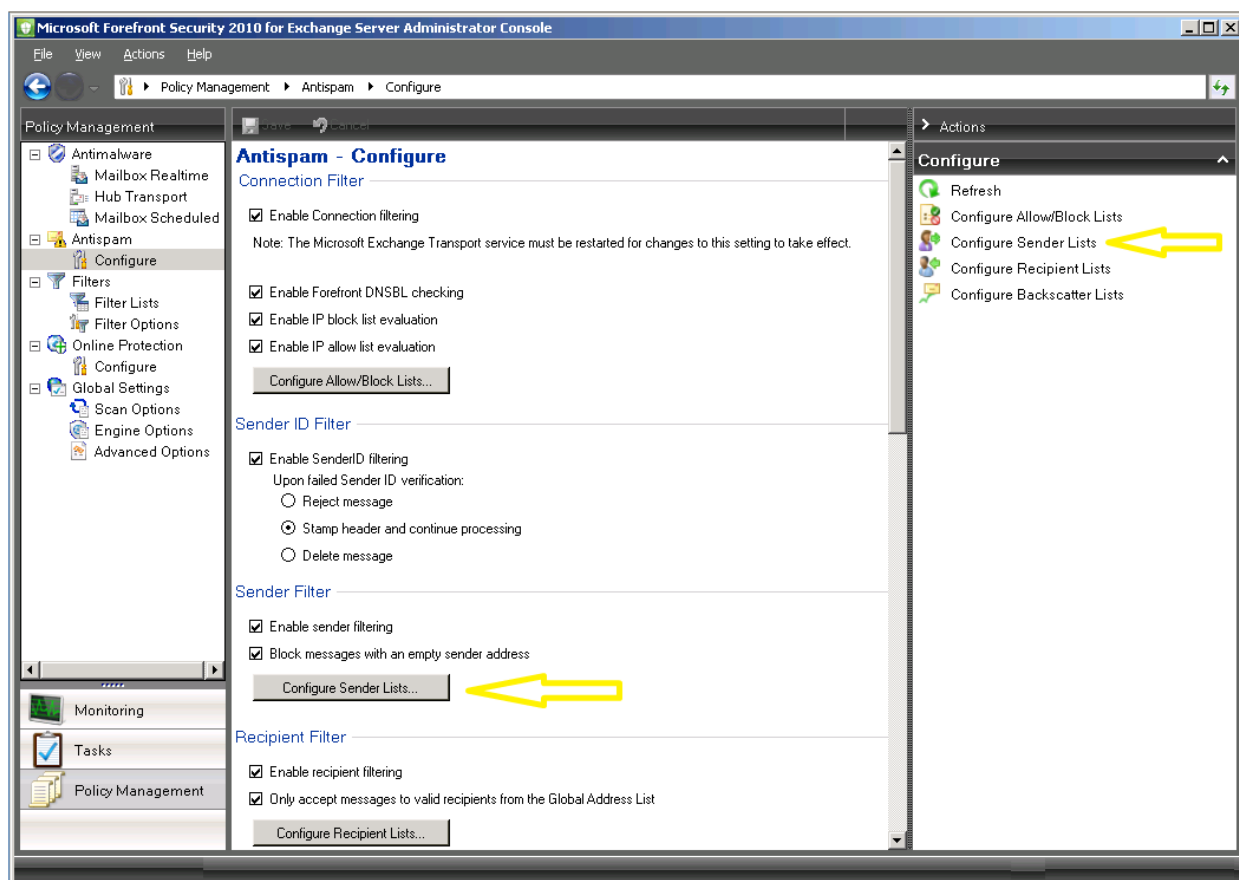
Sender Filtering

Sender Filtering is the first SMTP filtering agent that evaluates incoming messages. It aggregates the list of senders and domains blocked from submissions to FPE-protected Exchange organizations and a list of senders and domains that should be excluded from antispam scanning. During the SMTP transaction, the agent is triggered on the MAIL FROM: event and will reject the command if the sender or sender's domain exists on the Sender or Sender Domain block lists.

For example, if a FPE administrator adds the address “spammer@bogus.net” to the Sender Block List as shown below, all SMTP transactions from this sender are rejected by the Sender Filtering agent. Similarly, if an administrator enters the domain “Contoso.com” to the Sender Domain Block list, all messages initiated from that domain are rejected by the Sender Filtering agent. The last available option FPE administrator can enforce on the Sender Filter is to block messages with an empty Sender address. One of the known spammers’ tactics involves trying to confuse Exchange servers by sending messages with an empty Sender address. The FPE Sender Filtering agent will protect from this attack vector by monitoring the integrity of SMTP conversations and enforcing compliance with governing SMTP RFCs. If a malicious user tries to send a message without a valid Sender’s address on RFC2822, the FPE Sender Filtering agent will reject it. The RFC compliance is enforced for both RFC2821 and RFC2822.



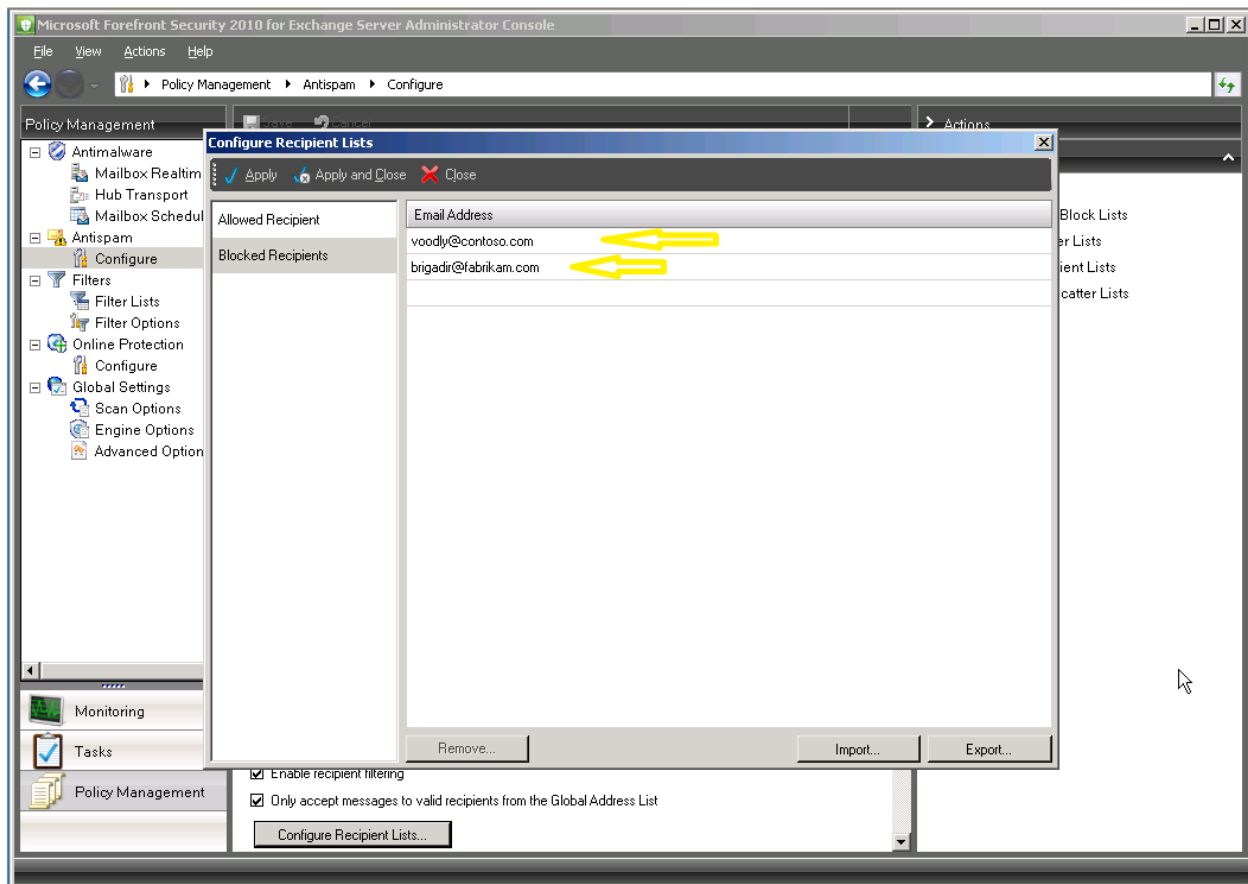
Additionally, the Sender Filtering agent will evaluate the Global Exceptions List for explicitly safe-listed senders and sender domains. If a match between the sender and an entry on the Global Exceptions List is found, the mail item is immediately given an SCL:-1 stamp from the agent and is extracted from the antispam processing pipeline. The mail item is then handed off to the antimalware scanning layer. The image below shows the location in the FPE for Exchange administrator console where FPE administrators can specify the Global Exceptions List.



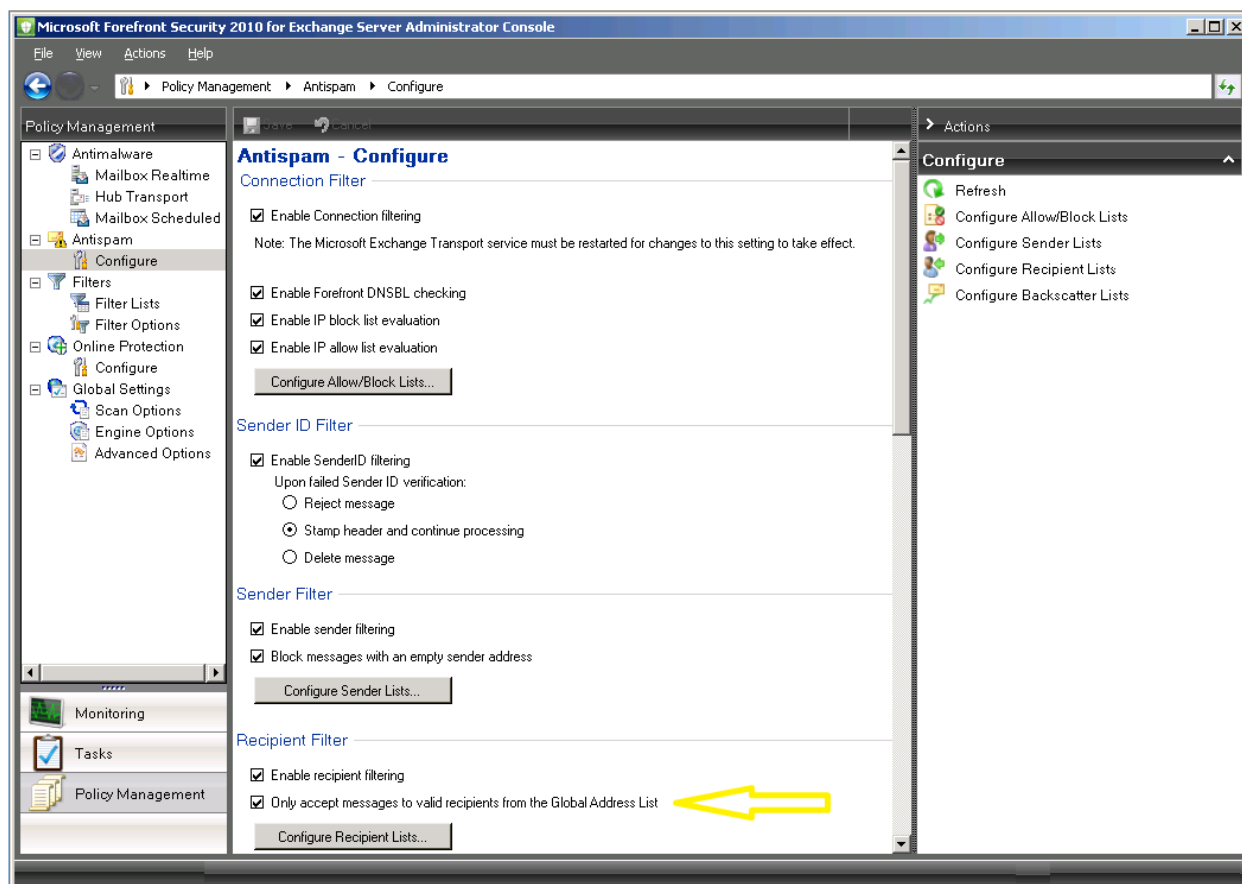
It is important to understand that regardless of the list positioning, any list consuming protection technology will be able to access it and take advantage of it.

Recipient Filtering

Recipient Filtering happens after Forefront's Sender Filtering and enforces administrator-defined actions on the recipients of messages. The Recipient Block list, same as Sender Block lists, is Global. In this regard global means it will be applied across the entire FPE-protected Exchange organization regardless of individual recipients settings. The list is static and requires manual population. In the example shown below, a FPE administrator enters "voodly@contoso.com" and "brigadir@fabrikam.com" onto the Recipient Block list. From that time on all e-mail coming from the Internet to these addresses will be rejected. Internally circulated e-mail won't be affected by this as this functionality applies only to the inbound mail stream (e-mail coming from the Internet).



Another important piece of functionality that FPE Recipient Filtering exposes is the ability to reject e-mail messages for non-existent recipients. During an SMTP session, if a submitting party issues the RCPT TO: command to a recipient that does not exist, the FPE agent will initiate Active Directory lookup to verify whether the recipient exists in the Exchange organization, and if it does not exist, the agent will reject the transaction for that recipient. The rest of the recipients won't be affected by this and mail will be delivered to them properly. Having this option enabled allows for the early rejection of non-existent recipients and saves network bandwidth and computing resources on generating DSNs. The image below shows where the Recipient Filtering functionality needs to be enabled and managed in the FPE administrator console.



Backscatter Filtering

Backscatter, also known as bogus NDRs or collateral spam, is a side effect of spamming attacks carried out with a spoofed sender address. The forged SMTP RFC2821 MAIL FROM: address points to a legitimate sender and in the event of delivery failure the receiving MTA will send a bounce to the unsuspecting victim referenced on the spoofed P1 Mail From: address. Backscatter represents not only a spam or annoyance issue but a serious security problem as it can carry malevolent payload and easily trick an unsuspecting recipient into opening it. FPE for Exchange is pioneering the backscatter protection for its customers in the FPE 2010 release of the product. Protection against backscatter will be exposed via an Anti-Backscatter filter that will tag all outbound (leaving the Exchange organization) mail with a BATV token. This agent is implemented as a Routing Agent.

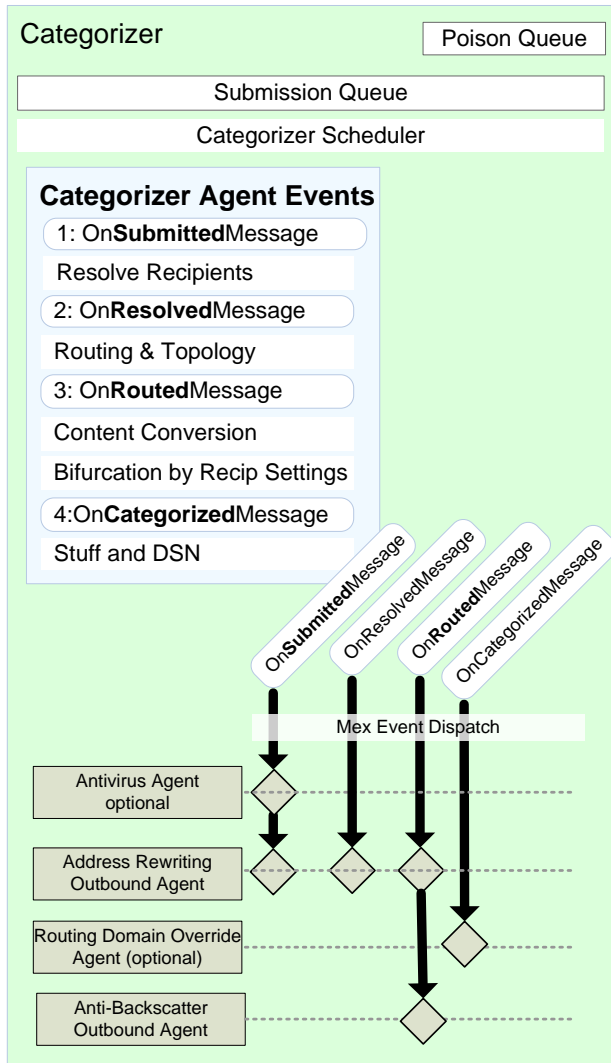
Here is how the backscatter filter works. It has been implemented as two agents:

- An Outbound agent that works inside the CAT (Categorizer)
- An Inbound agent that works inside the SMTP Receive pipeline of Exchange server.

On the outbound side, the agent stamps outgoing messages with a token (it will add it to the P1.MailFrom address) and on the inbound side, it verifies if the DSN has the token attached and whether the token computes correctly. Every outbound token contains information about the original sender

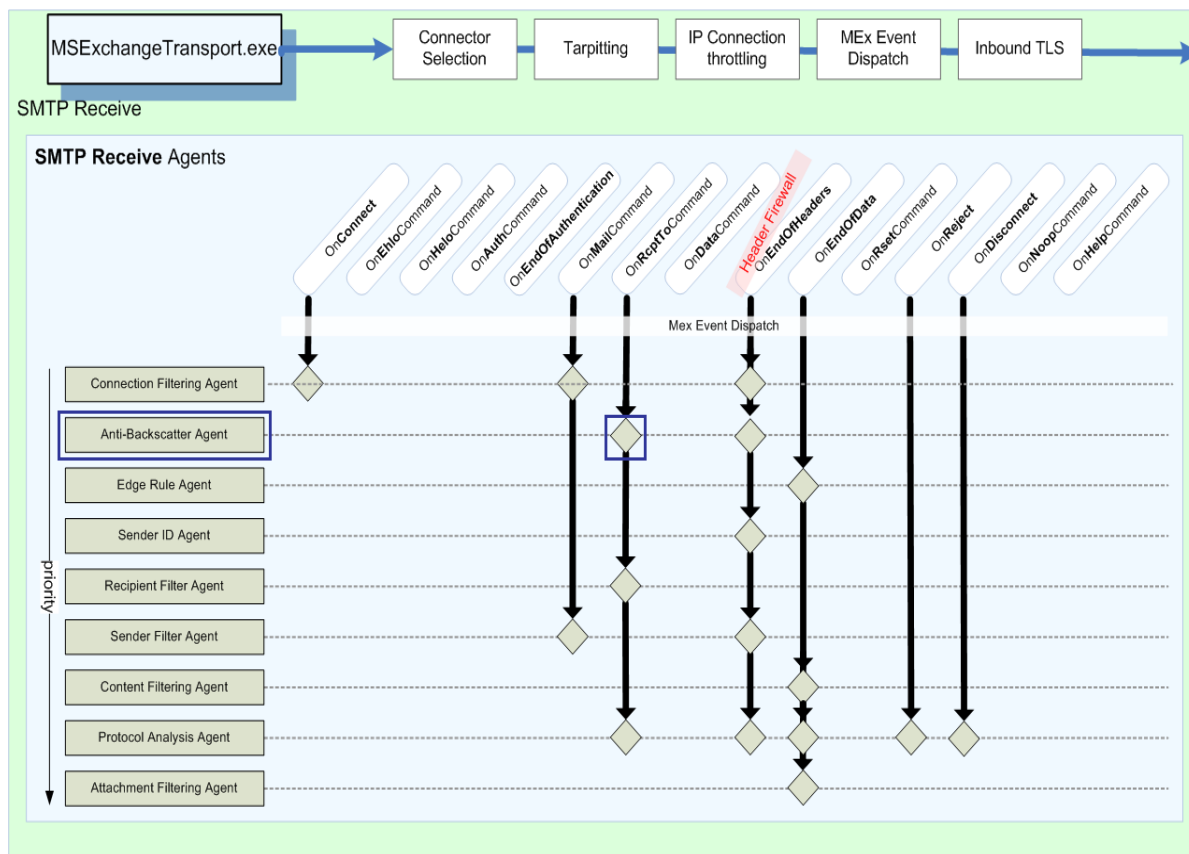
(sender's address), the key used to compute the token, and the time interval when the token is valid. This means that even if the spammers get the token, they won't be able to re-use it because the inbound agent will verify the token for integrity and if it does not compute correctly, will reject the e-mail transaction.

The diagram below shows the agent's execution place in the outbound Transport pipeline:



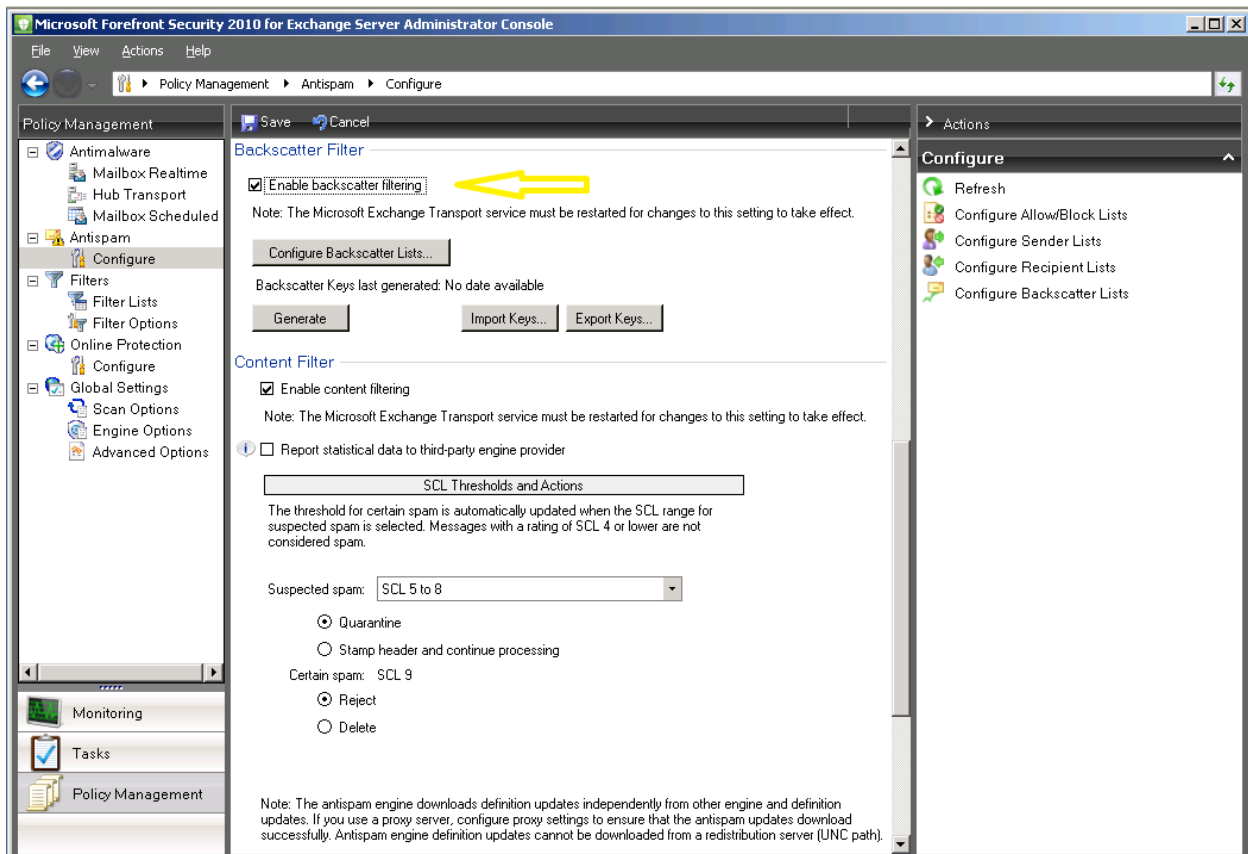
The anti-backscatter outbound agent, implemented as a routing agent in the Categorizer, runs on the OnRoutedMessage event. The agent will compute a token using a special key, the sender's address, the time of the event, and attach the token to the SMTP RFC2821 MAIL FROM: address.

On the inbound mail stream, backscatter filtering will be performed via the anti-backscatter agent implemented as an SMTP Receive agent. The diagram below shows how backscatter filtering works on the inbound Transport pipeline:



The anti-backscatter SMTP Receive agent is triggered on the RCPT TO: command (RFC2821) and resolves whether the incoming mail is a DSN, and if it is a DSN, it evaluates the fidelity of the attached token. If a remote MTA sends a bounce back to the Exchange recipient, the agent will verify the token's existence and its integrity. If the token is missing or does not compute correctly to the proper values, the DSN is rejected. If the BATV token computes correctly, the agent removes the BATV part from the address and hands the mail item off to the next antispam agent in the FPE pipeline.

The image below shows how Backscatter functionality has been exposed in the FPE administrator console:



The filter has been implemented based on BATV technology in a very simple, flexible, and secure way. In order to use the filter, an administrator needs to:

1. Enable the filter
2. Create a set of keys
3. Transfer the keys to all servers that participate in sending/receiving inbound/outbound mail (with the release of FPE Protection Manager this will be taken care of automatically)

The key to successful backscatter protection are the keys generated for the filter to use. While there is no need to re-generate the keys on a daily/weekly or even monthly basis, if the key set gets compromised, an administrator will be able to regenerate the set (only one new set of keys is allowed during a 24 hour period). An administrator can also specify an exemption list to exempt certain domains from being backscatter tagged/scanned. If an administrator, for example, adds "contoso.com" to the "Excluded domains" list, all mail sent to this domain won't be tagged with BATV tokens on the outbound and on the inbound DSNs will bypass the backscatter verification agent as well.

Protocol Analysis Layer Summary

Protocol Analysis Layer is based on:

- E-mail volume throttling per sending identity and receiving MTA
- Evaluation of fidelity and integrity of SMTP transaction
- Sender and Recipient filtering
- End user settings aggregation by the Forefront Protection 2010 for Exchange Server

Protocol Analysis Layer Benefits

Among the most visible benefits the Protocol Analysis Layer provides are:

- Ability to accept clean e-mail transactions bound for legitimate recipients only including evolving spam and DDoS attacks via DSNS
- Enforced user-defined filtering translates into richly rendered client messaging experience
- Great flexibility in selecting the right combination of individual filtering features complimented by the core transport technologies to harden the antispam defense at the SMTP session level
- Backscatter spam protection allows for early detection and rejection of incoming NDRs spam

Content Analysis Layer

After incoming e-mail has been scanned by the Source and Protocol Analysis filtering technologies, it will be transferred to the next layer – Content Filtering. This is the first time a third-party vendor's antispam engine has been delivered as a part of a Microsoft on-premise solution to protect Exchange server. Not surprisingly, FPE integrates one of the most effective content scanning engines available today, the Cloudmark Authority Engine.

The engine has been incorporated into the FPE pipeline and co-exists with other antispam features. The engine is a highly efficient content processor and compliments the end-to-end FPE antispam solution, making it a stellar performer. The engine continues the antispam processing after the SenderID filter and produces the final verdict on the value of the message. The SCL ratings assigned by the engine are also backwards-compatible, so they are understood by previous releases of Exchange server and Outlook client. All custom applications, developed by ISVs around SCL utilization, will continue to work without the need to re-factor third-party products to recognize the output of the Cloudmark Authority Engine.

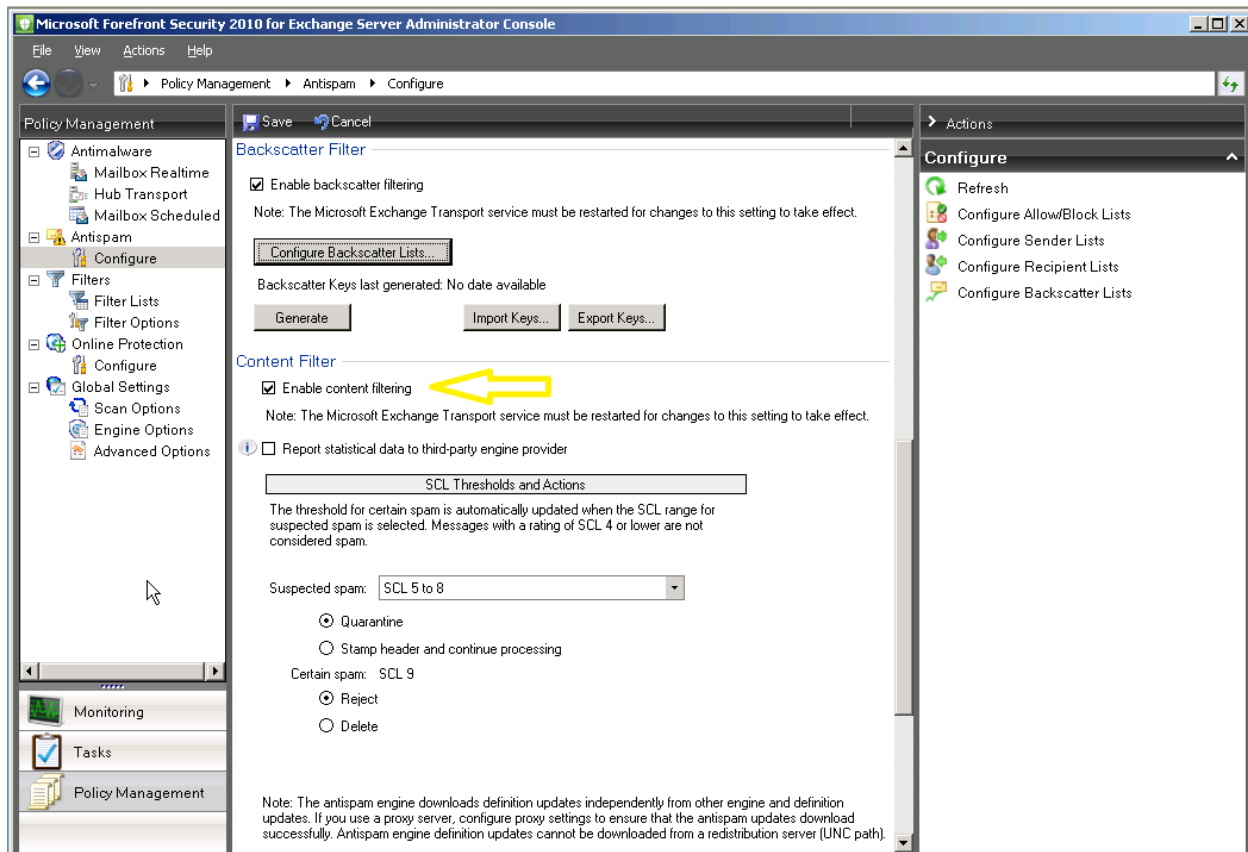
There are, however, a few major differences in the way FPE handles content filtering compared to the previous content scanning engines delivered as a part of Exchange 2007 server. First, and the most important, is the way the Cloudmark engine distributes the spam score weights. It tends to be very polarized to both lower and upper ends with very small score distribution in the mid-range. This translates into vast majority of SCLs having weights of either SCL:-1 for legitimate correspondence or SCL:8 and above for spam. Less than 10% of messages will receive an SCL of -1 (legitimate good mail) and more than 90% will receive an SCL of 9 (clear spam). This distribution logic covers almost 100% of all score assignments.

As you can see, the number of SCL scores distributed along the mid-range (specifically between SCL:5 to SCL:8) is very small, the total is less than a fraction of 1%. At a first glance, for administrators who are accustomed to more even statistical score distribution across the entire SCL range in Exchange 2007, this might look unusual. However, the key benefit from the new FPE SCL normalization logic is the fact that the amount of mail Quarantined for both the end user's and administrator's review has been significantly reduced to a degree that combing through the Junk E-mail folder trying to find and rescue False Positives is almost unnecessary. Simply put, the end users won't see any junk even in their Junk E-mail folders! Accordingly, Exchange administrators will be able to manage the Quarantine database much more easily and efficiently as the anticipated volume of mail stored in the database will be at most very modest. For example, based on internal Microsoft IT numbers, the number of quarantined items is less than 50 per 1 million incoming mails. This immediately translates into increased end-user productivity and unblocks Exchange administrators to allocate their time to other important projects. So it is normal for the end users to see empty Junk E-mail folder, for administrators to see only occasional mail in the Quarantine database, and have the bulk of SCL scores distribution to SCL:-1 and SCL:9.

Another difference in handling the content scanning is that FPE enables statistical data feedback to third-party providers allowing Cloudmark Authority Engine to receive more accurate and precise fingerprints via micro updates. These updates happen automatically so there is no need to explicitly enable content filter updates. Again, less clicks – better productivity! And probably the most important part for the updates is their frequency – the micro updates happen approximately every 45 seconds, so the fingerprints stored in local cache are always fresh and have the latest spam campaign characteristics to allow FPE content filter to react to the latest spam attacks almost instantaneously.

The feedback loop integrated into the Cloudmark Authority engine spans all continents and over 145 countries so the international coverage of content scanning enabled by FPE is now much stronger and reliable. Of course, due to local regulations and law challenges in certain regions of the world, the feedback loop from some countries is better than from others and FPE and Cloudmark are working relentlessly to enable superior, unremitting end-user contributions to the engine from across the globe.

Now, let's see what's under FPE's Content Filtering hood. The Content Filtering layer has been exposed via the FPE administrative console as shown below:



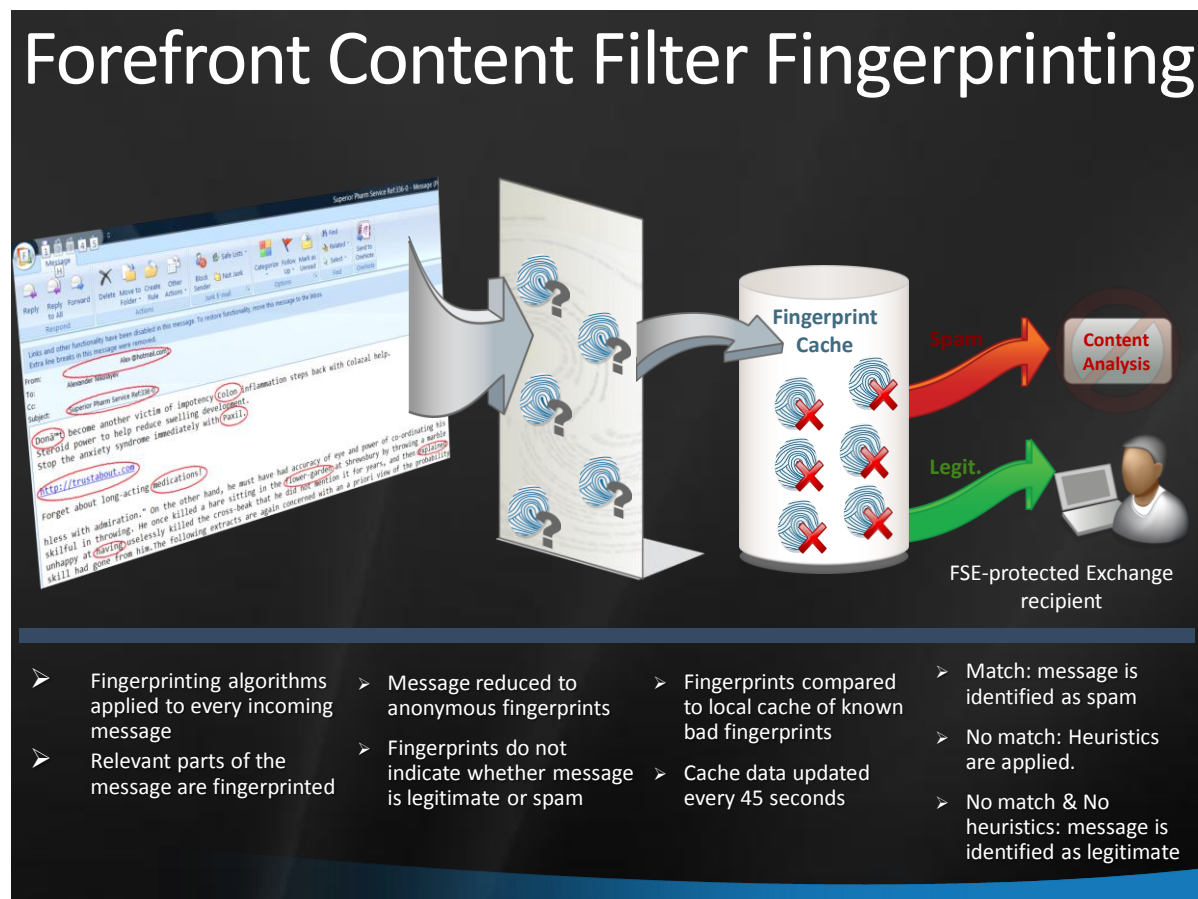
As you can see it's quite a departure from the previous Content Filter UI in Exchange 2007 server. There are very few options to configure so you are only a click (two clicks at most) away from enabling the strongest content filtering engine released by Microsoft to date.

The "Suspected spam" dropdown box enables administrator to achieve the desired balance of filtering – what should go into the Quarantine and what should be rejected or deleted. Rejection happens inside of the SMTP transaction so FPE Content Filter provides a reason in the SMTP reject response to enable corrective actions on the sending end if needed. Deletion of spam within the Certain Spam range is silent, no indication will be sent back to the sender. For the Suspect Spam management, an administrator can select all suspect spam (there are very few mails in this range) to be deposited into the Quarantine database or release it to the end users. In this case, all Suspect Spam mail will be deposited into the Junk E-mail folder so the end users can triage deposited spam themselves.

Frankly, there is little virtue in dumping spam into end users' Junk E-mail folder, however, local administrators need to make that decision themselves – whether to triage spam in a centralized Quarantine database or to make possible for end users to triage spam. In any case, the amount of spam entering a FPE-protected Exchange organization with the SCLs in range SCL:5 to SCL:8 will be very insignificant and will hardly result in any productivity loss, and we can position the decision of spam management as an end user satisfaction issue. If the end users prefer to manage spam themselves, the administrators will be able to redirect messages stamped with appropriate SCLs into end users Junk E-mail folders.

Content Filter Engine Scanning

Now, let's see how the actual content scanning is performed by the FPE content filter. When a message arrives, the FPE Content Filter Agent will first shred the MIME stream to allow extraction of relevant actionable parts by the Cloudmark Authority engine. The Cloudmark engine will apply fingerprints to the extracted content, aggregate them as anonymous fingerprints, and compare them to the cache of bad fingerprints maintained locally. This cache is continuously updated every 45 seconds to ensure the most accurate content fingerprinting and resolution. If a match for fingerprints derived from the message is found in the local cache, the message is identified as spam. If there is no match found, the engine will apply heuristics. After this, if there is no match found in both the fingerprints and the heuristics, the message is identified as legitimate. The Cloudmark Authority Engine will provide a raw spam score to the FPE content filter agent which in turn will normalize the raw spam score to SCL, stamp the message with the appropriate SCL header, and deliver the message to the next security layer in the FPE pipeline. The diagram below outlines the main steps and logic in FPE content filter processing.



Very often customers ask the following question – What is the preferred, or recommended, action for the content filter for Certain Spam – is it Reject OR Delete? While the FPE team can't enforce any recommendations, the most logical option would be to Reject the message inside of the SMTP transaction. This way, in a rare potential case of False Positive, the sender will receive a Non-Delivery Receipt notifying him that mail delivery failed because the message was identified as spam. This allows

the sender to take corrective action to resolve the problem. If Delete is used, the message is accepted by the Content Filter and then silently deleted without issuing any Delivery Status Notification back to the sender. Nevertheless, under certain circumstances an administrator can make a decision to Delete rather Reject. This will prevent any accidental information disclosure back to the spam sender.

The FPE Content Filter natively integrates anti-phishing protection; however, the Content Filter Agent won't stamp an anti-phishing header on the message as phishing protection is incorporated into the fingerprinting process and is a part of content scanning via the Cloudmark Authority engine.

FPE Content Filter Analysis Layer Summary:

- Based on Cloudmark Authority Engine with industry-leading performance metrics
- Embedded into the FPE pipeline via agents framework
- Scans MIME stream only – body + headers of the message
- Fingerprints-based engine that incorporates anti-phishing protection
- Enables feedback loop for better engine accuracy
- Seamlessly integrates into end-to-end antispam filtering

FPE Content Filter benefits:

- Reduced spam and phishing penetration
- Enhanced server performance
- Increased IT Pro and IW productivity
- Simplifies administration and management
- Improved end user satisfaction

In conclusion, it is important to understand that every action taken by the FPE antispam filtering agents is recorded in the FPE Agent Log and is retrievable via the appropriate PowerShell cmdlet.

FPE Protection 2010 for Exchange server integrates industry-leading content scanning engines and Microsoft internal technologies to benefit Exchange organizations by providing powerful and robust malware protection that includes state of the art antispam filtering.