Sun Java™ Enterprise System Technical Note: Sizing Disks for MTA Messaqe Queues

2005Q1

Part Number 819-2313-10

The *Sun Java Enterprise System 2005Q1 Technical Note: Sizing Disks for MTA Messaqe Queues* describes how to plan for Sun Java™ System Messaging Server 6 2005Q1 MTA network connectivity issues and how to tune the MTA for reattempts of mail delivery.

The component product affected by this technical note is:

- Sun Java System Messaging Server 6.x

This technical note contain the following sections:

Technical Note Revision History

Table 1 Revision History

Date	Description of Changes
March 4, 2005	Initial release of this technical note.

Disk Sizing for MTA Message Queues

Introduction

The behavior of the Messaging Server MTA Queue is to provide a transient store for messages waiting to be delivered. Messages are written to disk in a persistent manner to maintain guaranteed service delivery. If the MTA is unable to deliver the message, it will retry until it finally gives up and returns the message to the sender.

Message Queue Performance

Sizing the MTA Queue disks are an important step for improving MTA performance. The MTA's performance is directly tied to disk I/O first above any other system resource. This means that you should plan on disk volume that consists of multiple disk spindles, which are concatenated and stripped by using a disk RAID system.

End users are quickly affected by the MTA performance. As users press the SEND button on their email client, the MTA will not fully accept receipt of the message until the message has been committed to the Message Queue. Therefore, improved performance on the Message Queue results in better response times for the end-user experience.

Message Queue Availability

SMTP services are considered a guaranteed message delivery service. This is an assurance to end users that the messaging server will not lose messages that the service is attempting to deliver. When you architect the design of the MTA Queue system, all effort should be made to ensure that messages will not be lost. This guarantee is usually made by implementing redundant disk systems through various RAID technologies.

Message Queue Available Disk Sizing

The queue will grow excessively if one of the following conditions occurs:

- The site has excessive network connectivity issues
- The MTA configuration is holding on to messages too long
- There are valid problems with those messages (not covered in this document)

The following sections address these issues.

Planning for Network Connectivity Issues

Occasionally the MTA is unable to deliver messages due to network connectivity issues. In these cases, the messages will be stored on the queue until the next time the MTA is able to attempt to deliver (as defined by the retry interval).

Planning on disk space for these outages is based on a simple rule, the "General Rule for Message Queue Sizing:"

1. Determine average number of messages / minute expected to be delivered (N)
2. Determine average size (kb) of messages (S)
3. Determine maximum duration (minutes) of typical network connectivity outages (T)

Thus, the formula for estimating the Disk Queue Size is:

```
Disk Queue Size (kb) = N x S x T
```

Tuning MTA for Reattempts of Delivery

Occasionally, the system will not be able to deliver any messages. In this state, messages will reside on the message queue while the MTA attempts to set aside the messages for a period of time (retry interval) until it reattempts the delivery. This will continue until the MTA gives up and returns the message to the sender. The reason a message is undeliverable is fairly unpredictable. A number of reasons such as network connectivity, busy destination server, network throttles, and so on, could explain why the message is undeliverable.

On a busy server, these temporarily stored messages can build up during periods of high volume activities. Such a build-up can potentially cause problems with disk space. To avoid these build-ups, tune the MTA to retry delivery at a faster rate.

The retry interval is set within the Channel Block configurations of the `imta.cnf` file. The structure of this file consists of two parts: rewrite rules and channel blocks. The channel blocks define the behavior of a particular disk queue and related processes. This document refers to the `tcp_local` channel. The `tcp_local` channel provides delivery to sites outside an enterprise's local network, in other words, to places over the Internet.

The retry interval setting of the `tcp_local` channel is initially set by the default channel block. The default channel block allows settings to be duplicated to avoid having repeated settings.

The following is the default channel block:

```
defaults notices 1 2 4 7 copywarnpost copysendpost postheadonly noswitchchannel
immnonurgent maxjobs 7 defaulthost red.siroe.com red.siroe.com
```

First, the structure of the channel block consists of the channel name. In the example above, this is the default channel block, which will be applied to channels without these settings. The second part is a list of channel keywords.

The `notices` keyword specifies the amount of time that can elapse before message delivery notices (MDNs) are sent back to the sender. This keyword starts with the `notices` keyword followed by a set of numbers, which set the retry period. By default, the MTA will attempt delivery and send notices back to the sender. These notices come from “postmaster” to end-user inboxes.

In this example, the MTA will retry at a period of 1 day, 2 days, and 4 days. At 7 days, the MTA will return the message and regard the message as a failed delivery.

In many cases, the default setting of the MTA provides adequate performance. In some cases, you need to tune the MTA to avoid potential resource exhaustions, such as running out disk space for message queues. This is not a product limitation, but a limitation of the total Messaging Server system, which includes hardware and network resources.

In consideration of these possible disk size issues, deployments with a large number of users may not want to attempt message deliveries for much shorter intervals. If this is the case, study the documentation listed below.

Further Readings

Refer to the following documentation for more information.

- “To Set Notification Message Delivery Intervals,” in the *Sun Java System Messaging Server 6 2005Q1 Administration Guide*:

  `http://docs.sun.com/source/819-0105/mta.html#wp24778`

- “Chapter 12, Configuring Channel Definitions,” in the *Sun Java System Messaging Server 6 2005Q1 Administration Guide*:

  `http://docs.sun.com/source/819-0105/channel.html#wp997898`

Known Issues and Limitations

See the Java Enterprise System Release Notes Collection at the following URL to find out about known problems:

`http://docs.sun.com/app/docs/coll/entsysrn_05q1`

How to Report Problems and Provide Feedback

If you have problems with Sun Java Enterprise System, contact Sun customer support using one of the following mechanisms:

- Sun Software Support services online at
  http://www.sun.com/service/sunone/software

  This site has links to the Knowledge Base, Online Support Center, and ProductTracker, as well as to maintenance programs and support contact numbers.
- The telephone dispatch number associated with your maintenance contract

So that we can best assist you in resolving problems, please have the following information available when you contact support:

- Description of the problem, including the situation where the problem occurs and its impact on your operation
- Machine type, operating system version, and product version, including any patches and other software that might be affecting the problem
- Detailed steps on the methods you have used to reproduce the problem
- Any error logs or core dumps

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions.

To share your comments, go to http://docs.sun.com and click Send Comments. In the online form, provide the document title and part number. The part number is a seven-digit or nine-digit number that can be found on the title page of the book or at the top of the document. For example, the title of this book is *Sun Java Enterprise System 2005Q1 Technical Note: Sizing Disks for MTA Messaqe Queues,* and the part number is 819-2313-10.

Additional Sun Resources

Useful Sun Java System information can be found at the following Internet locations:

- Sun Java System Documentation
  http://docs.sun.com/prod/java.sys
- Sun Java System Professional Services
  http://www.sun.com/service/sunps/sunone
- Sun Java System Software Products and Service
  http://www.sun.com/software
- Sun Java System Software Support Services
  http://www.sun.com/service/sunone/software
- Sun Java System Support and Knowledge Base
  http://www.sun.com/service/support/software
- Sun Support and Training Services
  http://training.sun.com
- Sun Java System Consulting and Professional Services
  http://www.sun.com/service/sunps/sunone
- Sun Java System Developer Information
  http://developers.sun.com
- Sun Developer Support Services
  http://www.sun.com/developers/support
- Sun Java System Software Training
  http://www.sun.com/software/training
- Sun Software Data Sheets
  http://wwws.sun.com/software