MailGate Export Log Entity/Attribute Definitions

1. **Summarized Appliance Data (appliance_stats)** - Each row in the Summarized Appliance Data table represents one piece of information about the appliance as a specific point in time. The row data is essentially a name/value pair. The other fields in the row give additional context to the name/value pair. Data collected in this table includes appliance and system level information, appliance status, report checkpoint calculations, and other information. About 150,000 records per day are collected in a typical usage pattern.
    a. Id – Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster
    b. Name – name of the variable in the name/value pair
    c. Value – the value of the variable in the name/value pair
    d. bucket_time – Unix timestamp which represents the time what the name/value pair's data was generated
    e. Business_unit_id – the ID of the business unit which the variable represents. If business units are not used, this value will represent the default, which means the variable applies to the entire appliance. This field is a foreign key to the Business Units Data table.
    f. Last_update – the date of last modification of the row. Generally, this will be the same as the bucket_time
    g. appliance_id – duplicate of the ID of the appliance

2. **Business Units Data (business_unit)** – Each row represent a uniquely defined business unit within the policy structure. At least one global business unit row will exist in this table
    a. Id – Unique ID (primary key) for the business unit
    b. Name – Business unit name
    c. Description – Business unit description
    d. Last_update – the date of last modification of the row.

3. **Alerts/Audit log data (event_log)** - This log export includes all alerts and all audit log entries created by the MailGate appliance.
    a. Id – Unique ID (primary key) for each audit log entry
    b. Type – This value indicates whether the data in this row is replicated to just the local cluster, or to a remote central cluster if available (1=local, 2-=remote)
    c. event_id – the name of the mailgate service and internal MailGate event which generated the audit/alert record creation
    d. category_id –the type of the event
        i. 1=Anti-spam alert category,
        ii. 2=Anti-virus alert category,
        iii. 3=Backup alert category,
        iv. 4=Edge defense alert category,
        v. 5=License alert category,
        vi. 6=Mail Queue alert category,
        vii. 7=Network Connectivity alert category,
        viii. 8= System alert category,
        ix. 21=System alert category,
        x. 101=Admin login audit,
        xi. 102=AV settings audit,

        xii. 103=Edge defense settings audit,
        xiii. 104=General administration settings audit,
        xiv. 105=Message tracking audit.
        xv. 106=Policy manager settings audit,
        xvi. 107=Relay settings audit.
        xvii. 108=Reports config settings audit,
        xviii. 109= User login audit
        xix. 110= User and Group Management audit
        xx. 111=RBAC audit,
        xxi. 112= Exports config settings audit

    e. source – username or system name of the user/system that triggered the event. System events are generally cataloged under MailGate or Admin
    f. severity – Event severity (1=critical, 2=moderate, 3=low, 4=information)
    g. log_time – Unix timestamp for the event
    h. arg0 – 1st value for the event description. This can be an error description, an IP address, or other descriptive information about the event or alert. Content varies based on the event_id field.
    i. arg1 – 2nd value for the event description. This contains additional info not included in Arg0. It will often be empty.
    j. arg2plus – additional description for the event description, if necessary.
    k. alert_resolved – status of the event (1=resolved, 0=Not Resolved)
    l. Business_unit_id – the ID of the business unit for which the variable is for. A value of -1 means it is not associated with any specific business unit
    m. appliance_id – Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

4. **Inbound Messages Data (inbound_mail_log)** – Each entry in this log represents one policy evaluation process of an inbound mail message. Depending on the policy engine rules, there often can be one entry per recipient.

    a. db_id – Unique ID (primary key) for each inbound mail log entry
    b. message_id – The policy engine generated a unique identified called the Policy Engine ID (PE ID) for each email received for a recipient. This ID can be seen in the tracker, and can be used for correlation from log to live data.
    c. received – Unix timestamp when the inbound messages is received
    d. from_email – email address of the message sender
    e. from_display – display name of the sender, if known
    f. from_domain – sender's domain name MailGate (use REVERSE function, characters are in reverse order)
    g. subject –subject line of the message
    h. size – total message size, in bytes
    i. wl_size – word list size. Value is used internally. May be useful in limited debug situations.
    j. classification – Anti-spam classification, if AS features are turned on. Values are as follows: 0 – Unclassified, 1 – Legit, 2 - Bulk ,3 – Junk
    k. av_classification – this field displays the types of virus found, if one is found: (9 - virus detected by Virus Outbreak Detection, 7 – Clean, 3 – Virus, Null – no attachment) Note: There may be other values, but these have not been detected recently.
    l. reason – SPAM header information. All the information generated by the anti-spam evaluation.
    m. orig_disp – Processing action taken after policy evaluation
        i. 0=None
        ii. 1=Drop
        iii. 2=Return
        iv. 4=Admin Quarantined

v. 8=User Quarantined
vi. 16-Delivere
vii. 32=Deferred
viii. 64=Detained
ix. 128=Delete
x. 256=Redirect to Secure Message

n. spam_rating – spam category, based on the type of spam (8=Adult, 16=Scam, 0=Broadcast or Legit)
o. stored – indicates if message is stored on the box or not. Generally, messages will only be stored on the box if they are quarantined (1=yes, 0=no)
p. store_type – Internal use only
q. client_ip –IP address of the machine forwarding this email to the MailGate appliance
r. virus_name – name of the virus found by the Anti-Virus engine
s. av_info – additional Anti-Virus information from the AV engine
t. av_error_on_scan – indicates if there was an error during AV scanning (1=yes, 0=no)
u. av_down – Indicates that the AV engine was down, so MailGate was unable to perform a virus scan (1=yes, 0=no) Note: We have never seen this happen
v. av_user_notified – this flag indicates whether a notification was sent back to the sender of the email about the virus. User notification to external senders is a configurable feature of MailGate (1=yes, 0=no)
w. mta_appliance_id – ID of the appliance which processed the message. It is generally the same as the appliance_id below
x. receiving_conn_id – foreign key to the receiving connection ID from which this mail item was received
y. appliance_id – Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster.

5. **Inbound Messages Data per Recipient (inbound_mail_recipient)** – This log contains one entry per recipient of an inbound message. It is a child record of the Inbound Message Data table.
a. db_id – unique database ID for each inbound email message recipient
b. message_db_id – Foreign key to the Inbound Message data parent row
c. recipient_index – An index to distinguish each recipient entry, starting with 0, if there was a combined evaluation process in the inbound_mail_log
d. recipient_address –email address of one recipient of the email
e. recipient_domain – the domain name of the recipient of the email (use REVERSE function, characters are in reverse order)
f. recipient_primary_address – recipient's primary email address, This will be different than the recipient address if the email was sent to an alias
g. group_id – ID of the recipient's group as defined user MailGate's User and Groups options
h. user_action – Action taken by user on quarantined mail through User UI
i. time_action - Time that user action was taken
j. marker – hash of the message used for spam auto-classification.
k. message_id – The policy engine generated a unique identified called the Policy Engine ID (PE ID) for each email received for a recipient. This ID can be seen in the tracker, and can be used for correlation from log to live data
l. received – Unix timestamp when the message is received
m. from_email –email address of the sender
n. from_domain – sender's domain (use REVERSE function, characters are in reverse order)
o. from_display – sender's display name, if known
p. subject – subject line of the message
q. size – total message size, in bytes

r. classification – Anti-spam classification, if AS features are turned on. Values are as follows: 0 – Unclassified, 1 – Legit, 2 - Bulk ,3 – Junk

s. av_classification – this field displays the types of virus found, if one is found: (9 - virus detected by Virus Outbreak Detection, 7 – Clean, 3 – Virus, Null – no attachment) Note: There may be other values, but these have not been detected recently.

t. virus_detected – indicates whether a virus was detected (1=yes, 0=no)

u. original_disposition – Processing action taken after policy evaluation

  i. 0=None
  ii. 1=Drop
  iii. 2=Return
  iv. 4=Admin Quarantined
  v. 8=User Quarantined
  vi. 16-Delivere
  vii. 32=Deferred
  viii. 64=Detained
  ix. 128=Delete
  x. 256=Redirect to Secure Message

v. spam_rating – spam category, based on the type of spam (8=Adult, 16=Scam, 0=Broadcast or Legit)

w. stored – indicates if message is stored on the box or not. Generally, messages will only be stored on the box if they are quarantined. (1=yes, 0=no)

x. mta_appliance_id – ID of the appliance which processed the message

y. business_unit_id – The business unit of the email recipient, as defined via MailGate groups and business unit association. If not using LDAP recipient verification or named users, value will always equal default business unit.

z. appliance_id – Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

6. **Inbound Messages Attachment Data per Recipient (inbound_mail_attachment)** – An entry that indicates an attachment was received for an inbound participant in an email. There is one email per attachment per user receiving the attachment. *Note: report will contain data only if "Collect message attachment data" from "Mail Policies" – "Settings"-> "Message Attachment Data" is enabled.*

a. db_id – Unique ID value for this row

b. recipient_db_id – database id of inbound Message Key per Recipient row. This is a foreign key to its parent table

c. attachment_modification – indicates if attachment was modified or not, by stripping the attachment or removing a virus

d. attachment_name – the name of the attachment . This would be the file name without the extension

e. attachment_extension – the extension of the attachment

f. attachment_type – the file type of the attachment. The true type of the file is determined by inspecting the content of the attachment to determine if it matches a known file type.

g. attachment_encoded_size - base64 encoded size of the attachment in bytes

h. attachment_size – size of the attachment in bytes

i. received - Unix timestamp when the messages is received

j. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

7. **Summarized Inbound Policy Execution Data (inbound_policy_execution_summary)** – This log contains summary information on the execution of each defined inbound policy over the archive time period in the MailGate configuration
    a. id – Unique ID value for this row
    b. policy_name – name of the inbound policy defined in the filter and delivery policies
    c. policy_phase – policy phase (10=filter, 20=delivery)
    d. bucket_time – Unix timestamp for the row
    e. count – number of times policy was triggered, by meeting its IF condition, in the processing of email
    f. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
    g. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

8. **Summarized Inbound Policy Tag Data per Recipient (inbound_recipient_policy_summary)** – This log summarized all policy tag values which were set on emails intended for a specific recipient over the time period configured in the MailGate export definition.
    a. Id – Unique ID value for this row
    b. Recipient – recipient email address which is being summarized
    c. tag_id – internal ID of the tag. To get the mapping of internal tags to their names, especially for custom tags, this must be retrieves using the local version of remote support option.
    d. bucket_time – Unix timestamp for the row
    e. count – number of times policy execution triggered the setting of this specific tag on any incoming record
    f. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
    g. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

9. **Summarized Spam Report Data per Recipient (inbound_recipient_summary)** – Summary of all email traffic received by a particular recipient over the time period configured in the MailGate export definition.
    a. id – Unique ID value for this row
    b. recipient – recipient email address being summarized
    c. bucket_time - Unix timestamp for the row
    d. virus – Number of emails with viruses destined for recipient
    e. junk – Number of emails classified as junk intended for recipient
    f. bulk - Number of emails classified as bulk mail intended for recipient
    g. legit - Number of emails classified as legitimate intended for recipient
    h. total – Total Number of emails intended for recipient
    i. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
    j. appliance_id - the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

10. **Summarized inbound Spam Report Data per Domain (inbound_spam_by_domain_summary)** – Summary of all email traffic sent to a specific internal domain over the time period configured in the MailGate export definition.
    a. id – Unique ID value for this row

b. domain – domain to which the mail was sent, and is being summarized. The characters of the domain name are reversed in the text, so you should use the REVERSE function in any SQL statement accessing this data
c. bucket_time - Unix timestamp for this row
d. virus – Number of emails with viruses received for a specific internal domain
e. junk – Number of emails classified as junk received for a specific internal domain
f. bulk – Number of emails classified as bulk mail received for a specific internal domain
g. legit – Number of emails classified as legitimate received for a specific internal domain
h. total – Total Number of emails received for a specific internal domain
i. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
j. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

11. Summarized Inbound Spam Report Data per Domain (inbound_sender_domain_summary) - Summary of all email traffic received from a specific device in a domain sent inbound to MailGate over the time period configured in the MailGate export definition.

a. id – Unique ID value for this row
b. ip –IP address of the external email sending device being summarized
c. domain –domain of the inbound email sender
d. bucket_time – Unix timestamp for this row
e. virus – Number of emails with viruses received from the device/domain
f. junk - Number of emails classified as junk received from the device/domain
g. bulk – Number of emails classified as bulk mail received from the device/domain
h. legit - Number of emails classified as legitimate received from the device/domain
i. total – Total Number of emails received from the device/domain
j. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
k. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

12. Summarized Inbound Sender Data (inbound_sender_summary) - Summary of all email traffic received from a specific sender address in a domain over the time period configured in the MailGate export definition.

a. id – Unique ID value for this row
b. sender –email address of the sender being summarized
c. domain –domain of the sender
d. bucket_time – Unix timestamp for the row
e. virus – Number of emails with viruses received from the specific email address
f. junk – Number of emails classified as junk received from the specific email address
g. bulk – Number of emails classified as bulk mail received from the specific email address
h. legit - Number of emails classified as legitimate received from the specific email address
i. total – Total Number of emails received from the specific email address
j. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
k. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

13. Summarized Virus Infections Data (inbound virus summary) – Summarized information collected by the anti-virus scanning process over the time period configured in the MailGate export definition. It is useful for finding virus trends.

a. id – Unique ID value for this row
b. virus_name – Name or type of the virus that was identified in at least one email
c. av_classification - this field displays the types of virus found, if one is found: (9 - virus detected by Virus Outbreak Detection, 7 – Clean, 3 – Virus, Null – no attachment) Note: There may be other values, but these have not been detected recently.
d. bucket_time - Unix timestamp for the row
e. total – total viruses found of this type/name
f. most_recent – most recent date/time that this specific infection was found
g. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
h. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

14. Outbound Message Data (outbound_mail_log) - Each entry in this log represents one policy evaluation process of an outbound mail message. Depending on the policy engine rules, there often can be one entry per recipient.

a. db_id – Unique ID (primary key) for each outbound mail log entry
b. message_id – The policy engine generated a unique identified called the Policy Engine ID (PE ID) for each email received for a recipient. This ID can be seen in the tracker, and can be used for correlation from log to live data
c. received – Unix timestamp when outbound message was received by MailGate
d. from_email – email address of the message sender
e. from_display – display name of the sender
f. from_domain – sender's domain name MailGate (use REVERSE function, characters are in reverse order)
g. subject – subject line of the message
h. size - total message size, in bytes
i. wl_size – word list size, value is used internally. May be useful in limited debug situations
j. classification - Anti-spam classification, if AS features are turned on. Values are as follows: 0 – Unclassified, 1 – Legit, 2 - Bulk ,3 – Junk
k. av_classification – this field displays the types of virus found, if one is found: (9 - virus detected by Virus Outbreak Detection, 7 – Clean, 3 – Virus, Null – no attachment) Note: There may be other values, but these have not been detected recently.
l. reason – SPAM header information. All the information generated by the anti-spam evaluation.
m. orig_disp – Processing action taken after policy evaluation
   i. 0=None
   ii. 1=Drop
   iii. 2=Return
   iv. 4=Admin Quarantined
   v. 8=User Quarantined
   vi. 16-Delivere
   vii. 32=Deferred
   viii. 64=Detained
   ix. 128=Delete
   x. 256=Redirect to Secure Message
n. spam_rating – spam category, based on the type of spam (8=Adult, 16=Scam, 0=Broadcast or Legit)

o. stored – indicates if message is stored on the box or not. Generally, messages will only be stored on the box if they are quarantined (1=yes, 0=no)
p. store_type – internal use only
q. client_ip - IP address of the machine forwarding this email to the MailGate appliance
r. virus_name – name of the virus found by the Anti-Virus engine
s. av_info – additional Anti-Virus information from the AV engine
t. av_error_on_scan - indicates if there was an error during AV scanning (1=yes, 0=no)
u. av_down – Indicates that the AV engine was down, so MailGate was unable to perform a virus scan (1=yes, 0=no) Note: We have never seen this happen
v. av_user_notified – if user was notified about virus this flag indicates whether a notification was sent back to the sender of the email about the virus. User notification to external senders is a configurable feature of MailGate (1=yes, 0=no)
w. mta_appliance_id – ID of the appliance which processed the message. It is generally the same as the appliance_id below
x. receiving_conn_id foreign key to the receiving connection ID from which this mail item was received
y. appliance_id – Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

15. Outbound Messages Data per Recipient (outbound_mail_recipient) - This log contains one entry per recipient of an outbound message. It is a child record of the Outbound Message Data table.

a. db_id – database ID for each inbound email message recipient
b. message_db_id – Foreign key to the Outbound Message data parent row
c. recipient_index – An index to distinguish each recipient entry, starting with 0, if there was a combined evaluation process in the inbound_mail_log
d. recipient_address – email address of this recipient of the email
e. recipient_domain – the domain name of the recipient of the email (use REVERSE function, characters are in reverse order)
f. recipient_primary_address - recipient's primary email address, This will be different than the recipient address if the email was sent to an alias, and the recipient is an internal user references through an LDAP lookup which provides alias information
g. group_id – ID of the recipient's group as defined user MailGate's User and Groups options, if the recipient is from an internal domain
h. user_action – Action taken by user on quarantined mail through User UI
i. time_action – Time that user action was taken
j. marker – hash of the message used for spam auto-classification.
k. message_id – The policy engine generated a unique identified called the Policy Engine ID (PE ID) for each email received for a recipient. This ID can be seen in the tracker, and can be used for correlation from log to live data
l. received – Unix timestamp when the message was received
m. from_email – email address of the sender
n. from_domain`- sender's domain (use REVERSE function, characters are in reverse order)
o. from_display – sender's display name, if known
p. subject - subject line of the message
q. size – – total message size, in bytes
r. classification – Anti-spam classification, if AS features are turned on. Values are as follows: 0 – Unclassified, 1 – Legit, 2 - Bulk ,3 – Junk
s. av_classification - this field displays the types of virus found, if one is found: (9 - virus detected by Virus Outbreak Detection, 7 – Clean, 3 – Virus, Null – no attachment) Note: There may be other values, but these have not been detected recently.
t. virus_detected – indicates whether a virus was detected (1=yes, 0=no)

u. original_disposition – Processing action taken after policy evaluation
    i. 0=None
    ii. 1=Drop
    iii. 2=Return
    iv. 4=Admin Quarantined
    v. 8=User Quarantined
    vi. 16-Delivere
    vii. 32=Deferred
    viii. 64=Detained
    ix. 128=Delete
    x. 256=Redirect to Secure Message

v. spam_rating – spam category, based on the type of spam (8=Adult, 16=Scam, 0=Broadcast or Legit)

w. stored – indicates if message is stored on the box or not. Generally, messages will only be stored on the box if they are quarantined. (1=yes, 0=no)

x. mta_appliance_id – ID of the appliance which processed the message

y. business_unit_id – The business unit of the email recipient, if the recipient is internal user, as defined via MailGate groups and business unit association. If not using LDAP recipient verification or named users, value will always equal default business unit.

z. appliance_id – Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster.

16. **Outbound Message Attachment Data per Recipient (outbound_mail_attachment)** - An entry that indicates an attachment was received for an outbound participant in an email. There is one email per attachment per user receiving the attachment. *Note: report will contain data only if "Collect message attachment data" from "Mail Policies" – "Settings"-> "Message Attachment Data" is enabled.*

a. db_id – Unique ID value for this row

b. recipient_db_id – database id of inbound Message Key per Recipient row. This is a foreign key to its parent table

c. attachment_modification indicates if attachment was modified or not, by stripping the attachment or removing a virus

d. attachment_name – the name of the attachment . This would be the file name without the extension

e. attachment_extension – the extension of the attachment

k. attachment_type – the file type of the attachment. The true type of the file is determined by inspecting the content of the attachment to determine if it matches a known file type.

f. attachment_encoded_size - base64 encoded size of the attachment in bytes.

g. attachment_size - size of the attachment in bytes

h. received - Unix timestamp when the messages is received

i. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

17. **Summarized Outbound Policy Execution Data (outbound_policy_execution_summary)** - This log contains summary information on the execution of each defined outbound policy over the time period configured in the MailGate export definition.

a. id – Unique ID value for this row

b. policy_name – of the inbound policy defined in the filter and delivery policies being summarized here

c. policy_phase – policy phase (10=filter, 20=delivery)

d. bucket_time – Unix timestamp when policy was executed

e. count – number of times policy was triggered, by meeting its IF condition, in the processing of email

f. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
g. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

18. Summarized Outbound Policy Tag Data per Sender (outbound_sender_policy_summary) - This log summarized all policy tag values which were set on emails sent by a specific sender over the time period configured in the MailGate export definition.

a. id – Unique ID value for this row
b. sender – sender's email address which is being summarized
c. tag_id – internal ID of the tag. To get the mapping of internal tags to their names, especially for custom tags, this must be retrieves using the local version of remote support option.
d. bucket_time – Unix timestamp for this row
e. count – number of times policy execution triggered the setting of this specific tag on any outgoing email
f. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
g. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

19. Summarized Outbound Sender Domain Data (outbound_sender_domain_summary) - Summary of all email traffic received from a specific device in a domain sent outbound through MailGate over the time period configured in the MailGate export definition.

a. Id – Unique ID value for this row
b. ip – ID address of the internal email sending device being summarized
c. domain – domain of the outbound email sender
d. bucket_time - Unix timestamp for this row
e. virus – Number of emails with viruses received from the device/domain
f. junk – Number of emails classified as junk received from the device/domain
g. bulk – Number of emails classified as bulk mail received from the device/domain
h. legit – Number of emails classified as legitimate received from the device/domain
i. total – Total Number of emails received from the device/domain
j. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
k. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

20. Relay Receiving Connections/Messages Data (receiving_connection_event_log) – Each entry in this log represents one connection created to the MailGate appliance.   Note: This is one of two logs produced by the Relay Receiving Connections/Messages Data export log option

a. id – Unique ID value for this row
b. receiving_conn_id – receiving connection ID. This is a self referencing foreign key to the first entry in this log for the established connection
c. receiving_conn_msg_id – This is a reference to the Policy Engine ID (PE ID) for the email received in this instance of the connection.   This serves as a foreign key to the mail log table.
d. mta_ip –IP address of the machine connecting to the MailGate appliance to send it mail.

e. mta_host_name – hostname of the machine connecting to the MailGate appliance to send it mail. (use REVERSE function, characters are in reverse order)
f. from_email –email address of the sender
g. outbound – if the connection is outbound or not (1=outbound, 0=inbound)
h. tracking_status – Internal use only. Value has no meaning after mail processing is complete.
i. state – Edge Defense failure category. This field is a numeric representation of the failure description in the field below. See additional notes for list of values.
j. failure_description – text description of the connection failure, if the connection failed, otherwise empty. Most types of connection failures are due to Edge Defense features.
k. log_time - Unix timestamp for the row
l. tls_used – indicates if connection is received over TLS or not (1=yes, 0=no)
m. tls_require_option – indicates if TLS was required by MailGate (1-yes, 0=no)
n. throttled – indicates if the connection was throttled due to high mail flow (1-yes, 0=no)
o. msg_priority – message priority (1=high, 2-normal, 3=low (good???)
p. ip_rep_score – IP reputation score. Possible values are between -10 to +10. The lower the number, the more confidence that the IP is a spam site. For example, -10 would be a known span site, while +10 would be an explicitly trusted site.
q. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

21. Relay Receiving Connections/Messages Data (receiving_connection_rcpt_log) – Each entry in this log represents one message received over a connection to the MailGate appliance. There can be multiple messages delivered with a single connection, so this is a child table to the receiving_connection_event_log. Note: This is one of two logs produced by the Relay Receiving Connections/Messages Data export log option

a. Id – Unique ID value for this row
b. receiving_conn_msg_id – messages ID of the message received with this connection foreign key to the inbound or outbound mail log)
c. address – address of the recipient
d. state – Edge Defense failure category. This field is a numeric representation of the failure description in the field below. See additional notes for list of values.
e. failure_description – text description of the connection failure, if the connection failed, otherwise empty. Most types of connection failures are due to Edge Defense features.
f. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

22. Relay Sending Connections/Messages Data (sending_connection_message_log) - Each entry in this log represents one connection created by the MailGate appliance to another mail-capable device. Note: This is one of two logs produced by the Relay Sending Connections/Messages Data export log option

a. id – Unique ID value for this row
b. mta_ip – IP address of the machine being send mail using this connection
c. mta_host_name – – hostname of the machine the MailGate appliance is connecting to in order to send it mail. (use REVERSE function, characters are in reverse order)
d. message_id – This is a reference to the Policy Engine ID (PE ID) for the email sent in this instance of the connection. This serves as a foreign key to the mail log table.
e. recipient_domain – domain of the recipient (use REVERSE function, characters are in reverse order)
f. outbound – if the connection is outbound or not (1=outbound, 0=inbound)
g. state – Edge Defense failure category. This field is a numeric representation of the failure description in the field below. See additional notes for list of values.
h. tls_used – indicates if connection is sent over TLS or not (1=yes, 0=no)

i. tls_verification_used – indicates if the TLS certificate was verified by MailGate (1=yes, 0=no)
j. tls_require_option – indicates if TLS was required by MailGate (1-yes, 0=no)
k. failure_description – text description of the connection failure, if the connection failed, otherwise empty. Most types of connection failures are due to Edge Defense features.
l. log_time - Unix timestamp for the row
m. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

23. Relay Sending Connections/Messages Data (sending_connection_rcpt_log) - Each entry in this log represents one message sent over a connection from the MailGate appliance. There can be multiple messages delivered with a single connection, so this is a child table to the sending_connection_event_log. Note: This is one of two logs produced by the Relay Sending Connections/Messages Data export log option

a. Id – Unique ID value for this row
b. sending_conn_msg_id – message ID of the message received with this connection (foreign key to the inbound or outbound mail log)
c. recipient_address –email address of the recipient
d. state – Edge Defense failure category. This field is a numeric representation of the failure description in the field below. See additional notes for list of values.
e. failure_description – text description of the connection failure, if the connection failed, otherwise empty. Most types of connection failures are due to Edge Defense features.
f. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

24. Summarized Relay Connections Data (mta_connection_summary) – This log summarizes connection information between MailGate and other MTA relay across multiple connections over time, capturing information on reasons for connection rejection summarized over the time period configured in the MailGate export definition.

a. id – Unique ID value for this row
b. mta_ip – IP address of the Mail Transfer Agent connecting to MailGate
c. mta_host_name –hostname of the Mail Transfer Agent connecting to MailGate (use REVERSE function, characters are in reverse order)
d. bucket_time – Unix timestamp for the connection
e. outbound – If the connection is outbound or not (1=outbound, 0=inbound)
f. receiving_established_count – total number of successful connections
g. receiving_dropped_dha_count - number of dropped connections because of suspected DHA (Directory Harvest Attack
h. receiving_dropped_dos_count - number of dropped connections because of suspected DoS (Denial on Service) attack
i. receiving_dropped_dnsbl_count - number of dropped connections because of DNSBL (sender is on a DNS Block Lists)
j. receiving_dropped_manual_count - number of dropped connections because of manual block policy
k. receiving_dropped_ip_reputation_count - number of dropped connections because of IP Reputation of the connecting machine
l. receiving_dropped_batv_count - number of dropped connections because of BATV (Bounce Address Tag Validation)
m. receiving_dropped_dkim_count - number of dropped connections because of DKIM (Domain Keys Identified Mail)
n. receiving_dropped_general_defense_count - number of dropped connections because of General Defense settings
o. receiving_dropped_mail_spoof_count - number of dropped connections because of spoofing

p. receiving_dropped_other_count - number of dropped connections because of unknown reason (dropped from remote side)
q. receiving_dropped_tls_count - number of dropped connections because of failed attempt to connect via TLS
r. receiving_accepted_msg_count – total number of received messages across all connections
s. receiving_rejected_msg_count – number of rejected messages across all connections
t. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
u. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

25. Summarized Inbound/Outbound TLS Data per Domain (mail_domain_message_summary) - This log summarizes connection information between MailGate and other MTA relay across multiple connections over the archive time period in the MailGate configuration. It primarily collects information about usage of TLS, although all connections will be summarized in this table.

a. Id – Unique ID value for this row
b. mta_ip – IP address of the Mail Transfer Agent connecting to MailGate
c. mta_host_name – hostname of the Mail Transfer Agent connecting to MailGate (use REVERSE function, characters are in reverse order)
d. mail_domain – domain name of the Mail Transfer Agent connecting to MailGate (use REVERSE function, characters are in reverse order)
e. bucket_time – Unix timestamp for this row
f. outbound – connection is outbound or not (1=outbound, 0=inbound)
g. receiving_tls_require_option – indicates if TLS was required when receiving from the domain (1=yes, 0=no)
h. sending_tls_require_option – indicates if TLS was required when sending to the domain (1=yes, 0=no)
i. received_msg_count – total number of messages received
j. received_tls_msg_count– number of messages received over TLS
k. receiving_smtp_error_count – number of smtp errors when receiving mail
l. receiving_tls_error_count – number of TLS errors received
m. sent_msg_count – total number of messages sent
n. sent_tls_msg_count – number of messages sent over TLS
o. sending_smtp_error_count – number of smtp errors while sending mail
p. sending_tls_not_offered_count – number of messages not sent because there was no TLS was offered by the other MTA
q. sending_tls_no_response_count – number of messages not sent because TLS connection was not established
r. sending_tls_invalid_cert_count – number of TLS verification failures encountered
s. sending_tls_cert_name_not_match_count – number of TLS failures because of certificate name mismatches
t. sending_tls_error_count – number of errors sending using TLS
u. type – Internal indicator on whether this row needs to not be replicated, replicated in local cluster only or replicated in global cluster
v. appliance_id – the Unique ID of the appliance within the local and remote cluster. This value is generated as part of the license key creation. It is guaranteed to be unique across your MailGate cluster

Additional Notes:

1. The exports associated with Secure Mailboxes are not addressed in the document. They are not used when using MailGate as a pure Email flow security solution, only when using the Secure Messaging feature with allows retrieval of encrypted content directly from the appliance.

2. As mentioned with reach entry, host names and domain names are reversed in their text. For example, “yahoo.com” will be stored as “moc.oohay”, and a host name of “10.20.1.250” will be stored as “052.1.02.01” – To address this, you can use the REVERSE function in a SQL statement after putting the data in a staging table:
    a. Example: REVERSE(recipient_domain) as recipient_domain” clause needs to be used

3. In general, once an export schedule is defined, the summary records produced by the export process cover the time period specified in the export definition. So if an export is set to occur once per day, each record will summarize activity for the previous 24 hours, and the next days export will cover the next day’s 24 hours of data. If the exports are hourly, each export will cover only an hour’s worth of data, as shown on the screenshot below –

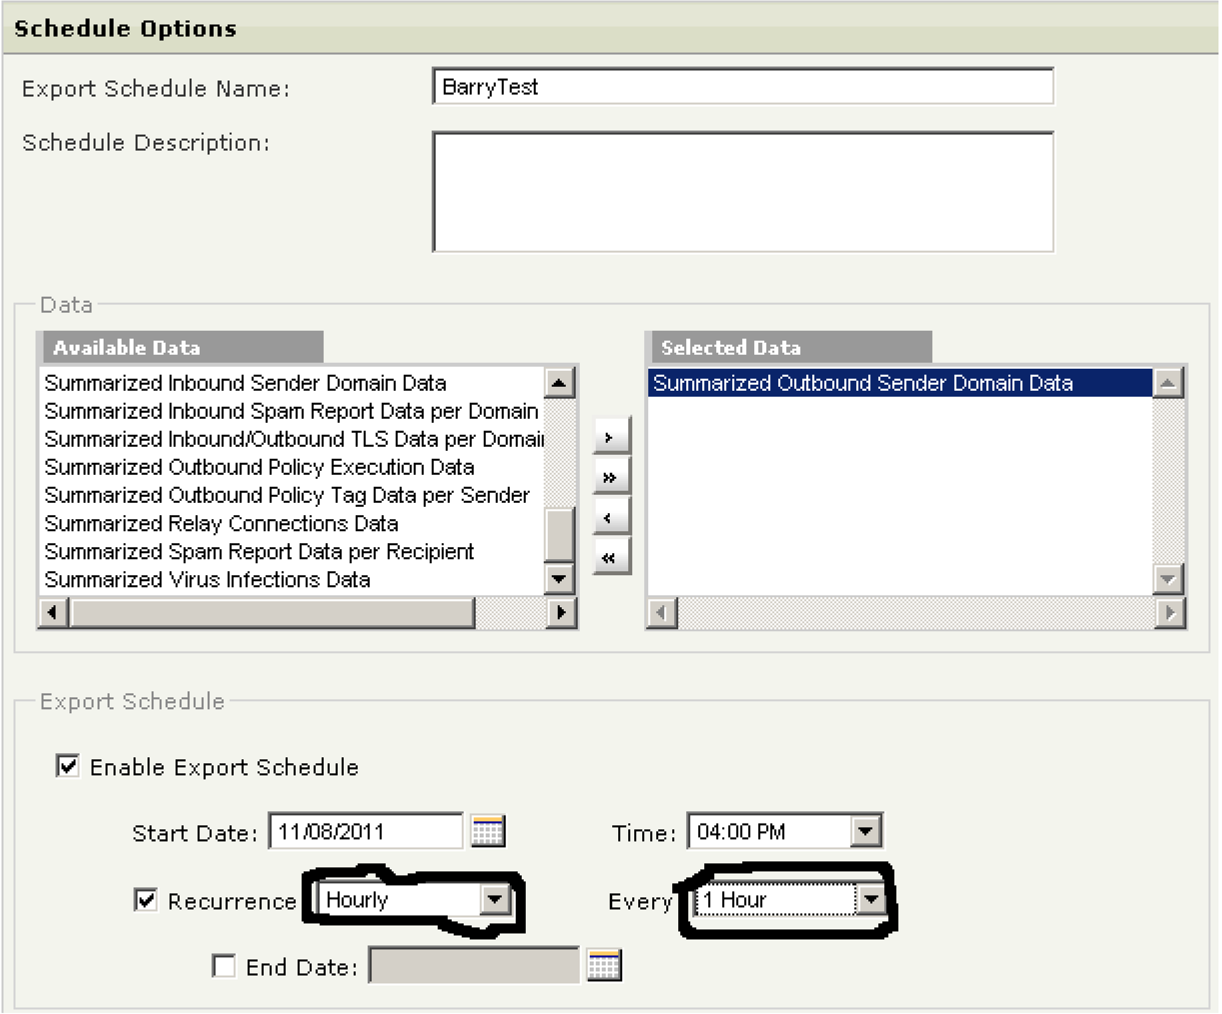

Schedule Options

Export Schedule Name: BarryTest

Schedule Description:

Data

Available Data:
- Summarized Inbound Sender Domain Data
- Summarized Inbound Spam Report Data per Domain
- Summarized Inbound/Outbound TLS Data per Domai
- Summarized Outbound Policy Execution Data
- Summarized Outbound Policy Tag Data per Sender
- Summarized Relay Connections Data
- Summarized Spam Report Data per Recipient
- Summarized Virus Infections Data

Selected Data:
- Summarized Outbound Sender Domain Data

Export Schedule

☑ Enable Export Schedule

Start Date: 11/08/2011 Time: 04:00 PM

☑ Recurrence Hourly Every 1 Hour

☐ End Date:

4. The only exception to the time period described in the last entry is the first export after the scheduled export is defined. This export will cover all data collected and retained up to that point in time. For example, given the setting in the screen below, the first archive would represent 7 days of data.

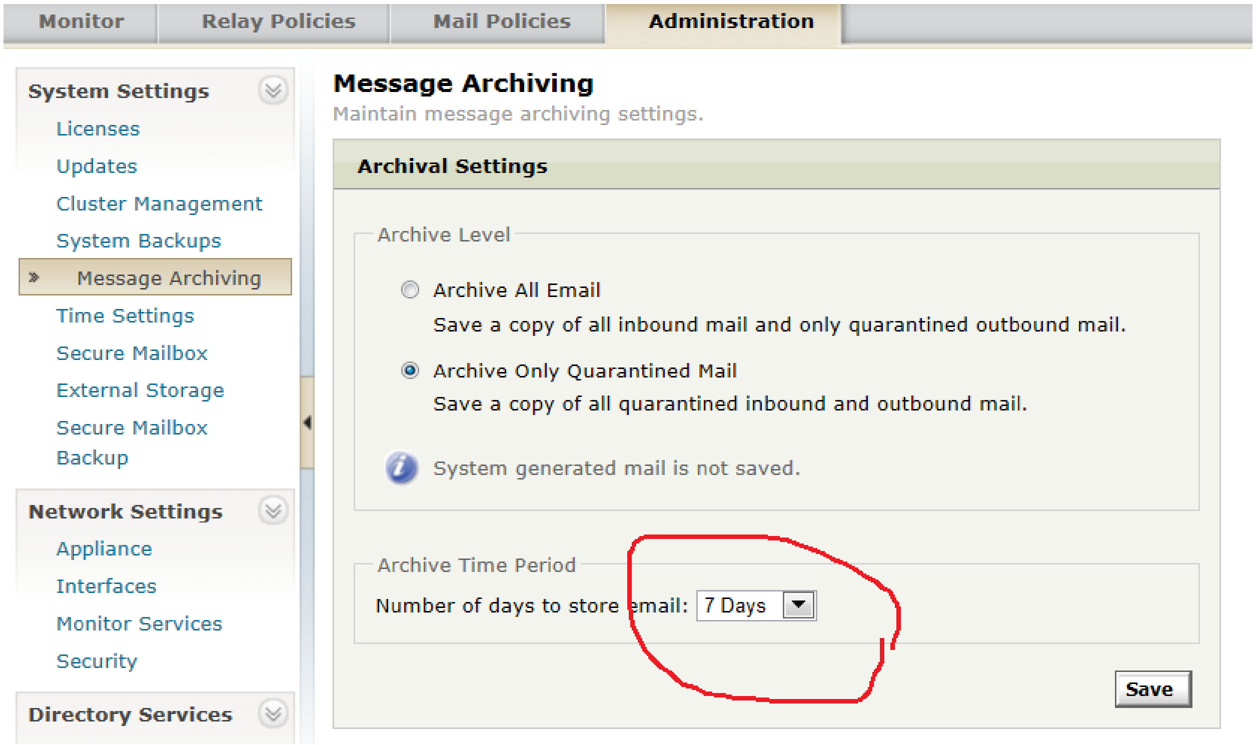

Message Archiving

Maintain message archiving settings.

Archival Settings

Archive Level

○ Archive All Email
Save a copy of all inbound mail and only quarantined outbound mail.

● Archive Only Quarantined Mail
Save a copy of all quarantined inbound and outbound mail.

System generated mail is not saved.

Archive Time Period

Number of days to store email: 7 Days

5. List of valid values for the State field in the Connection Log exports

a. STATE_SUCCEEDED = "0";
b. STATE_DROP_DHA = "201";
c. STATE_DROP_DOS = "202";
d. STATE_DROP_DOA = "203";
e. STATE_DROP_DNSBL = "204";
f. STATE_DROP_GENERAL_DEFENSE = "205";
g. STATE_DROP_MANUAL = "206";
h. STATE_DROP_SPOOF = "207";
i. STATE_DROP_TIMEOUT = "208";
j. STATE_DROP_LOST_CONNECTION = "209";
k. STATE_DROP_LIMIT_EXCEEDED = "210";
l. STATE_DROP_CRATE_LIMIT = "211";
m. STATE_DROP_MSGS_LIMIT = "212";
n. STATE_DROP_NON_SMTP_COMMAND = "213";
o. STATE_DROP_MSG_SIZE_EXCEEDED = "214";
p. STATE_DROP_IP_REPUTATION_BLOCK = "215";
q. STATE_DROP_IP_REPUTATION_THROTTLE = "216";
r. STATE_DROP_POLICY_BLOCK = "217";
s. STATE_DROP_POLICY_THROTTLE = "218";
t. STATE_DROP_UNKNOWN = "220";
u. STATE_INTERNAL_ERROR = "221";
v. STATE_TLS_HANDSHAKE_ERROR = "250";
w. STATE_TLS_NOT_AVAILABLE = "251";
x. STATE_TLS_MUST_ISSUE_STARTTLS = "252";
y. STATE_MESSAGE_SUCCEEDED = "300";

z. STATE_MESSAGE_DROP_TIMEOUT = "308";

aa. STATE_MESSAGE_DROP_LOST_CONNECTION = "309";

bb. STATE_MESSAGE_DROP_LIMIT_EXCEEDED = "310";

cc. STATE_MESSAGE_DROP_CRATE_LIMIT = "311";

dd. STATE_MESSAGE_DROP_MSGS_LIMIT = "312";

ee. STATE_MESSAGE_DROP_NON_SMTP_COMMAND = "313";

ff. STATE_MESSAGE_DROP_MSG_SIZE_EXCEEDED = "314";

gg. STATE_MESSAGE_DROP_UNKNOWN = "320";

hh. STATE_MESSAGE_INTERNAL_ERROR = "321";

ii. STATE_MESSAGE_TOO_LARGE = "322";

jj. STATE_MESSAGE_TOO_MANY_HOPS = "323";

kk. STATE_MESSAGE_POLICY_ERROR = "324";

ll. STATE_MESSAGE_QFILE_WRITE_ERROR = "325";

mm. STATE_MESSAGE_SOFTWARE_ERROR = "326";

nn. STATE_MESSAGE_SPOOFED_MAILFROM = "327";

oo. STATE_MESSAGE_INCOMPLETE = "349";

pp. STATE_MESSAGE_TLS_HANDSHAKE_ERROR = "350";

qq. STATE_MESSAGE_TLS_NOT_AVAILABLE = "351";

rr. STATE_MESSAGE_TLS_MUST_ISSUE_STARTTLS = "352";

ss. STATE_RCPT_SUCCEEDED = "0";

tt. STATE_RCPT_NOT_INITIALIZED = "0";

uu. STATE_RCPT_TOO_MANY_RECIPIENTS = "370";

vv. STATE_RCPT_GENERIC_ERROR = "371";