

# Optimizing Microsoft Exchange Database Replication Performance in Mid-to-Large Scale Enterprise Deployments

<sup>1</sup>Qasim Turab Ali Zaidi, <sup>2</sup>Bader Sulaiman Rumaikhany

<sup>1</sup>Massachusetts Institute of Technology SLOAN School of Management

<sup>1</sup>Saudi Arabian Oil Company, Dhahran, Kingdom of Saudi Arabia

<sup>2</sup>Saudi Arabian Oil Company, Dhahran, Kingdom of Saudi Arabia

---

**Abstract:** This research paper is directed at improving the Replication Overhead of Microsoft Exchange Mailbox Databases in a Database Availability Group for Mid-to-Large sized organizations while not compromising on the High Availability and resiliency of the email service. There are many aspects of a healthy Microsoft Exchange email landscape, yet the one factor which is often ignored is the number of Mailbox Database copies that are managed and maintained by an Organization. The larger an organization, the more complex the email environment, resulting in dependency on Hardware and constant stress due to increase in demand. Validating the Replication Threads for each Exchange Mailbox Server and reducing the number of Copies may be the solution until a new environment is commissioned that addresses the stress pain points. The primary goal of this research paper is to improve upon the reliability and stability of an enterprise email environment, which increases user productivity through improved system response and availability.

**Keywords:** Replication Performance, Microsoft Exchange Database.

---

## I. INTRODUCTION

Microsoft Exchange Server is the most widely, if not the only, used email platform for mid-size to large organizations especially for the on-premises deployments. This paper provides an insight into the performance optimization for mid-to-large scale deployments of Microsoft Exchange Server by utilizing the most optimal hardware configuration. Hardware in terms of CPU and Memory can either provide a consistent experience or negatively impact the email service if not appropriately planned and tested before production.

This research paper aims to provide the most optimal hardware resources required for different on-premises deployments of Microsoft Exchange while not compromising on the High Availability aspect of Mailbox Databases in a Database Availability Group; it looks at one of the factors which is the number of Mailbox Database Copies on each Microsoft Exchange Server in a Database Availability Group. Client connectivity, such as RPC/HTTPS, MAPI and ActiveSync protocols, are not considered in this research. Microsoft provides Exchange Calculator and JetStress tool which is highly recommended to arrive at proper hardware (CPU, Memory, IOPS, and Disk Drives) prior to deploying a highly reliable and scalable Microsoft Exchange environment. These two critical components are not part of this research paper; yet the concept presented here looks at optimum performance based on the replication threads for each Mailbox Database.

*Explanation of the terms used extensively in this paper*

Term	Explanation
Microsoft Exchange	Microsoft Exchange is an email software installed on a Windows Server to provide organizations with corporate email service. Read more <a href="#">here</a>
Mailbox Database	A Mailbox Database is an Exchange database which hosts user's email contents in their mailboxes. The <a href="#">link</a> provides information on what a Mailbox Database is and how to utilize it.
Database Availability Group (DAG)	A DAG is a cluster of Microsoft Mailbox Servers hosting Active and Passive Copies of Mailbox Databases. This cluster is responsible for managing the replication state and health of each Mailbox Database. More <a href="#">here</a>
High Availability (HA)	Mailbox Databases can be deployed with copies on other Microsoft Exchange servers to provide <a href="#">High Availability</a> in case of failures or server being inaccessible.
Disaster Recovery (DR)	Disaster Recovery (DR) is the process of restoring service to normal functionality. While <a href="#">DR</a> is a huge topic in its own, it is not touched in entirety as scope of this research document.
Active / Passive Copy	Active copy of a Mailbox Database is where users connect in order to access the service. Passive copy is a replica of the Active copy which can be made Active automatically by the service or manually by administrators. Microsoft documentation can be found <a href="#">here</a> .

## II. HIGH AVAILABILITY, RELIABILITY, AND IMPACT

### *High Availability (HA)*

Email service must function appropriately round-the-clock without delay. Downtime of email service can be financially devastating for businesses which rely heavily on email service and timely responses. Unavailability or delay in email communications can no longer be tolerated as it impacts the reputation as well as hindering the time-sensitive nature of a corporation to be competitive. If the service degrades, it cannot provide optimal opportunities to a business, thereby severely impacting the ability to compete in today's global economy. Organizations strive to provide uninterrupted email service which is always available, even during the scheduled maintenance timeframes

HA is thus critical to email service; while no organization can prevent outages completely, it is highly imperative to have redundancy and fail-over so that the impact of unavailability of the service or server can be minimized. Microsoft has invested heavily in this area and unavailability of a mailbox database or a mailbox server or the entire Data Centre is built into the Microsoft Exchange Server, however, the HA depends solely on the number of Mailbox Database Copies which are geographically dispersed on Microsoft Windows Servers in Geo-redundant Data Centres.

### *HA Mailbox Database Copies & Redundancy*

Below is a table showing 4 Mailbox Databases hosted on 4 different Microsoft Exchange Servers in a Database Availability Group. 2 of the Microsoft Exchange servers are in Data Centre 1 and the other 2 are in Data Centre 2. If any of the Microsoft Exchange Server is not accessible or the entire Data Centre has communication problems, the email service is not impacted. Note that each Mailbox Database on a Microsoft Exchange Server has 3 copies, 1 each on another Microsoft Exchange Server. In this deployment, the organization can afford a maximum of 3 failures and the 4<sup>th</sup> failure will have impact on the email service.

**TABLE I: 4 HA MAILBOX DATABASES FOR 4 NODE DAG**

Data Centre 1		Data Centre 2	
Server 1	Server 2	Server 3	Server 3
Database1 (Active)	Database2 (Active)	Database3 (Active)	Database4 (Active)
Database3 (Passive)	Database3 (Passive)	Database4 (Passive)	Database1 (Passive)
Database3 (Passive)	Database4 (Passive)	Database1 (Passive)	Database2 (Passive)
Database4 (Passive)	Database1 (Passive)	Database2 (Passive)	Database3 (Passive)

For the simple deployment in Table I, the number of replication threads each Microsoft Exchange Server maintains, for the Mailbox Databases on itself can be explained using the formula:

$$\text{Replication\_Threads\_Server} = (\text{No\_Active\_Copies} \times \text{No\_of\_Passive\_Copies}) + \text{No\_of\_Passive\_Copies}$$

Since the deployment is consistent across all servers, the Replication threads for each Server is

$$\text{Replication\_Threads\_Server} = (1 \times 3) + 3 = 6$$

The Replication Threads means that any change to these Mailbox Databases require Microsoft Exchange Server to continuously monitor and replicate the changes with all member servers in the Database Availability Group. This requires the resources in terms of CPU and Memory available to the Microsoft Exchange Server to be adequately available thus not impacting the performance which would have adverse implications on the users' email service. As the number of users in an organization grow, the number of interactions increase thereby putting extra stress on the resources consumed by the email service. While the hardware cannot often be upgraded, it is highly important to constantly monitor the usage of the email service as well as resources utilized to provide the most available and reliable user experience.

***Impact of Number of Users per Database***

For mid-to-large size organizations, the number of users with the storage allocated for them and HA of the Mailbox Databases must be constantly monitored. As the organization hires more employees, the number of email users grow as well as the size of each Mailbox Database. Among other things, the Database Availability Group of Microsoft Exchange Servers has to provide consistent replication of changes to each replicating partner while maintaining the HA of each Mailbox Database.

Table II shows the same number of Microsoft Exchange Servers in a Database Availability Group with 16 Active Mailbox Databases and 3 Passive Copies each:

**TABLE II: 12 HA MAILBOX DATABASES FOR 4 NODE DAG**

Data Centre 1		Data Centre 2	
Server 1	Server 2	Server 3	Server 3
Database1 (Active)	Database2 (Active)	Database3 (Active)	Database4 (Active)
Database2 (Passive)	Database3 (Passive)	Database4 (Passive)	Database5 (Passive)
Database3 (Passive)	Database4 (Passive)	Database5 (Passive)	Database6 (Passive)
Database4 (Passive)	Database5 (Passive)	Database6 (Passive)	Database7 (Passive)
Database5 (Active)	Database6 (Active)	Database7 (Active)	Database8 (Active)
Database6 (Passive)	Database7 (Passive)	Database8 (Passive)	Database9 (Passive)
Database7 (Passive)	Database8 (Passive)	Database9 (Passive)	Database10 (Passive)
Database8 (Passive)	Database9 (Passive)	Database10 (Passive)	Database11 (Passive)
Database9 (Active)	Database10 (Active)	Database11 (Active)	Database12 (Active)

Database10 (Passive)	Database11 (Passive)	Database12 (Passive)	Database13 (Passive)
Database11 (Passive)	Database12 (Passive)	Database13 (Passive)	Database14 (Passive)
Database12 (Passive)	Database13 (Passive)	Database14 (Passive)	Database15 (Passive)
Database13 (Active)	Database14 (Active)	Database15 (Active)	Database16 (Active)
Database14 (Passive)	Database15 (Passive)	Database16 (Passive)	Database1 (Passive)
Database15 (Passive)	Database16 (Passive)	Database1 (Passive)	Database2 (Passive)
Database16 (Passive)	Database1 (Passive)	Database2 (Passive)	Database3 (Passive)

For this deployment, the Mailbox Database Replication threads for each Server is


$$\text{Replication\_Threads\_Server} = (4 \times 3) + 12 = 24$$

While this does not seem a lot, it surely takes up extra resources; the CPU and Memory on each Microsoft Exchange Server must be adequate for all transactions which are replicated amongst partners of the same Mailbox Database in the DAG. If the servers had physical specifications to meet the requirements of Table I then the deployment of Table II would have severe negative impact to the email service impacting the organization's ability to conduct timely business communications.

For a similar design with extra number of Mailbox Database Copies, say 96 in total (1 HA and 5 Passive) and 12 node DAG, the Mailbox Database Replication threads for each Server becomes:

$$\text{Replication\_Threads\_Server} = (24 \times 5) + 72 = 192$$

This new calculation of exponential increase in replication threads per server reveals below graph of CPU and Memory consumed on an average work day. In Graph 1, it is clear that the servers are under constant pressure and Microsoft Exchange processes are saturated while maintaining a healthy environment. In this scenario, if one or two servers are to go offline, this would increase the load of not only client connectivity but also pressure on the backend services which will ultimately impact the available resources in terms of CPU and Memory, resulting in user communication failures or at best, performance degradation of the corporate email environment.

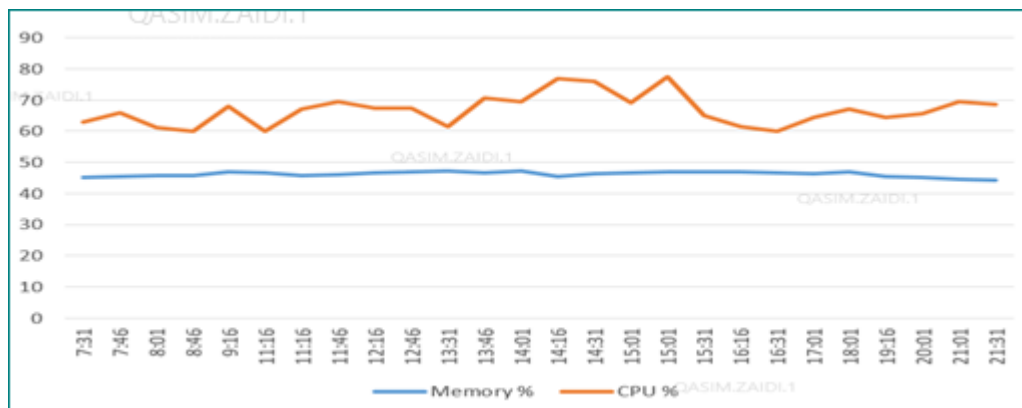


**GRAPH I: CPU / MEMORY CONSUMPTION**

While the best option at this stage is to migrate Microsoft Exchange Servers to newer hardware, tested and proven for the demand using Microsoft Exchange Calculator and JetStress, reducing the number of Mailbox Database Copies may be an option for the time being. Migration is costly and time consuming hence if the organization reduces the number of Mailbox Database Copies from 6 to 5, then it results in below Replication Threads per Server and this may be the workaround for some large organizations.

$$\text{Replication\_Threads\_Server} = (24 \times 4) + 48 = 144$$

With reduction of 1 Mailbox Database Copy in the organization, the number of threads has decreased from 192 to 144, resulting in reduction of 48 Replication Threads per Server, with better consumption of CPU and Memory, resulting in better performance for Microsoft Exchange environment.



**GRAPH II: CPU / MEMORY CONSUMPTION**

Graph II shows consistent and greatly improved performance of Exchange Server resources. The utilization achieved here has direct positive impact on user experience and proper functioning of a healthy email service in the organization.

### III. CONCLUSION

Microsoft Exchange is a business-critical application. Employees rely heavily on email service for their day-to-day business communications. Proper functioning of the email service depends on the availability of resources and meeting the constantly increasing demands of mid-to-large size organizations. This research paper looks at improving the performance of Microsoft Exchange landscape while maintaining the High Availability of the corporate email infrastructure. HA of the email environment depends on the number of Microsoft Exchange Database Copies which in turn can have negative performance impact. Validating these copies and adjusting to numbers which can be sustained is a continuous effort and the user experience must always take precedence; which results in optimizing resources as well as consumption of resources on Microsoft Exchange Servers.

While reducing the copies of Mailbox Databases decrements the fault-tolerance; yet a balance must be achieved between the availability of email service and the performance demands of an organization. The formula presented in this research paper can be utilized to find this delicate balance to provide the best-in-class email service with the most optimal performance of Microsoft Exchange backend services.

### REFERENCES

- [1] Microsoft, "Manage mailbox databases in Exchange Server," <https://docs.microsoft.com/en-us/exchange/architecture/mailbox-servers/manage-databases?view=exchserver-2019>.
- [2] Microsoft, "Database availability groups", <https://docs.microsoft.com/en-us/exchange/high-availability/database-availability-groups/database-availability-groups?view=exchserver-2019>.