



Monitoring Postfix Mail Server

eG Innovations Product Documentation

About eG Innovations

eG Innovations provides intelligent performance management solutions that automate and dramatically accelerate the discovery, diagnosis, and resolution of IT performance issues in on-premises, cloud and hybrid environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, from virtualization to cloud, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

To contact eG Innovations sales team, email sales@eginnovations.com.

Copyright © 2018 eG Innovations Inc. All rights reserved.

This document may not be reproduced by any means nor modified, decompiled, disassembled, published or distributed, in whole or in part, or translated to any electronic medium or other means without the prior written consent of eG Innovations. eG Innovations makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The information contained in this document is subject to change without notice.

All right, title, and interest in and to the software and documentation are and shall remain the exclusive property of eG Innovations. All trademarks, marked and not marked, are the property of their respective owners. Specifications subject to change without notice.

Table of Contents

ABOUT EG INNOVATIONS	2
CHAPTER 1: MONITORING THE POSTFIX MAIL SERVER	1
1.1 How does eG Enterprise Monitor the Postfix mail server?	3
1.2 The Postfix Queues layer	3
1.2.1 Active Queue Test	4
1.2.2 Deferred Queue Test	8
1.2.3 Hold Queue Test	11
1.2.4 Incoming Queue Test	15
1.2.5 MailDrop Queue Test	18
CHAPTER 2: CONCLUSION	22

Table of Figures

Figure 1.1: The Postfix mail server architecture	1
Figure 1.2: The layer model of Postfix	2
Figure 1.3: The Postfix Queues layer	3

Chapter 1: Monitoring the Postfix mail server

Postfix is a free and open-source mail transfer agent (MTA) that routes and delivers electronic mail to various destinations. Postfix consists of a combination of server programs that run in the background, and client programs that are invoked by user programs or by system administrators (See Figure 1).

The Postfix core consists of several dozen server programs that run in the background, each handling one specific aspect of email delivery. Examples are the SMTP server, the scheduler, the address rewriter, and the local delivery server. For damage-control purposes, most server programs run with fixed reduced privileges, and terminate voluntarily after processing a limited number of requests. To conserve system resources, most server programs terminate when they become idle.

Client programs run outside the Postfix core. They interact with Postfix server programs through mail delivery instructions in the user's `~/.forward` file, and through small "gate" programs to submit mail or to request queue status information.

Other programs provide administrative support to start or stop Postfix, query status information, manipulate the queue, or to examine or update its configuration files.

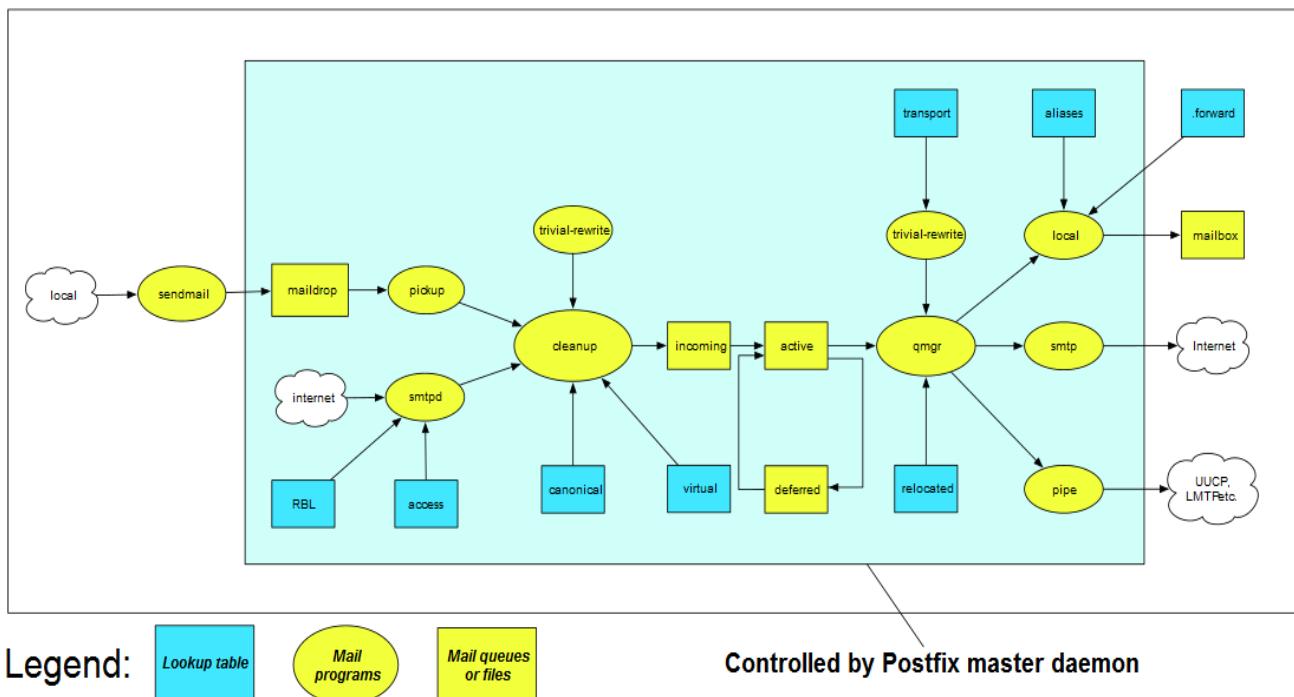


Figure 1.1: The Postfix mail server architecture

Postfix manages pipelines of processes that pass the responsibility for message delivery and error notification from one process to the next. All message and notification "state" information is persisted in the file system. The processes in a pipeline operate mostly without centralized control; this relative autonomy simplifies error recovery. When a process fails before completing its part of a file or protocol transaction, its predecessor in the pipeline backs off and retries the request later, and its successor in the pipeline discards unfinished work. This approach makes Postfix highly resilient, as long as the operating system or hardware don't fail catastrophically.

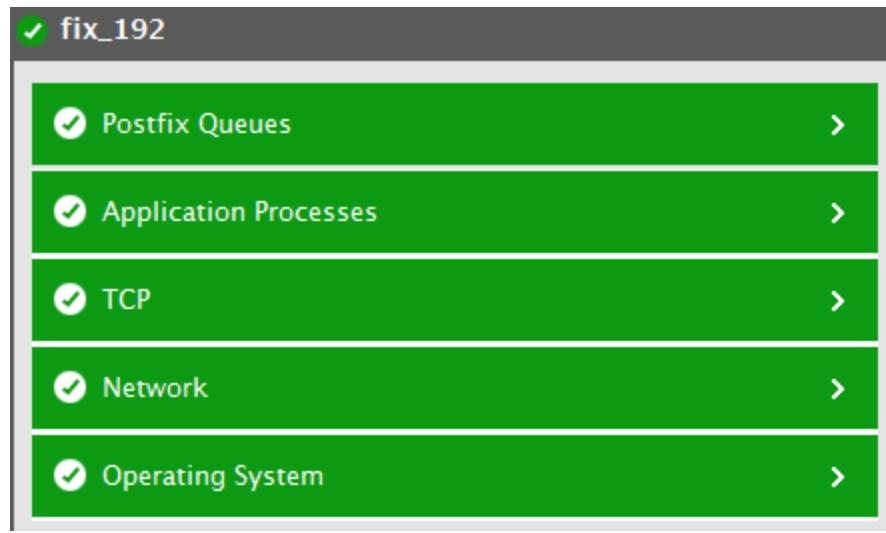


Figure 1.2: The layer model of Postfix

Each layer of Figure 1.2 is mapped to a variety of tests each of which report a wealth of metrics related to the queues of the Postfix mail server. Using these metrics administrators can find quick and accurate answers to the following queries:

- What is the total size of the active queue?
- In which time duration, the number of messages in the active queue were the highest - is it less than 5 minutes? between 5 and 10 minutes? between 10 and 20 minutes? between 20 and 40 minutes? between 40 and 80 minutes? between 80 and 160 minutes? between 160 and 320 minutes? between 320 and 640 minutes? between 640 and 1280 minutes? or is it above 1280 minutes?
- What is the total size of the incoming queue?
- In which time duration, the number of messages in the incoming queue were the highest - is it less than 5 minutes? between 5 and 10 minutes? between 10 and 20 minutes? between 20 and 40 minutes? between 40 and 80 minutes? between 80 and 160 minutes? between 160 and 320

minutes? between 320 and 640 minutes? between 640 and 1280 minutes? or is it above 1280 minutes?

- What is the total size of the hold/deferred/maildrop queues?
- In which time duration the number of messages in the hold/deferred/maildrop queues were the highest - is it less than 5 minutes? between 5 and 10 minutes? between 10 and 20 minutes? between 20 and 40 minutes? between 40 and 80 minutes? between 80 and 160 minutes? between 160 and 320 minutes? between 320 and 640 minutes? between 640 and 1280 minutes? or is it above 1280 minutes?

1.1 How does eG Enterprise Monitor the Postfix mail server?

eG Enterprise employs both *agentless* and *agent based* approach to monitor the target Postfix mail server.

To collect the metrics of interest from the Postfix mail server, it is essential to provide the credentials of a user who is authorized to access the Postfix mail server i.e., the eG installed user should be able to communicate with the Postfix mail server and collect the required metrics.

Since the **Operating System**, **Network**, **TCP** and **Application Processes** layers are already discussed in the *Monitoring Unix and Windows servers* document, let us now focus on the Postfix Queues layer alone in the forthcoming section.

1.2 The Postfix Queues layer

This layer tracks the messages in various queues of the Postfix mail server and reports the total size of each queue and the number of messages that were in the queue for a specified time period.



Figure 1.3: The Postfix Queues layer

1.2.1 Active Queue Test

The queue manager is a delivery agent scheduler which ensures fast and fair delivery of mail to all destinations within designated resource limits. By default, the active queue accommodates a maximum of 20000 messages. The messages to the active queue are from the incoming queue where new mails arrive and the deferred queue where old mails are retried. The messages in the active queue are ready to be sent, but are not necessarily in the process of being sent. In reality, the active queue is a set of data structures within the memory of the queue manager process. The envelope information for messages in the active queue is managed in memory, allowing the queue manager to do global scheduling, allocating available delivery agent processes to an appropriate message in the active queue. Within the active queue, the messages are broken up into groups of recipients that share the same transport/nexthop combination; the group size is capped by the transport's recipient concurrency limit. Multiple recipient groups (from one or more messages) are queued for delivery grouped by transport/nexthop combination. An active queue will be empty only when a queue manager marks a destination as dead (this may be due to the destination being down) and defers all the mails assigned to that destination to the deferred queue. An active queue may be overflowing if the destination is slow or if there is a problem in the destination which has caused excessive arrival of mails leading to a sudden surge in the active queue as well as congesting the destination. A destination may also be congested if the entire deferred queue undergoes an unwarranted flush. When messages reside in the active queue for a prolonged period, then the active queue may grow indefinitely and may not hold the space to accommodate all the messages. This will lead to a drastic degradation of the performance of the Postfix mail server. To avoid such unpleasant degradation and to keep track of the performance of the Postfix mail server, it is necessary to continuously track the messages in the active queue. The Active Queue test helps administrators in this regard!

This test periodically monitors the active queue of the target Postfix mail server and reports the total size of the queue as well as the split up of the message count in terms of time duration i.e., the number of messages that were in the queue for a specified time duration.

Target of the test : A Postfix mail server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target Postfix mail server being monitored

Configurable parameters for the test

Parameters	Description
Test period	How often should the test be executed

Parameters	Description
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is <i>NULL</i> .
UseSUDO	By default, this flag is set to False indicating that the test does not collect the queue related statistics, by default. If this flag is set to True , then the test uses the sudo command to collect the queue related statistics.
Timeout	Specify the duration (in seconds) beyond which this test should time out in the TIMEOUT text box. The default is 30 seconds.
Detailed Diagnosis	<p>To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measures made by the test

Measurement	Description	Measurement Unit	Interpretation
Queue size	Indicates the total size of the queue.	Number	<p>Ideally, a low value is desired for this measure.</p> <p>An unusually high number of messages in the queue is indicative of any problem with the corresponding queue or its end points.</p> <p>The Detailed diagnosis of this measure lists the message count</p>

Measurement	Description	Measurement Unit	Interpretation
			for each domain to which the mails are destined to.
Less than 5 mins	Indicates the number of messages that were in the queue for less than 5 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 5-10 mins	Indicates the number of messages that were in the queue for a time duration between 5 to 10 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 10- 20 mins	Indicates the number of messages that were in the queue for a time duration between 10 to 20 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 20- 40 mins	Indicates the number of messages that were in the queue for a time duration between 20 to 40 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 40- 80 mins	Indicates the number of messages that were in the queue for a time duration between 40 to 80 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 80- 160 mins	Indicates the number of messages that were in the queue for a time duration between 80 to 160 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 160- 320	Indicates the number of	Number	Under normal conditions, the active

Measurement	Description	Measurement Unit	Interpretation
mins	messages that were in the queue for a time duration between 160 to 320 minutes.		queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 320- 640 mins	Indicates the number of messages that were in the queue for a time duration between 320 to 640 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
Between 640-1280 mins	Indicates the number of messages that were in the queue for a time duration between 640 to 1280 minutes.	Number	Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.
More than 1280 mins	Indicates the number of messages that were in the queue for more than 1280 minutes.	Number	<p>Under normal conditions, the active queue should be empty i.e., the mail should leave the system to the destination as quickly as it comes in.</p> <p>An unusually high value of this measure may indicate that the active queue is congested. Administrators should therefore investigate the real reason for congestion and should take necessary steps to send the mails to the deferred queue or should set a rule intimating the users to try sending mails to the problematic domain later. Some of the possible reasons for congestion may be either the unavailability of the domain to which the mails are destined to or the domain may be slow.</p>

1.2.2 Deferred Queue Test

If a message is delivered to most of the deliverable recipients and for some recipients delivery failed due to a transient reason (may succeed in delivery later), then such messages are placed in the deferred queue. The queue manager scans the deferred queue periodically and during each scan, a fraction of the deferred queue is brought back to the active queue for retry. Each message in the deferred queue will have a cool-off time limit set beyond which the message will be retried for delivery. One of the common causes of large deferred queues is the failure to validate recipients at the SMTP input stage. This is due to spammers routinely launching dictionary attacks from unreliable sender addresses following which the invalid recipient addresses bounce and clog the deferred queue. Therefore, recipient validation is strongly recommended. Another common cause of congestion is unwarranted flushing of the entire deferred queue. The deferred queue holds messages that are likely to fail to be delivered and are also likely to be slow to fail delivery (i.e., time out). As a result the most common reaction to a large deferred queue is to flush out the deferred queue which may ease congestion to an extent. The deferred queue should not be flushed until and unless most of its content has recently become deliverable (e.g. relayhost back up after an outage)! If the deferred queue grows endlessly, then the messages will often be retried for delivery which may sometimes flood the active queue and cause a brief congestion of the queues. To avoid this, administrators should continuously monitor the deferred queue and figure out at what time the messages in the deferred queue started increasing manifold. Administrators should also identify the domain to which most of the messages failed to be delivered so that legitimacy of that domain can also be examined. To help administrators in these tasks, eG Enterprise Suite provides you with the **Deferred Queue** test.

This test periodically monitors the deferred queue of the target Postfix mail server and reports the total size of the deferred queue as well as the split up of the message count in terms of time duration i.e., the number of messages that were in the deferred queue for a specified time duration.

Target of the test : A Postfix mail server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target Postfix mail server being monitored

Configurable parameters for the test

Parameters	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.

Parameters	Description
Port	The port at which the specified host listens. By default, this is <i>NULL</i> .
UseSUDO	By default, this flag is set to False indicating that the test does not collect the queue related statistics, by default. If this flag is set to True , then the test uses the sudo command to collect the queue related statistics.
Timeout	Specify the duration (in seconds) beyond which this test should time out in the TIMEOUT text box. The default is 30 seconds.
Detailed Diagnosis	<p>To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measures made by the test

Measurement	Description	Measurement Unit	Interpretation
Queue size	Indicates the total size of the queue.	Number	<p>An unusually high number of messages in the queue is indicative of any problem with the corresponding queue or its end points.</p> <p>The Detailed diagnosis of this measure lists the message count for each domain.</p>
Less than 5 mins	Indicates the number of messages that were in the queue for less than	Number	

Measurement	Description	Measurement Unit	Interpretation
	5 minutes.		
Between 5-10 mins	Indicates the number of messages that were in the queue for a time duration between 5 to 10 minutes.	Number	
Between 10- 20 mins	Indicates the number of messages that were in the queue for a time duration between 10 to 20 minutes.	Number	
Between 20- 40 mins:	Indicates the number of messages that were in the queue for a time duration between 20 to 40 minutes.	Number	
Between 40- 80 mins	Indicates the number of messages that were in the queue for a time duration between 40 to 80 minutes.	Number	
Between 80- 160 mins	Indicates the number of messages that were in the queue for a time duration between 80 to 160 minutes.	Number	
Between 160- 320 mins:	Indicates the number of messages that were in the queue for a time duration between 160 to 320 minutes.	Number	
Between 320- 640 mins	Indicates the number of messages that were in	Number	

Measurement	Description	Measurement Unit	Interpretation
	the queue for a time duration between 320 to 640 minutes.		
Between 640-1280 mins	Indicates the number of messages that were in the queue for a time duration between 640 to 1280 minutes.	Number	
More than 1280 mins	Indicates the number of messages that were in the queue for more than 1280 minutes.	Number	<p>A low value is desired for this measure.</p> <p>When a host with lots of deferred mail is down for some time, it is possible for the entire deferred queue to reach its retry time simultaneously. This can lead to a very full active queue once the host comes back up. The phenomenon can repeat approximately every <i>maximal_backoff_time</i> seconds if the messages are again deferred after a brief burst of congestion. Since the messages are retired constantly, it is important for the administrators to keep a constant vigil on the value of this measure. If this measure is at a high always, then the messages will always be retried leading to congestion.</p>

1.2.3 Hold Queue Test

Whenever a Postfix mail server detects spam messages/malware, then those messages are automatically placed in the hold queue. By default, the administrator defines *smtpd* access policies, or cleanup header/body checks that may cause messages to be automatically diverted from normal processing and be placed indefinitely in the hold queue. The messages remain in the hold queue

until the administrator intervenes. The messages in the hold queue are barred from periodic delivery attempts. Messages can stay in the hold queue longer than the maximum lifetime specified for the hold queue. The old messages in the hold queue can be released only by moving those messages into the maildrop queue. Once the messages are moved into the maildrop queue, the messages will get a new timestamp and will be attempted delivery. The messages that are relatively new in the hold queue can be moved directly to the deferred queue. If the hold queue consists of too many messages, then the hold queue may grow indefinitely and at some point, the queue may not be able to accommodate new messages resulting in a bottleneck of the queue. To avoid such bottlenecks and figure out the exact time duration for which messages are in the hold queue, it is necessary for the administrator to periodically monitor the hold queue. The **Hold Queue** test helps administrators to identify the time duration from which the hold queue was bombarded with messages.

This test periodically monitors the hold queue of the target Postfix mail server and reports the total size of the queue as well as the split up of the message count in terms of time duration i.e., the number of messages that were in the queue for a specified time duration.

Target of the test : A Postfix mail server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target Postfix mail server being monitored

Configurable parameters for the test

Parameters	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is <i>NULL</i> .
UseSUDO	By default, this flag is set to False indicating that the test does not collect the queue related statistics, by default. If this flag is set to True , then the test uses the sudo command to collect the queue related statistics.
Timeout	Specify the duration (in seconds) beyond which this test should time out in the TIMEOUT text box. The default is 30 seconds.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of

Parameters	Description
	<p>this test for a particular server, choose the On option. To disable the capability, click on the Off option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measures made by the test

Measurement	Description	Measurement Unit	Interpretation
Queue size	Indicates the total size of the queue.	Number	<p>An unusually high number of messages in the queue is indicative of any problem with the corresponding queue or its end points.</p> <p>The Detailed diagnosis of this measure lists the message count for each domain.</p>
Less than 5 mins	Indicates the number of messages that were in the queue for less than 5 minutes.	Number	
Between 5-10 mins	Indicates the number of messages that were in the queue for a time duration between 5 to 10 minutes.	Number	
Between 10- 20 mins	Indicates the number of messages that were in the queue for a time duration between 10 to 20 minutes.	Number	

Measurement	Description	Measurement Unit	Interpretation
Between 20- 40 mins	Indicates the number of messages that were in the queue for a time duration between 20 to 40 minutes.	Number	
Between 40- 80 mins	Indicates the number of messages that were in the queue for a time duration between 40 to 80 minutes.	Number	
Between 80- 160 mins:	Indicates the number of messages that were in the queue for a time duration between 80 to 160 minutes.	Number	
Between 160- 320 mins	Indicates the number of messages that were in the queue for a time duration between 160 to 320 minutes.	Number	
Between 320- 640 mins	Indicates the number of messages that were in the queue for a time duration between 320 to 640 minutes.	Number	
Between 640-1280 mins	Indicates the number of messages that were in the queue for a time duration between 640 to 1280 minutes.	Number	
More than 1280 mins	Indicates the number of messages that were in the queue for more than 1280 minutes.	Number	A high value for this measure is a cause of concern as this may be an indication of spam messages and malware which may have been filtered out.

1.2.4 Incoming Queue Test

The cleanup service of the Postfix mail server writes all new mails into the incoming queue. Under normal conditions the incoming queue is nearly empty (has only mode 0600 files), with the queue manager able to import new messages into the active queue as soon as they become available. The incoming queue grows rapidly when the message input rate spikes above the rate at which the queue manager can import messages into the active queue. The main factors that are slowing down the queue manager are disk I/O and lookup queries to the trivial-rewrite service. If the incoming queue grows abnormally and is not emptied, then the new messages from the maildrop queue cannot be accommodated. In such cases, administrators may not know when exactly the incoming queue started growing and to which particular domain those messages were to be delivered. To have a fair knowledge on these issues, administrators can use the **Incoming Queue** test.

This test periodically monitors the incoming queue of the target Postfix mail server and reports the total size of the queue as well as the split up of the message count in terms of time duration i.e., the number of messages that were in the queue for a specified time duration.

Target of the test : A Postfix mail server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target Postfix mail server being monitored

Configurable parameters for the test

Parameters	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is <i>NULL</i> .
UseSUDO	By default, this flag is set to False indicating that the test does not collect the queue related statistics, by default. If this flag is set to True , then the test uses the sudo command to collect the queue related statistics.
Timeout	Specify the duration (in seconds) beyond which this test should time out in the TIMEOUT text box. The default is 30 seconds.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when

Parameters	Description
	<p>specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measures made by the test:

Measurement	Description	Measurement Unit	Interpretation
Queue size	Indicates the total size of the queue.	Number	<p>An unusually high number of messages in the queue is indicative of any problem with the corresponding queue or its end points.</p> <p>The Detailed diagnosis of this measure lists the message count for each domain.</p>
Less than 5 mins	Indicates the number of messages that were in the queue for less than 5 minutes.	Number	
Between 5-10 mins	Indicates the number of messages that were in the queue for a time duration between 5 to 10 minutes.	Number	
Between 10- 20 mins	Indicates the number of messages that were in the queue for a time duration between 10 to 20 minutes.	Number	

Measurement	Description	Measurement Unit	Interpretation
	20 minutes.		
Between 20- 40 mins	Indicates the number of messages that were in the queue for a time duration between 20 to 40 minutes.	Number	
Between 40- 80 mins	Indicates the number of messages that were in the queue for a time duration between 40 to 80 minutes.	Number	
Between 80- 160 mins	Indicates the number of messages that were in the queue for a time duration between 80 to 160 minutes.	Number	
Between 160- 320 mins	Indicates the number of messages that were in the queue for a time duration between 160 to 320 minutes.	Number	
Between 320- 640 mins	Indicates the number of messages that were in the queue for a time duration between 320 to 640 minutes.	Number	
Between 640-1280 mins	Indicates the number of messages that were in the queue for a time duration between 640 to 1280 minutes.	Number	
More than 1280 mins	Indicates the number of messages that were in	Number	A high value for this measure is a cause of concern as it may indicate

Measurement	Description	Measurement Unit	Interpretation
	the queue for more than 1280 minutes.		an impending problem with the domain to which the messages are destined to. The domain may either be too slow to accommodate all the messages or the domain may be down for a long period of time.

1.2.5 MailDrop Queue Test

Messages that have been submitted via the Postfix sendmail command, but not yet brought into the main Postfix queue by the pickup service, await processing in the maildrop queue. Messages can be added to the maildrop queue even when the Postfix system is not running. These messages will be processed only after the Postfix is started. All mails that enter the Postfix queue have to pass through the cleanup service which is responsible for envelope and header rewriting, header and body regular expression checks, automatic bcc recipient processing, milter content processing, and reliable insertion of the message into the incoming queue.

Since the pickup service is single threaded, only one message can be delivered at a time at a rate that does not exceed the disk I/O latency of the cleanup service. If the cleanup service consumes excessive CPU or if the queue is nearly full or if the local message submission rate is high in the maildrop queue, then it indicates congestion of the queue. Also, if the active queue is full, the pickup service may slow down the message injection to the active queue which may also lead to the congestion of the maildrop queue. If the maildrop queue is congested for a longer duration, then new mails may not be accommodated to the queue resulting in slowdown in the reception of mails by the Postfix server. To avoid frequent congestion of the maildrop queue and to figure out the domain to which a large number of messages have arrived in the maildrop queue, administrators can use the **Maildrop Queue** test.

This test periodically monitors the maildrop queue of the target Postfix mail server and reports the total size of the queue as well as the split up of the message count in terms of time duration i.e., the number of messages that were in the queue for a specified time duration.

Target of the test : A Postfix mail server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target Postfix mail server being monitored

Configurable parameters for the test

Parameters	Description
Test period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the specified host listens. By default, this is <i>NULL</i> .
UseSUDO	By default, this flag is set to False indicating that the test does not collect the queue related statistics, by default. If this flag is set to True , then the test uses the sudo command to collect the queue related statistics.
Timeout	Specify the duration (in seconds) beyond which this test should time out in the TIMEOUT text box. The default is 30 seconds.
Detailed Diagnosis	<p>To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option.</p> <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measures made by the test

Measurement	Description	Measurement Unit	Interpretation
Queue size	Indicates the total size of the queue.	Number	<p>An unusually high number of messages in the queue is indicative of any problem with the corresponding queue or its end points.</p> <p>The Detailed diagnosis of this measure lists the message count for each domain.</p>

Measurement	Description	Measurement Unit	Interpretation
Less than 5 mins	Indicates the number of messages that were in the queue for less than 5 minutes.	Number	
Between 5-10 mins	Indicates the number of messages that were in the queue for a time duration between 5 to 10 minutes.	Number	
Between 10-20 mins	Indicates the number of messages that were in the queue for a time duration between 10 to 20 minutes.	Number	
Between 20-40 mins	Indicates the number of messages that were in the queue for a time duration between 20 to 40 minutes.	Number	
Between 40-80 mins	Indicates the number of messages that were in the queue for a time duration between 40 to 80 minutes.	Number	
Between 80-160 mins	Indicates the number of messages that were in the queue for a time duration between 80 to 160 minutes.	Number	
Between 160-320 mins:	Indicates the number of messages that were in the queue for a time duration between 160 to 320 minutes.	Number	
Between 320-640 mins	Indicates the number of messages that were in the queue for a time duration between 320 to 640 minutes.	Number	
Between 640-1280	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
mins	messages that were in the queue for a time duration between 640 to 1280 minutes.		
More than 1280 mins	Indicates the number of messages that were in the queue for more than 1280 minutes.	Number	

Chapter 2: Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to the **Postfix** mail servers. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact support@eginnovations.com. We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to feedback@eginnovations.com.