

Performance and Scalability Guidelines

To ensure continuous worry-free operation of MailStore Server, the impact of infrastructure components such as network, storage and other hardware as well as the influence of different configuration options on the overall performance need to be understood.

Please understand that this document can only explain basic coherences without providing explicit instructions applicable for every possible setup. If uncertain, please do not hesitate to contact our technical support for further questions.

These recommendation are based on an average yearly email volume of 10,000 emails per user. The maximum recommended number of users per MailStore Server should not exceed 500 users.

Contents

- [1 Hardware Sizing](#)
 - [1.1 Processor \(CPU\)](#)
 - [1.2 Main Memory \(RAM\)](#)
 - [1.3 Storage](#)
- [2 Archiving Profiles](#)
 - [2.1 Archiving Individual Mailboxes](#)
 - [2.2 Archiving Journal or Multidrop Mailboxes](#)
 - [2.3 Archiving Email Clients or Files](#)
- [3 Other](#)
 - [3.1 Archive Stores](#)
 - [3.2 Indexing of Attachments](#)
 - [3.3 Searching](#)

Hardware Sizing

Providing suitable hardware is the most important prerequisite for high-performance operation as well as persistent scalability of MailStore Server. Below, further details are provided on how hardware components help to ensure the level performance required to satisfy end-users needs.

Processor (CPU)

High CPU usage occurs during email archiving when bodies and attachments are indexed. With an increasing number of concurrent operations, the number of CPU cores has to be raised accordingly.

Keep in mind that an improper configuration of virtual CPU sockets and cores may lead to a decreased guest performance if virtualization technologies are used. Please consult the documentation of your

virtualization solution for further details about the socket-to-core ratio of physical and virtual CPUs recommended for running multiple multi-threaded processes inside virtual machines.

Main Memory (RAM)

Along with the size of the archive, the average memory requirement is slightly increased when archiving new email, browsing the archive or performing searches. Additional main memory is temporarily required for every email to be archived or when a search returns a large number of hits. Therefore it is recommended to leave at least 20% of main memory unused.

To calculate the minimum memory requirements use the following simple formula:

```
REQUIRED MEMORY BY OS + 256 MB + USERS × 5 MB + 20%
```

Storage

MailStore Server is a very I/O intensive process and due to its storage technology comparable with a database server application like Microsoft SQL Server or similar. For that reason the used storage has a massive impact on the performance of MailStore Server, affecting all areas from archiving email to end user access. The following topics should receive special attention:

Connectivity

Generally, direct attached storage (DAS) is preferred due to its low latency. When using network attached storage (NAS), the minimum required disk throughput is 5 megabyte per second. Thus it is recommended to establish a dedicated 1 Gbit/s link between MailStore Server and the NAS. Due to variable available bandwidth as well as unsteady latency, the storage must not be accessed over a WiFi, PowerLine or similar connection. Also using iSCSI or FibreChannel is to be preferred over other protocols such as CIFS/SMB.

Allocation

It is recommended to allocate dedicated disk arrays for MailStore Server to minimize the risk of negative influence caused by other I/O intensive systems such as database or email servers. Additionally, different disk arrays for system (operating system, program and temporary files) and other data should be used.

Hard Disks & RAID Level

Because the overall performance mainly depends on the maximum number of random disk I/O operations per second (also known as *random disk IOPS*), RAID arrays of level 1 or 10 that consist of multiple smaller disks are highly recommended.

Due to the nature of RAID levels 5, 50, 6 or 60, which on average do not offer any write speed gain, their usage would lead to decreased overall performance and thus cannot be recommended for disk I/O intensive applications such running MailStore Server.

For setups a high number of users or very high email volume the use of fast ($\geq 10,000$ RPM) SAS drives is highly recommended.

Adding solid state disks (SSD) for read-write caching or even offloading frequently requested data can result in extra performance, allowing a higher number of concurrent archiving threads or daily email volume.

If uncertain about the optimal configuration of your storage, please consult the storage system's documentation or vendor support to find out about the recommend configuration for disk I/O intensive database applications.

Archiving Profiles

The choice of the archiving strategy may also have an impact on overall performance. Therefore it is important to know how the different archiving profiles provided by MailStore influence the usage of system resources.

Archiving Individual Mailboxes

Obviously, with every running single mailbox or multiple mailbox archiving profile, the workload on email servers and MailStore Server is increased.

Synchronizing the mailbox content with already archived emails primarily puts disk I/O load on the underlying storage, whereas archiving new emails also consumes CPU time and main memory for the purpose of indexing.

To minimize system load and limit the overall execution time of archiving profiles, no more than 5,000 emails should remain in the archived mailboxes. This can be achieved by enabling email deletion by MailStore, e.g. delete emails older than a certain age from the mailboxes. Furthermore, the number of concurrent archiving threads should be kept as low as possible.

The workload created by these types of archiving profiles has to be considered as high, therefore their usage can only be recommended for environments with small mailboxes and a fairly low email volume.

Archiving Journal or Multidrop Mailboxes

When archiving emails from journal or multidrop mailboxes, emails are processed sequentially. Following the recommendation to let MailStore delete emails from such mailboxes after they have been archived successfully, no synchronization overhead exists, contrary to archiving individual mailboxes. Therefore system resources like disk, processor and main memory are only impacted when new emails are archived.

For that reason these archiving profiles are suitable for any environment, irrespective of mailbox sizes or email volume within given limits.

Archiving Email Clients or Files

When archiving via email clients or directly from email files the details in *Archiving Individual Mailboxes* apply to this kind of archiving profiles as well.

Additionally, the execution of client side archiving profiles cannot be controlled centrally by MailStore Server, which increases the risk of too many concurrent archiving threads putting an unexpected high workload on the server.

Therefore client side archiving profiles are suitable for one-time archiving of old emails stored locally on computers as well as for continuous archiving of email in environments with a small number of users and/or low email volume.

Other

Archive Stores

MailStore Server's storage technology consists of individual units (archive stores) with each having its own database and search indexes. In the default configuration MailStore Server automatically creates a new archive store every 5 million emails and activates it to immediately store newly archived emails in it.

In practice, this value has proven to be suitable for most environments. Thus administrators changing this setting, should take care to not select a value that leads to the creation of new archive stores too often.

Information about creating and managing archive stores is available in the chapter [Managing Storage Locations](#).

Indexing of Attachments

A list of too many file types may lead to decreased index performance and therefore has a negative impact on the maximum archiving throughput and additionally increases the CPU load while archiving.

Enter only the file types that are really useful for indexing. Further details regarding indexing of attachments can be found in the [Search Indexes](#) chapter of the MailStore Server Manual

Searching

The impact of search operations depends on their scope, the number of archive stores and found items. The more items are found, the more main memory is required for storing the search result, which means that searching across all user archives typically consumes much more main memory than searching in a single user archive. Also, every full text index file in every archive store must be accessed when a search across all users is performed, resulting in a high number of random reads from the storage. Please keep in mind that it is very common that searches across all user archives containing several tens of thousands of emails can take several minutes.