MxVault

Questions and Answers

GLOBAL MICRO
Intelligent Technology

Table of Contents

MxVault

Questions and Answers

GLOBAL MICRO
Intelligent Technology

Incoming Service

1. Incoming Filtering Service

KNOWLEDGEBASE:

"Bounce spam" can be an annoying problem. The email SMTP protocol is a very simple protocol that was defined in 1982. Spam was not yet a problem and to keep things as simple as possible, no security measures were implemented in the protocol itself. The result of this is that there is no verification whatsoever that the "From:" address in an email message actually belongs to the sender.

To try and avoid spamfilters, spammers will typically use random email addresses as fake senders. This way they can avoid any simple spamfilter that blacklists based on the sender email address. It is important however that the email address they use as a sender does exist, since spamfilters can apply a "sender verification check" to ensure that the sending address itself exists.

MxVault applies advanced methods to identify and block "bounce-spam".

CUSTOMER QUESTION:

What causes Bounce?

YOUR ANSWER:

Properly set up mail servers will not cause bounce spam and directly reject the message with a 5xx error code when the spammer tries to deliver it. Unfortunately there are many legitimate mail servers that are incorrectly set up. The spammer tries to deliver a spam message with your email address in the “From” field to an unknown address, the bad mail server accepts the messages for delivery, it then finds out that the destination user does not exist, and it will send a bounce email to your email address because it (wrongly!) believes you are the originating sender. Because these bounces do not come from spamming servers, but from legitimate servers, they are very hard to block by any spam filters.

MxVault

Questions and Answers

CUSTOMER QUESTION: What is SPF /DKIM / DMARC?

YOUR ANSWER: We have setup a SPF record for your domain, this will reduce the attractiveness for spammers to use your domain for sending out email. Also signing your emails with a DKIM certificate should further reduce the attractiveness to spoof your domain name for outgoing spam.

SPF

Sender Policy Framework (SPF) is an email validation system, designed to prevent unwanted emails using a spoofing system. To check this common security problem, SPF going to verify the source IP of the email and compare it with a DNS TXT record with a SPF content.

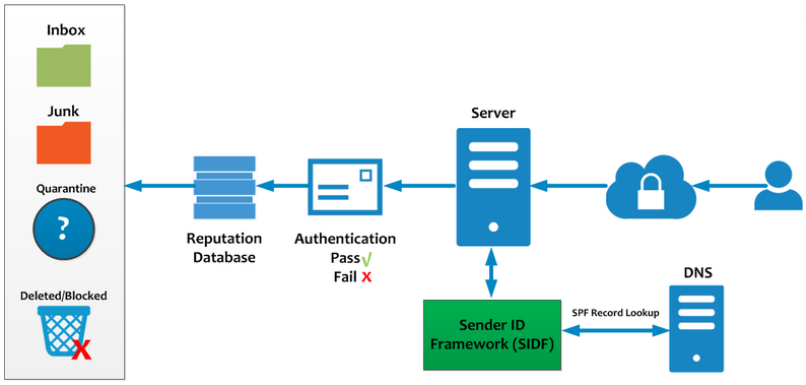

DKIM

DomainKeys Identified Mail (DKIM), is a method to associate the domain name and the email, allowing to a person or company assume the responsibiltity of the email.

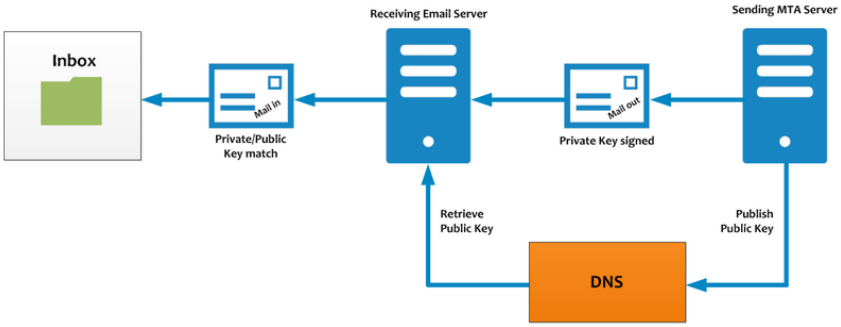

DMARC

DMARC, which stands for "Domain-based Message Authentication, Reporting & Conformance", is a technical specification created by a group of organizations that want to help reduce the potential for email-based abuse by solving a couple of long-standing operational, deployment, and reporting issues related to email authentication protocols.

DMARC standardizes how email receivers perform email authentication using the well-known SPF and DKIM mechanisms. This means that senders will experience consistent authentication results for their messages at AOL, Gmail, Hotmail, Yahoo! and any other email receiver implementing DMARC. We hope this will encourage senders to more broadly authenticate their outbound email which can make email a more reliable way to communicate.

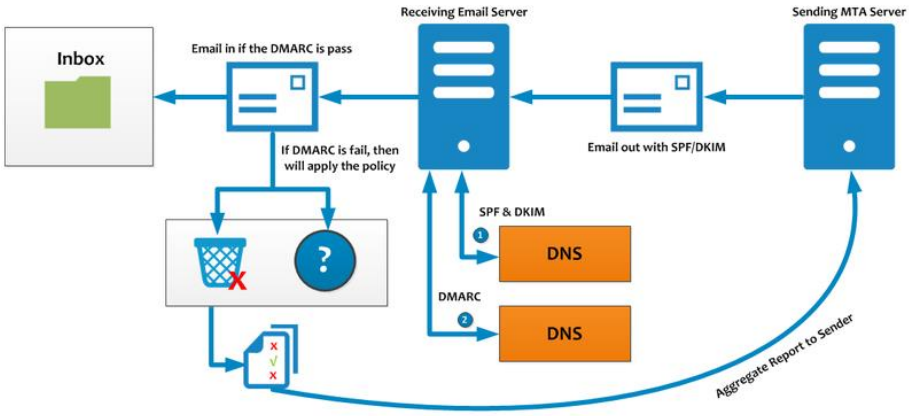

2. Domain Aliasing

CUSTOMER QUESTION: Do you offer domain aliasing?

YOUR ANSWER: For users with multiple domains we offer our free domain aliasing feature where domains can be added directly in the SpamPanel. This means that every email sent to one of your domain aliases will be sent to the same user on your main domain.

Users can access the Domain Aliases feature from the Domain Level (Domains > Overview > Click on domain) by going to Incoming > Domain Aliases.

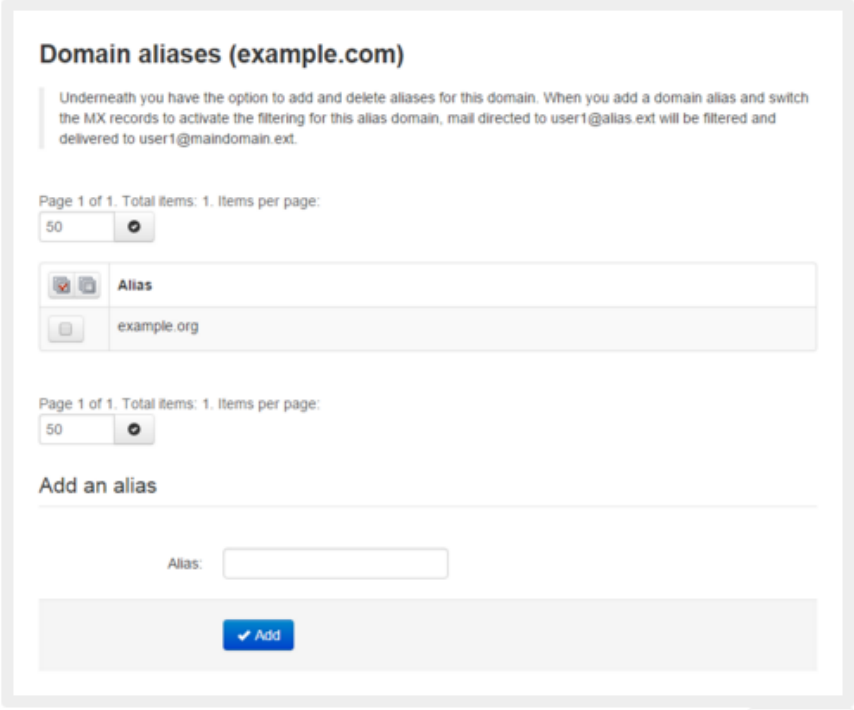

Assuming your main domain is example.com, the alias is example.org and your destination mail server is mail.example.org we will deliver any email sent to myemail@example.org to myemail@example.com on the mail.example.com SMTP server.
Do note, the email headers will still show the original recipient, myemail@example.org.
Control panel

Alias domains don't have separate access to the control panel, so you must continue to log in as you normally do.
All SMTP traffic to the alias domain is automatically rewritten to the main domain, so any changes or lookups on the main domain will include the alias' traffic as well.
Consequently, if you are searching for a specific email sent to a domain alias using the log search, the recipient will therefore show as user@maindomain.

3. Encryption TLS

CUSTOMER QUESTION: Does MxVault system supports TLS?

YOUR ANSWER: The MxVault system fully supports incoming connections protected using TLS. Deliveries are always made over TLS when supported by the destination mailserver. This way all email is securely transmitted.

4. Reporting SPAM

CUSTOMER QUESTION: How do I report SPAM via my browser?

YOUR ANSWER: When using a browser based email client, it's sometimes needed to be able to report spam to our systems. This can now be done by using a simple script that can push the message to our systems if you can view the whole source of the email via your browser.

Prerequisites:
Firefox or Chrome (Windows/Linux/OSX)
Addon - Greasemonkey
Please note this currently will only run on "Google Mail (Google Apps)", "Horde" & "RoundCube". If you have other web browsing mail-clients that you would like to include, please contact support@MxVault.com for more information.

Installation steps:
Install and enable addon from above

Download the script here (this should automatically install the script - if it does not, please open the file, and select all, and copy to your clipboard).
If automatic installation did not work, please click "Greasemonkey > New User Script > Use script from Clipboard".
Once the script is installed, please verify that the "Greasemonkey" icon is enabled.

Reporting Steps:
Login to your mail client via your preferred browser
Select the message you wish to report
View source of the message
A "Report Spam" button should appear on the page on the right.
Click "Report Spam"
If successful a pop-up box will appear confirming it was sent.
Close window
This script will not move, delete or change this message in anyway. This will simply send those headers and body content to our training servers.

CUSTOMER QUESTION: What do I do if I still receive SPAM?

YOUR ANSWER: If you are still receiving spam, submitting the emails to the central MxVault training system will first of all help to reduce the spam you're seeing, and it will allow us to retrieve detailed information on what may be causing your problems. Emails that did not pass our filtering systems, or that passed the filters because the sender/recipient was whitelisted are automatically excluded from the training systems.

Spampanel webinterface

In the Spampanel web interface we have a "Report Spam" button. You can upload a spam message on the page to train the spamfilter. It is a 'drag 'n' drop' style feature, meaning you can save the SPAM email to your system, then drag and drop the email into the "Report Spam" area. Note that only emails are accepted that have been processed by the cluster you're reporting to. Currently only the .eml format is supported.

Outlook add-on
For .msg format you can use the free "Outlook Email client add-on" to report spam which was not correctly blocked by our systems.

Thunderbird add-on

If you're using Thunderbird you can also use the free "Mozilla Thunderbird client addon" to report spam.

MailApp Apple OSX
If you're using Mail.app for Apple OSX you can also use the following tool to report spam.

IMAP
Finally, it's possible to report spam via our special IMAP system. We have instructions for Thunderbird and Outlook.

5. Error code 500

CUSTOMER QUESTION: What does this error code mean, please advise?

YOUR ANSWER: 500 - The server could not recognize the command due to a syntax error.
501 - A syntax error was encountered in command parameters or arguments.
502 - This command is not implemented.
503 - The server has encountered a bad sequence of commands.
504 - A command parameter is not implemented.
550 - No mailbox by that name is currently available, for example because it was not found, or because the command was rejected due to policy reasons, such as a full mailbox. Please clear the callout cache after the mailbox has been emptied.
551 - The recipient is not local to the server. The server then gives a forward address to try.
552 - The action was aborted due to exceeded storage allocation.
553 - The command was aborted because the mailbox name is invalid.
554 - The transaction failed.

6. Greylisting

CUSTOMER QUESTION: How does Greylisting work?

YOUR ANSWER:

We apply an advanced form of greylisting to help and stop a significant amount of spam with minimal resource usage. Although greylisting is a controversial technology, it is still highly effective when applied properly.

First of all it's important to mention that all nodes within the cluster are synchronized, and aware of the connections made to each other. Thus for the greylisting technology it does not matter to what node a connection is made. We also keep track centrally of "reputable hosts" to avoid any greylisting delays from known legitimate servers.
Greylisting works based on the 'triplet' information consisting of "sending server IP /24 subnet"/"sender email address"/"recipient email address". Whenever we receive a connection from an unknown 'triplet', we will temporarily reject (SMTP code 4xx) the connection for 10 minutes after seeing the first attempt.

A temporary reject in this case means that the sending server is requested to temporarily queue the email, and automatically retry later. Any legitimate SMTP server is required by the RFC to support this, and it's a fully automatic process of which the original sender will not receive any notification. It does not matter how often the server retries within the 10-minute interval or to which node, only after the 10 minutes we will accept the email.
This will result in a short delivery delay, therefore there is an advanced automatic system to minimize such delays. After accepting the email from an previously unknown 'triplet', the 'triplet' becomes 'white' to avoid temporarily rejecting connections from such triplets in the future. Furthermore whenever we have seen (at least) 5 different successful (white) triplets from the same IP /24 subnet or (at least) 2 different successful (white) triplets from the same subnet and sender email address, the subnet or subnet+address is added to an internal "greylisting whitelist" system to avoid greylisting connections from that IP.

All active mail servers delivering email to the servers will therefore not be influenced by the greylisting technology, since they will be on the internal "greylist whitelist". The greylisting technology is only applied to new unknown servers. Servers that have been blacklisted for sending out spam, will lose their whitelisted entry again so may shortly be greylisted for new connections.

Greylisted triplets become white after 10 minutes.
IP subnets are added to the "greylisting whitelist" after 5 white triplets.
IP subnet + sender address pairs are added to the "greylisting whitelist" after 2 white subnet+address pairs.

Greylist grey entries are expired after 8 hours.
Greylist white entries are expired after 60 days (if they have not been seen again).

Greylist triplets only apply to individual recipient domains, but the "greylisting whitelist" is shared across all domains for a cluster.
Be Advised: Most support questions regarding temporarily rejected connections are because customers see the temporary reject log entries, and are not aware that the message was NOT blocked/identified as spam. The message was only shortly delayed to verify that the sending server is behaving correctly (in accordance with the requirement for SMTP servers).
The "sending server IP /24 subnet" is basically the first part of the sending server's IP address. For example, if the server's IP was 222.153.243.117, then the string used in the 'triplet' would be '222.153.243'. This includes up to 256 (.0 to .255) servers, almost always within the same organisation. This means that if an organisation has several sending servers (typically within the same subnet), it does not matter which sending server makes the second attempt.

7. Message queuing

CUSTOMER QUESTION: How is mail queued?

YOUR ANSWER: Generally emails are directly delivered to the destination server and not stored locally on the filtering machines. However if the destination server is unavailable, all email sent to known destination recipients are queued locally on the filtering servers for delivery. Emails which have been permanently rejected by the destination server with a 5xx error code, will NOT be queued and are rejected by the systems. We have a list available with SMTP Error Codes, which can be found here. If the destination mailserver acts as catch-all for the domain, no recipients will be cached and therefore email will only be queued if the valid recipients have been explicitly set as Local Recipient in the MxVault control panel.

You can access the email queue from the webinterface, from which you can also manually force-retry delivery of a queued message.
Automatic Retry Schedule
Messages queued for known valid recipients because of temporary problems with the destination route (for example network problems)

are automatically retried for delivery at the following approximate intervals:

During the first 2 hours, delivery is retried at a fixed interval of 15 minutes.

During the next 14 hours, delivery is retried at a variable interval, starting at 15 minutes and multiplying by 1.5 with each attempt (e.g. after 15 minutes, then 22.5 minutes, then 34 minutes, and so on).

From 16 hours since the initial failure, until 4 days have passed, delivery is retried at a fixed interval of every 6 hours.
After 4 days we generate a bounce to the sender. If the bounce cannot be delivered immediately, it will be frozen. After this time, delivery of the message will have permanently failed. Optionally via the Software API the 4 days can be overruled to a longer (or shorter) period.
When a message is frozen, (when a message can neither be delivered to its recipients nor returned to its sender) no more automatic delivery attempts are made. A super administrator can "thaw" (force retry) such messages when the problem has been corrected.

MxVault caches valid recipients up to 4 days. When such entry has expired, MxVault will not queue email for such recipients and instead temporarily reject the message so it's queued on the sending server. The sending server in such case will automatically retry delivery. When using the "Local Recipients" feature, no caching is involved and MxVault continues to accept and queue the emails for all specified recipients.

Messages Queued
The SMTP RFC 5321 specifies a sending server must queue messages which cannot be directly delivered because of a temporary failure at the receiving end. Therefore in case of temporary issues with the email infrastructure, emails will not be bounced immediately but are instead queued on the sending server(s) and automatically retried for delivery. In case of downtime of the destination mailserver, messages are only accepted for delivery by the filter cluster if the recipient is known to be valid. Valid destination recipients are cached up to 96 hours.
When a destination server cannot be reached for 4 days, all messages will be bounced after 4 days and no new email will be accepted/queued until the destination server is back online. This 4 day period is conform the SMTP RFC. The reason it's not longer than 4 days, is because it's important for the sender to be aware that delivery of their message has been failing for 4 days so they can try and contact the recipient in another way.
Your own fallback server(s)

Please note that when you specify multiple destination routes, the MxVault system will assume you run your own fallback system. If the specified fallback server is not responding to recipient callouts, there will be no database built up of valid recipients internally. We recommend not to specify any fallback server(s) unless you've specifically designed your infrastructure to handle outages of the main destination server. Troubleshooting messages in the queue
If there are messages stored in the queue, this is always to to a (temporary) error delivering to the destination server. To investigate the issue:

Verify you have set a correct destination route (also ensure there are not multiple destination routes specified, normally there should just be 1 route)
Check the logs on your destination server to investigate why it's not accepting the delivery attempts
Run a telnet test to check the response of your destination mailserver
If after following these steps you still have an issue, please contact MxVault support providing a sample sender/recipient/date to investigate

8. Outlook Email Addon

CUSTOMER QUESTION: Is there an Email Addon / Plugin?

YOUR ANSWER: MxVault provides email client add-ons to report Spam which was not correctly blocked to our central systems.

Outlook Plugins:

32-bit:
http://syndication.globalmicro.co.za/downloads/software/SpamReport_Outlook_32bit.msi

64-bit:
http://syndication.globalmicro.co.za/downloads/software/SpamReport_Outlook_64bit.msi

Since version 2010, Office is also available in 64 bits. If you have a 64-bits Windows installed, you probably also have Office in 64 bits. In order for the add-on to work, you should install the correct version that matches your Outlook version.

For Outlook 2007 you have to use the 32 bits version.

These are also compatible with the new Outlook 2013, both 32 bits & 64 bits. Currently Outlook 2016 is only compatible with the 32 bit version.

9. SPAM

CUSTOMER QUESTION: What do I do when I Receive SPAM?

YOUR ANSWER: If you are still receiving spam, submitting the emails to the central MxVault training system will first of all help to reduce the spam you're seeing, and it will allow us to retrieve detailed information on what may be causing your problems. Emails that did not pass our filtering systems, or that passed the filters because the sender/recipient was whitelisted are automatically excluded from the training systems.

In the Spampanel web interface we have a "Report Spam" button. You can upload a spam message on the page to train the spamfilter. It is a 'drag 'n' drop' style feature, meaning you can save the SPAM email to your system, then drag and drop the email into the "Report Spam" area. Note that only emails are accepted that have been processed by the cluster you're reporting to. Currently only the .eml format is supported.

10. Virus Scanning

CUSTOMER QUESTION: How does MxVault handle virus scanning?

YOUR ANSWER: Viruses, malware and other online threats often spread via email, therefore it is important to virus-scan emails before they arrive to the mail-client of a user. MxVault actively blocks both spam AND its malicious attachments such as viruses, malware, ransomware, spyware and so on.

Due to the fact that viruses generally try to spread as spam emails, the majority of email viruses are already blocked before they are scanned with our antivirus engine, because of our anti-spam technologies. Thanks to this resource efficient and intuitive setup, even viruses not yet known to virus scanners are generally, safely put away in quarantine or rejected outright.

As an additional measure, we run the open-source ClamAV antivirus framework, whose virus definitions are updated every 30 minutes. Besides using the ClamAV databases, we have also added more datasets specialized in email virus problems, provided by several external partners, to ensure real-time, optimal protection against the latest virus outbreaks. Our internal reputation systems also contribute

to virus scanning and ensure optimal protection against not only spam, but also malware, phishing, and viruses.

Outgoing Services

1. Outgoing Filtering Services

KNOWLEDGEBASE:

The filters are very effective at blocking a large percentage of outbound spam/viruses, to prevent issues with your network reputation. It is very important however to pro-actively suspend any spamming customers/accounts, to stop the abuse at its source. If such accounts are not suspended/blocked, eventually there will always be a spam run which our engines could miss. You can prevent any such spam escalations (or other type of attacks from abusive customer accounts), by ensuring the account is locked down before it starts to cause real issues. Our systems allow you to quickly and easily identify such abusive accounts, before any third-party issues occur.

There are a number of ways that spammers can be monitored via our systems, the best method depends per-customer.

Best practice for smarthost users:

- Ensure all your smarthost authentication users are grouped as part of a single administrative domain (e.g. out.yourcompany.com)
- Configure your sending MTA to always include an end-user identification header
- Set your outgoing MxVault user account to use this identity header
- Manually/automatically locate abusive identities and shutdown the main spam source (and temporarily lock down the identity via our identity management as an immediate measure)

CUSTOMER QUESTION: How do you monitor outbound spam monitoring?

YOUR ANSWER: The filters are very effective at blocking a large percentage of outbound spam/viruses, to prevent issues with your network reputation. It is very important however to pro-actively suspend any spamming customers/accounts, to stop the abuse at its source. If such accounts are not suspended/blocked, eventually there will always be a spam run which our engines could miss. You can prevent any such spam escalations (or other type of attacks from abusive customer accounts), by ensuring the account is locked down before it starts to cause real issues. Our systems allow you to quickly and easily identify such abusive accounts, before any third-party issues occur.
There are a number of ways that spammers can be monitored via our systems, the best method depends per-customer.

Best practice for smarthost users:

Ensure all your smarthost authentication users are grouped as part of a single administrative domain (e.g. out.yourcompany.com)
Configure your sending MTA to always include an end-user identification header
Set your outgoing MxVault user account to use this identity header
Manually/automatically locate abusive identities and shutdown the main spam source (and temporarily lock down the identity via our identity management as an immediate measure)

2. What are ARF reports

CUSTOMER QUESTION: What are ARF Reports and how do you handle them?

YOUR ANSWER: An ARF (Abuse Reporting Format) report is an email format abuse report generated every time an outgoing sender's message gets blocked.

This can be set by adding the email address you wish to use at the domain's settings page of each domain under the 'Administrator contact', or by setting the Clusters default Administrator contact.

Using these reports are extremely useful for closing down spammers in your network, giving you an alert each time a spam message is sent. Each report contains the message that was blocked as an attachment,

as well as information regarding the outbound sender account that was used and a time stamp.

We recommend adding the email address you are using to the recipient whitelist, as the AFR reports contains a copy of the spam message that was blocked, and this can cause the reports to be blocked by the incoming filter. Alternatively you can use an email address designed for this that does not have any filtering on this address.

3. Block port 25 outgoing traffic to force smarthost usage

CUSTOMER QUESTION: How do I block Port 25 outgoing traffic?

YOUR ANSWER: Spammers may abuse a script on your servers to send out spam directly to port 25 destination email addresses, therefore bypassing the MxVault smarthost filtering. There should be no legitimate reason for scripts to communicate to port 25 on external servers, as all email should be handled locally on port 25.

You can simply use the iptables firewall to block all outgoing connections to port 25 and force all traffic to pass the MxVault filtering:

IPv4

To add:

- iptables -A OUTPUT -p tcp --dport 25 -j DROP
- To remove:
- iptables -D OUTPUT -p tcp --dport 25 -j DROP

IPv6

To add:

- ip6tables -A OUTPUT -p tcp --dport 25 -j DROP
- To remove:
- ip6tables -D OUTPUT -p tcp --dport 25 -j DROP

Be Advised: You need to ensure the rules are automatically applied after a reboot again.

4. Classifications

CUSTOMER QUESTION: In the MxVault Control Panel we use different classifications to describe why a message was rejected or temporarily rejected.

YOUR ANSWER: Temporarily Rejected (4xx SMTP response)

Temporarily Rejected
Messages which have been temporarily rejected, stay stored on the sending mail server. Legitimate mail servers always automatically retry delivery of such messages. Depending on the reason of the temporary reject, the message could get accepted at a subsequent delivery attempt. It's always possible to whitelist the sender to disable any checks and to ensure that the message will get accepted as soon as it's retried by the sending server.

Greylisted
Temporary rejection due to greylisting. This technology is only applied to new IP addresses which do not have a (good) reputation yet in our global systems. We do not apply "classical greylisting" so this should not cause any delays on your legitimate traffic. For new Local Cloud installations please allow up to 72 hours for the systems to "learn" about your traffic.

You have been denied authentication
This means that you have used incorrect outgoing authentication details too often in a short period of time. To resolve this, use the correct authentication details and wait a few moments and try again. This is to protect against brute-force attacks on your SMTP credentials.

Unable to verify destination address
This means the destination server is unreachable or temporarily rejecting the email traffic. You'll have to check the destination route set to ensure delivery is attempted to the correct server. The logs on the destination server should show why it is not accepting the delivery attempts.

Internal error

An internal error occurred, this should automatically resolve. If not, please contact support.
Per-minute connection limit exceeded
The sender has exceeded his/her per-minute limit.
Too Many Connections
Too many connections from the sending server. Ratelimited.

Too Many Concurrent SMTP Connections
There is a hard-coded limit of 10 concurrent SMTP connections per IP to protect the systems against attack. Please ensure that the sending mail server only opens up a maximum of 10 concurrent connections to avoid hitting this limit.

Too many messages. Please wait for a while and try again.
This indicates that the outgoing user has exceeded the maximum amount of messages configured for that outgoing user to be sent.

Mail for this domain cannot be accepted right now; please retry (Unable to handle in active connection.)

Within a single SMTP connection, it is possible to deliver a message to different recipients. The SMTP protocol only allows you to either "accept" OR "reject" the email, without distinguishing between the different recipients. In case one of the recipients has different filtering settings, we cannot "accept" or "reject" the message as the classification may differ per-recipient. In such case we return a temporary rejection, so the sending server will retry delivery individually for the recipients allowing to classify each message separately. Most SMTP servers retry immediately, and hence there will be no delivery delay. If all recipients are sharing the same filtering settings, the message will be immediately accepted for all recipients (or rejected) without this temporary reject. In case a delay is experienced, the sender can instead configure their server to either immediately retry (to prevent such delay), or to open a separate delivery connection for each recipient.

Rejected (5xx SMTP response)

Messages which have been rejected are blocked by the system. Generally, these messages can be reviewed in the "Spam quarantine", from where they can be released. It's always possible to whitelist the sender to disable any checks and to ensure that the message will get accepted as soon as it's resent by the sender.

Lines in message were longer than user maximum

This means that line within the email is longer than the set maximum. The RFC 5322 (SMTP 5321) specifies a maximum line length of 998. Normal email clients always enforce this limit to avoid delivery problems. The problem should be resolved at the sender side, or the check can be disabled.

Message had more parts than the user maximum (Too many MIME parts)
This refers to the amount of MIME parts that a message contains. The default limit is set to 100. This can be de-passed and triggered with excessive amounts of attachments or other MIME parts.

Sending server used an invalid greeting
The sender has used an invalid HELO/EHLO. This could be either because an IP address is used for the HELO, or because the HELO contains an invalid character, for example: underscore (_). The RFC states that a FDQN (Fully Qualified Domain Name) MUST be used.

Considered spam
Our systems considered this message as SPAM and quarantined the message. Releasing the message from quarantine will report it as a classification mistake to correct our systems.

SPF failure
This means that the SPF (Sender Policy Framework) has been broken. If this is legitimate mail, then this could be due to a forwarding construction. Please see our SPF knowledgebase article for more information.

Pyzor
Pyzor is a content related classifier based on collected/reported data from our datasets. Releasing the message from quarantine will report it as a classification mistake to correct our systems directly.

Sending server is missing DNS records
The sending server is missing MX records or A records. Please note that any DNS changes only take effect after the initially set TTL has expired.

Destination address does not exist
The destination server is rejecting the connection with a 5xx permanent failure. The logs on the destination server will show why the message was rejected. You'll have to resolve the problem on the destination server to ensure it accepts the email.

Phishing attempt detected

Our systems detected a phishing attempt. Releasing the message from quarantine will report it as a classification mistake to correct our systems.

Date header far in the past or future.
This classification means that the date header of the email is more than the default 7 days in the past or future. Releasing this will only deliver the message to the recipient. This is something the sender will need to resolve.

Bad header count (Message incorrectly formed)
Emails should never contain duplicate headers such as "Subject" or "To". In case such duplicate headers are found, the message will be rejected until the underlying bug is fixed in the email sending software.

Blacklisted sending server
The sending server has been blacklisted on the IP blacklist.
Sending server listed on multiple DNSBL
The sending server has been found on multiple blacklists.

Sending server attempted too many invalid addresses
The email sending server has attempted to deliver email to too many invalid email addresses in a certain time period. Please retry again later.

Blacklisted sender
The sender was added to the custom sender blacklist.

URLBL
A URL within the email has been listed on several blacklists. The rejection message contains more information about the responsible list.

UCEPP
A token was detected in the message that has been seen in recent spam (e.g. URL, IP, phone number, or other specific details).

External Pattern Match
The layout & format of the email matches known spam emails already listed. The rejection message contains more information about the responsible list.

User-specified blackhole address
A user specified /dev/null Address. This email will not get delivered anywhere.

Combined Score
The "combined" result provides a weighted classification score of the different classifiers. Depending on the configured "quarantine threshold", the message will be rejected as spam or accepted. A quarantine threshold score of 0.9 is recommended. To be more tolerable for senders using a wrong HELO/PTR/IP configuration, a score of 0.91 can be set. The lower the quarantine threshold, the more messages will be quarantined as spam. The SMTP message returned for this classification is "High probability of spam" to the sender. Please ensure to release the message from quarantine if it's legitimate, this will adjust the scoring in our various databases.

CRM114
CRM114 is a statistical content check. When a message gets blocked by this classifier on our systems, then this mean there has been a close match within the email that corresponds to an already seen spam message.

Subject contains invalid characters.
When a message is rejected with "550 Subject contains invalid characters" the email subject will have non-ASCII characters, which is not allowed by the RFC. To include non-ASCII characters in subjects, the subject is required to be properly encoded, for example with UTF-8. Any normal mail client will automatically handle that for you, so it's likely a bug in a custom written script that generated the invalid subject. The evidence header for this classification will show "Badly formed Subject header".

Tokens
Global Tokens (Hosted cloud / Local Cloud)
These are statistical content checks that are built based on data collected from all our clusters and clients worldwide.

Cluster Tokens (Local Cloud Only)
This is similar to the global tokens, but based specifically on your Local Cloud traffic and reports.

Sanesecurity
We make use of certain datasets from Sanesecurity. To decode Sanesecurity signatures please check here.

Safebrowsing

In case your message has been rejected with "safebrowsing" in the rejection message, it means it has been (recently) listed by Google as hosting malicious files.

Header is too long
MxVault by default will reject emails with excessive large header values, as this is a common indicator for non-legit emails.

Restricted characters in address
In case your message has been rejected with "550 restricted characters in address" in the rejection message, it means that the recipient address contains a character that is not accepted by the system, for example: "&". You can control which characters are allowed for a domain on the "Domain settings" page.

Relay not permitted
In case your message has been rejected with "550 Relay not permitted!" in the rejection message, it means that delivery was attempted to the incoming filtering service on port 25 to a domain which has not (yet) been added to the filtering solution. To resolve this, please add the domain to the incoming filtering service. If you're trying to use the outgoing filtering service, please ensure to use the outgoing filtering service port 587 instead.

Message submission is for authorised users only!
This indicates you're attempting delivery via our outgoing email filter on port 465/587 (default). If you're receiving this response to an incoming email delivery attempt, your mail server is wrongly set up (and likely a misconfigured version of Lotus Domino). If you're trying to send outgoing email, please ensure to provide a valid username/password to authenticate.

Legitimate bounces are never sent to more than one recipient.
In case your message has been rejected with "Legitimate bounces are never sent to more than one recipient" in the rejection message, it means that the mail server was trying to deliver an email to multiple recipients with an empty "MAIL FROM:<>" (return-path). The SMTP RFC 5.3.2.1 indicates that null sender emails (bounces) can never be sent to multiple recipients, so there may be be a misconfiguration on the mailserver.

Destination address is not configured.
This usually means that the filtered domain is using 'Local Recipients' and that specific email address in not in their list of approved recipients.

The content of this message looked like spam.
This indicates the message has been blocked based on our content scanners, as similar messages have been reported as spam.

Unrouteable address
This error occurs if there is a (permanent) network error delivering to the destination mail server. This issue is unrelated to the MxVault software and indicates a network problem. Possibly the DNS servers of the domain are broken, or they cannot be reached from the filtering server. Alternatively it's possible the destination hostname or IP does not exist, or is unreachable because of a permanent issue. You can check for DNS errors on the following page: http://dnscheck.sidn.nl/. Please contact your network administrator to investigate any networking issues.

We do not accept mail from this address
This error occurs if the sender has been manually added to the "Sender blacklist" for the receiving domain.

We do not accept message/partial messages here
Before people had a permanent internet connection, sending larger emails was time-consuming and often failed. Therefore, older email clients sometimes still break up large emails into separate parts for delivery. This old email feature is not used anymore nowadays, and imposes a severe risk as it makes detection of viruses impossible (as viruses would be split over separate emails before being assembled again by the destination email client). Please ensure to resolve your email client settings to split up larger emails.

Failed sender's DMARC policy, REJECT
This error occurs if the sender's domain has a strict DMARC policy in place. If the sender's DMARC record is set to "REJECT" and the messages come from IP addresses that are not in the sender's SPF, then these are rejected and not quarantined.

Accepted (2xx SMTP response)
ACCEPTED

Messages that display the 'Accepted' response have not necessarily been delivered. It means the message has been accepted for delivery. If immediate delivery fails, the message will be automatically retried. If the destination server rejects the email, a bounce will be generated to the sender.

Message looked like non-spam

This message was accepted for delivery based on our content checks. Reporting the message as spam will help correct our systems.

Accepted, DNSWL
The sending server is listed on several DNS-Whitelists. This means no spam has been seen recently from this sending server. Reporting the message as spam will help correct our systems.

Accepted, whitelist
The sender has been placed on a manual whitelist by the recipient. Removing the sender/recipient from the whitelist will prevent spam getting through.

5. Filtering Technology

CUSTOMER QUESTION: What filtering methods are used?

YOUR ANSWER: The filtering methods and system of MxVault have been specifically designed to avoid false positives. For that reason, many different checks are performed to avoid making mistakes based on only a single classifier. Two levels of filtering can be distinguished. Filtering at the "SMTP level" and filtering at the "DATA level". Thanks to the combination of many different advanced filters and the compliance with the RFCs on how to handle connections, the technologies ensure email can never disappear. The sender is always informed by their sending server that the message was rejected - in addition, messages blocked at the "DATA level" are available in the quarantine system.

SMTP Level
As much as possible the incoming email connections are not blocked until after the "rcpt to:" SMTP command. This way it is ensured that the connection is properly logged belonging to the recipient domain in the logging server, to ensure keeping an easy overview of all connections made to a certain recipient. Before the "DATA level" is reached, the connection is checked to see if it follows the RFC standards, is not listed on internal and/or external blacklists, and several other things. If the connection appears to be coming from an unknown source that has not a good reputation yet in our systems, it may be temporarily rejected with a 4xx code. In that case the sending server will queue the email, and automatically retry delivery. After 10 minutes the connection will be accepted by the cluster (on any of the filtering nodes), and the internal whitelists are adjusted to avoid causing such a delay in the email delivery the next time. This concept is also known as greylisting, however the MxVault implementation is a lot smarter than traditional greylisting systems since all nodes are fully synchronized, and only connections from servers that are unknown in the MxVault network are temporarily delayed. Therefore, email delays because of greylisting on active filtering clusters are quite uncommon and generally do not cause any problems for the recipients. If the connection appears to originate from a spamming source, often the connection is also temporarily rejected with a 4xx code. This way even if the server would have been wrongly listed (e.g. on an external blacklist) as a spamming source, or if the spamming problem has been resolved on the sending server, the email still does not get lost and will be delivered to the final recipient. Only if the connection is from a known, spam-only source, or if the behavior is in direct conflict with the RFC standards, a connection may be permanently rejected with a

5xx error code. If that ever would happen for a legitimate sender, the sender will always receive a bounce notification from their sending server. This issue only occurs when there are serious problems with the sending server that should be resolved at the sender's side.

Data Level
After the "DATA level" is reached the system will scan the email content of the message based on a combination of advanced statistical filtering technologies, spam fingerprint databases, viruses, phishing, and spyware. Email detected as spam is either temporarily rejected (4xx error code) or permanently rejected (5xx error code) depending on the total score. Email which is permanently rejected at this level as spam is quarantined and available for release (except for viruses). In case a legitimate email would have been permanently blocked, the sending server will also always inform the sender that the email did not arrive.

6. Generate DKIM certificate

CUSTOMER QUESTION: Why should I use DKIM?

YOUR ANSWER: There are several advantages to using DKIM to sign your outgoing emails:
The recipient is able to verify that the message did come from the specified sender;
The recipient is able to verify that the message content (and important headers, like the subject) has not been altered;
It lowers the chance of the email being identified as spam, although this is not the primary reason to sign.

If a spammer is trying to abuse your domain or email address, using DKIM the changes of spam getting through will decrease. Many email servers (e.g. Yahoo!, Gmail, MxVault) will check for a valid DKIM signature on incoming email.

How does it work?
DKIM adds a special DKIM Signature to the email headers. This signature contains a hashed value of the content (both important headers and the body). When a server that is checking for DKIM receives an email, it will do the following:
Retrieve the public key from the DNS of the sending domain.
Use the key to decrypt the signature.

Verifies the content.

The exact actions a mail server takes when it discovers an invalid signature depend on the configuration of that server.

Contact your administrator to setup DKIM /DMARK and or SPF

7. SMTP error codes

CUSTOMER QUESTION: I receive 5xx error code

YOUR ANSWER: Below are a list of possible 5xx error codes:

- 500 - The server could not recognize the command due to a syntax error.
- 501 - A syntax error was encountered in command parameters or arguments.
- 502 - This command is not implemented.
- 503 - The server has encountered a bad sequence of commands.
- 504 - A command parameter is not implemented.
- 550 - No mailbox by that name is currently available, for example because it was not found, or because the command was rejected due to policy reasons, such as a full mailbox. Please clear the callout cache after the mailbox has been emptied.
- 551 - The recipient is not local to the server. The server then gives a forward address to try.
- 552 - The action was aborted due to exceeded storage allocation.
- 553 - The command was aborted because the mailbox name is invalid.
- 554 - The transaction failed.

8. Virus scanning

CUSTOMER QUESTION: How do you handle virus scanning?

YOUR ANSWER:

Viruses, malware and other online threats often spread via email, therefore it is important to virus-scan emails before they arrive to the mail-client of a user. MxVault actively blocks both spam AND its malicious attachments such as viruses, malware, ransomware, spyware and so on.

Due to the fact that viruses generally try to spread as spam emails, the majority of email viruses are already blocked before they are scanned with our antivirus engine, because of our anti-spam technologies. Thanks to this resource efficient and intuitive setup, even viruses not yet known to virus scanners are generally, safely put away in quarantine or rejected outright.

As an additional measure, we run the open-source ClamAV antivirus framework, whose virus definitions are updated every 30 minutes. Besides using the ClamAV databases, we have also added more datasets specialized in email virus problems, provided by several external partners, to ensure real-time, optimal protection against the latest virus outbreaks. Our internal reputation systems also contribute to virus scanning and ensure optimal protection against not only spam, but also malware, phishing, and viruses.

Archive Services

1. Searching the archive

CUSTOMER QUESTION: I am logged in as a user and searching my archived mail, but no results are showing. My primary domain is *primarydomain.com*. We have an alias domain called *aliasdomain.com*. My primary email address is *user@aliasdomain.com*. Why is this happening?

YOUR ANSWER: All emails that are archived are sent to the primary domain of the company. When users log in to search their archive, they will find all of their emails that are sent and received to and from all of their aliases as long as their primary email address is at their primary domain. If a user has their primary email address at a domain different to their primary domain, then as a user, they will not be able to search the archive for those mails. However, the admin of the domain **will** be able to search for all emails, including those sent and received to and from the alias domains for all users, regardless of the primary email addresses.