Synology®

Synology MailPlus High Availability White Paper

Last updated: Nov 17, 2022

Table of Contents

Executive Summary

Since the email system was invented in the 1980s and with the rapid development of the Internet, emails have become the most popular inter-organization communication channel in the business environment. Businesses must continue operating especially during moments of downtime or system failures. High availability is the ability for a system to operate continuously without failing during periods of loss of services. It is one of the basic requirements to consider when choosing an email system so that it can maintain availability and reduce mail loss.

If you choose to build a high-availability email system, you will encounter problems such as high setup costs, overly complex systems, and management issues. **There are two ways to achieve HA for our email services, MailPlus High Availability (HA) and Synology High Availability (SHA). MailPlus HA cluster architecture is specifically designed for mail services. It provides complete high-availability services and two-way mail replication.** No additional hardware load balancers and independent SAN storage devices are required. It also eliminates the limitation of the physical heartbeat connection. With merely two servers, you can construct a high-availability cluster system that is 100% controlled within the enterprise through a simple setup process to achieve uninterrupted services and complete data synchronization. SHA's protection focuses more on file-protocol services, such as SMB, iSCSI, Fibre Channel, and more.

In this white paper, we will introduce the architecture of MailPlus High Availability. We will cover topics such as the load balancing mechanism of core services, automatic failover, and recovery mechanism. We will also explain how MailPlus high availability avoids data loss caused by a split-brain.

We will be comparing this with our Synology High Availability package and explain why the MailPlus High Availability architecture is a more suitable choice for building a high-availability mail system.

Synology MailPlus High Availability (HA) Architecture

The Synology MailPlus HA architecture provides a complete service to optimize the availability of the mail system. These services include the following:

- Rapid backup and recovery mechanism
- Two-way synchronization of mail data
- Service-level error detection
- Automatic failover

The MailPlus HA cluster consists of two independent operating servers: a primary server and a secondary server. Each server includes a complete mail service and storage system, which handles mail coming in and out, spam detection, virus detection for email attachment, full-text search and indexing, and more.

The primary server runs on the main IP address of the MailPlus HA cluster and receives all mail service requests. These requests will then be assigned to either the primary or secondary server to be processed, and a two-way synchronization will be performed for mail data to be transferred from the primary server to the secondary server or vice versa. The primary and secondary servers may process different service requests, however, data will remain consistent and synchronized across both servers.

To set up a MailPlus HA, you can use Synology NAS models with different computing capabilities to form a MailPlus HA cluster. For example, an old model can be used with a new model with better computing capabilities to form a cluster with HA backup functions. However, it is recommended to use two identical Synology NAS.

Even though the two models can have different hardware configurations, the storage space of both needs to be configured with a large enough capacity to accommodate all mail and users' settings synchronization. When the storage space is exhausted, DSM will warn users that MailPlus needs more storage space. This problem needs to be addressed as soon as possible to ensure the normal operation of the MailPlus HA cluster.

Service Redundancy Mechanism

The core objective of a MailPlus HA cluster is to make sure that the services are uninterrupted when a failover occurs. When any server in the cluster malfunctions and cannot provide services, the other server needs to take over all tasks with all settings remaining the same. As a result, the overall service would not be interrupted to make the whole mail system operate normally again.

The MailPlus HA cluster uses the mechanism of a virtual network interface to create an overall external interface for the cluster. It has a specific unified external IP, which is called the external IP address and is assigned to the primary server. All mail service requests are accessed and received through this IP. These requests will then be assigned to either the primary or secondary server. The two servers communicate with each other to ensure that each server in the cluster is working.

Primary Server Malfunction

When the primary server malfunctions, the secondary server will detect the anomaly within a minute. The secondary server will automatically take over the cluster's external IP address and receive and process email service requests independently. Since the IP did not change, all client programs, such as the mail app on mobile devices and the desktop version of the email client, can continue operating without changing the settings.

Secondary Server Malfunction

When the secondary server malfunctions, it will not impact the mail services of the cluster. The primary server will process all email service requests independently. It will temporarily stop mail synchronization and data configuration, so it is recommended to recover your secondary server as soon as possible.

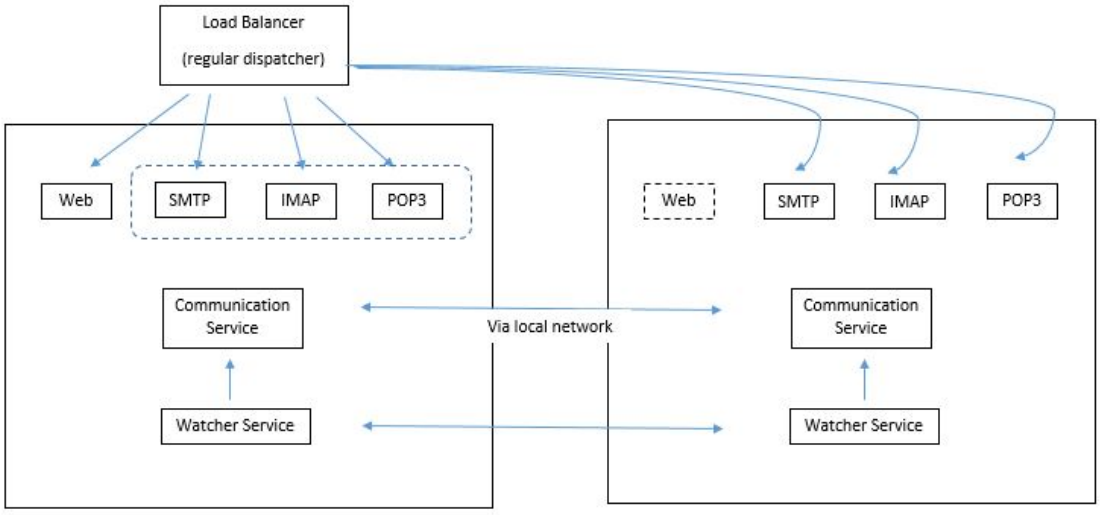

Mail Protocols

SMTP Service

The SMTP service is one of the core functions of the entire system that sends, receives, and forwards messages to external mail systems. After the cluster receives the SMTP request, it will dispatch the request to one of the servers. This server will start processing the mail. Emails have to pass through strict security measures before arriving in users' inboxes. DNSBL, SMTP

command behavior, spam, and virus detection will scan the mail for any issues. After scanning, it is stored on the corresponding server and sent to the recipient's mailbox. The whole mail processing finishes after the mail is synchronized between the two servers.

When one of the servers malfunctions, the remaining server will take over the primary server's role and the cluster's external IP address. That means all the SMTP requests would be received and processed as normal by the new primary server after the switchover without changing any settings. Mails that have been in a queue after failing to send will automatically process as normal without having to manually resend the mail.

IMAP/POP3 Service

In addition to the SMTP service, users will also use IMAP/POP3 to access mail. When the server receives mail, POP3 will download the mail from the server and users can choose to delete the mail from the server. On the contrary, IMAP will store the mail on the server and synchronize changes across multiple devices.

In the MailPlus HA cluster, the IMAP/POP3 requests will also dispatch to the two servers according to the server's processing capacity. When both servers operate normally, the respective computing resources can be effectively utilized by processing different requests mimicking load balancing. When one of the servers malfunctions, the working server will process all IMAP/ POP3 requests to maintain uninterrupted service.

Service Load Balancing

MailPlus leverages SMTP/POP3/IMAP services to each server and receives every request via the cluster's external IP address. After receiving the request, the cluster will dispatch the requests to each server according to their computing power. This can help reduce the workload of each server so that the two servers can work optimally and effectively to process mail and client requests.

Web Client Service

The Synology MailPlus package is a web client interface that allows users to view, manage, send messages, and perform related operations through the browser without installing additional desktop client software onto their operating system. The primary server processes all requests from the web client so it would have more loading under this cluster. Every time a user's settings are modified on the client-side, they are copied to the secondary server to maintain identical data.

Notes:

- This synchronization mechanism requires both servers to install and activate the MailPlus web client package.

Mail Data Replication

Mail and User Data

The mail data of each user of MailPlus, including personal settings, are stored in a separate folder. The user's data will not affect each other, thus improving the independence and security of the data. If a single user cannot access their mail data, other users' data can still be accessed normally, avoiding the risk of large-scale data damage caused by centralized stored data.

When the cluster receives messages, it will perform various scans to ensure the messages are safe to pass and dispatch to the users' mailbox through the mailbox manager. When the system has joined the cluster, it will notify the server responsible for managing the synchronization mechanism through the notification mechanism. A synchronization task will be created for the user. The server will be notified to synchronize the mail and user data.

Notes:

- MailPlus does not synchronize DSM settings, such as permissions, certificates, etc. These must be configured independently by each server and backed up by the Hyper Backup package.

Task Synchronization

The mail service maintains a task queue and processes synchronization tasks in order. When the task starts to process, it will call the synchronization program to connect with the other server. After the connection is established successfully, the two servers will start exchanging synchronization information mutually. They communicate which messages and folders need synchronization or updating. Once the synchronization is done, the two mailboxes are identical.

Mail Processing

When the cluster is operating normally, the user will access their mail from the primary server. If there are any changes to the user's mailbox, it will be processed by the primary server first and then replicated to the secondary server. In rare cases, such as network isolation and instability,

the same user might have access to both servers. If there are any updates or changes, the mail synchronization mechanism will take care of the differences.

In other cases, there might be different action for the same message on both servers. For example, while a message is moved to folder A on the primary server, the same message is dropped to folder B on the secondary server. In this situation, the time of this message and the processing sequence number in the mailbox would be used to determine which folder the mail will be stored.

Failover Events

There are two causes for a failover - power interruption or lost connection.

Power Interruption

When the power of the primary server is interrupted or restarted, failover will be triggered by the secondary server after it fails to detect the primary server within a minute. The data will remain on both servers during this period. After the cluster repairs itself, it will re-synchronize the data.

Lost Connection

When the network connection between the two servers cannot be detected, the secondary server will detect the cluster abnormality and trigger a failover. When the network connection recovers, the cluster will continue to synchronize the differences in data on both servers, reducing the possibility of data loss.

When MailPlus experiences a split-brain or failover, the average failover time is less than one minute because of the active-active server. If any emails fail to send during this period, it usually will be redelivered within 5 to 30 minutes. MailPlus will safely receive the resent mails after the servers restart.

Performance of MailPlus High Availability

In the daily operation environment of an enterprise, under the premise of highly stable services, providing high-performance task processing capabilities and a smooth end-user experience are important goals of MailPlus design. Under the protection of the MailPlus HA cluster, the system performance still needs to be at a considerable level. If we compare the maximum concurrent user number of the MailPlus HA cluster system and the number measured on a non-cluster system, it would be only about 10% lower than the latter one.

Best Practices for Deployment

The MailPlus HA cluster is designed to provide mail service load balancing and service redundancy for a single site. It does not require a physical direct connection for heartbeat detection. Therefore, the configuration is mainly based on the same local network environment. This chapter focuses on the best deployment practices for your MailPlus HA. It is broken up into several factors that you should consider when deploying MailPlus HA.

Service Performances

SSD Cache

MailPlus data types are mainly mail and database files. Therefore, the efficiency of randomly reading and writing small files will highly affect the operating efficiency of MailPlus itself. Compared with hard drives, SSD cache provides a very significant random access performance boost. We recommend mounting the SSD cache on the volumes of the primary and secondary servers where MailPlus is located. The recommended capacity can be configured according to the space expected to be used during planning. Usually, 5~10 % of the total used space is the estimated capacity of the SSD cache. For more information on SSD cache size, visit this page.

Memory Size

The various services in the MailPlus system require enough memory space to load the necessary data for effective operation. We recommend allocating sufficient memory space according to the number of users to maintain the system's performance. The cluster itself will also require some memory space to perform tasks such as mail synchronization. Compared to the memory required by a single-server MailPlus system, the cluster only occupies less than 3% of space. So you can configure the primary and secondary servers according to the memory space requirements of single server deployment. For more information on memory size, visit this page.

Networking

The two servers need to be built in the same local network during configuration so that the cluster operation can provide a virtual network interface to quickly take over the external cluster's IP during failover.

Since the MailPlus HA cluster uses the same network interface, it doesn't need a physical connection for the heartbeat connection. When the MailPlus system is set up, the network interface specifies which connection is used for the cluster's communication. This prevents system service abnormalities that would occur if the cluster used a heartbeat connection.

Reliability

The advantage of having MailPlus HA is that when the primary server malfunctions, the secondary server will temporarily take over all mail service requests. After the primary server recovers, data modifications that have been processed during the failover period will be synced back to the primary server. When the secondary server malfunctions, the primary server will assume all workload and data modifications processed during that period will also be synchronized to the secondary server after it recovers.

Remote Backup

We recommend synchronizing all user mail and configuration data via Shared Folder Sync. We also recommend using Hyper Backup to regularly back up the server's configuration files to a remote location. In abnormal conditions, you can use the backup configuration file to restore the service. The service can be re-launched by adjusting it according to the network environment in the remote location. Since the MailPlus shared folder is replicated by Shared Folder Sync to the remote location, the data and permissions are completely the same as the source. So when the MailPlus settings restore, you can start using it without waiting for a large amount of data to be restored. Starting from MailPlus version 2.2.0, MailPlus is compatible with SHA.

Choosing MailPlus HA or Synology High Availability

The MailPlus HA and the Synology High Availability (SHA) packages provide failover and anomaly detection. This chapter will explain the differences between the two mechanisms and the applicable scenarios to use them.

MailPlus HA operates at the application layer, not at the operating system level or file system level. The servers will synchronize mail data and the settings of all users. Since MailPlus HA uses the active-active mode, only the data differences are synchronized regularly.

We recommend users choose MailPlus HA if they have a large user base (for example, businesses that may use a dedicated NAS for MailPlus Server). In case of a system failure, the MailPlus Servers will recover and operate back to normal quickly.

On the other hand, SHA operates at the systems level and file system level. Since SHA uses the active-passive mode, the system and application data from the active server is overwritten continuously to the passive server.

We recommend users choose SHA if they wish to run multiple services on one single server since SHA offers system-level protections such as when the NAS experiences power outages, storage crashes, file service failures, and more.

Notes:

- For more details, refer to the Synology High Availability White Paper under the **Data Replication** section.

The following is an analysis of the differences between the two architectures:

	MailPlus HA	Synology HA
Scenario	This is recommended for business users who dedicate a NAS to MailPlus Server. MailPlus HA is designed for MailPlus failure detection and mail data synchronization.	This is recommended for users who use multiple packages or services on their NAS, but mainly for file backup purposes.
Operation Mode	Active-active mode	Active-passive mode

Heartbeat Detection Mechanism	The heartbeat connection is established through the general network interface, which is also used for data transfer and services.	The heartbeat connection is established through a dedicated heartbeat interface, which is different from the one used for data transfer and services.
Service Operating Status	All services operate on each server independently.	All services operate only on the active server.
Data Synchronization	Designed for mail data, only the difference of mail data is synchronized.	All system and application data on the active server will be overwritten to the passive server continuously.
Failover	The secondary server will take over the cluster's IP and all services automatically without restarting other DSM services.	The passive server will take over all services automatically, but it needs a longer time to restart all services. Once it's under split-brain status, the whole cluster will be unbinded, and the cluster must be repaired manually.
Anomaly Detection	After a malfunction, the two servers will communicate and select one as the primary server. The cluster would then sync its data and go back to operating normally.	After a malfunction, the cluster needs to be set up manually. The active server will overwrite all data to the passive serve before operating. If the data has been overwritten while malfunctioning, it cannot be rescued.

Conclusion

This white paper provides an overview of the MailPlus HA cluster and explains the characteristics of mail data, load balancing of core services, and two-way backup of services and data settings.

The differential data judgment of the two-way synchronization mechanism greatly improves the stability of the cluster. If it encounters a split-brain exception, which causes the two servers to have new mail, it can prevent data from being overwritten. Therefore, it greatly reduces the probability of mail data loss, and the ability to automatically repair allows administrators to resolve the cluster without manually comparing the different data. This greatly reduces the IT burden and allows IT to focus on other issues.

We compared the Synology High Availability and MailPlus High Availability providing a table comparison of the two and recommended the best scenarios to use either of the packages. MailPlus HA, which is specifically built for mail services, is a more suitable choice for constructing high-availability system protection for the mail services.