



WHITE PAPER

The Microsoft Large Mailbox Vision

Giving users large mailboxes without breaking your budget

Introduction

Giving your users the ability to store more email has many advantages. Large mailboxes keep email on the Exchange Server instead of allowing it to be scattered in Outlook Data Files (.PST files). That helps reduce the risk of data loss, improve regulatory compliance, and increase productivity among both workers and IT staff. The main barrier to implementing large mailboxes is the perceived cost and complexity of storing large amounts of email data.

Microsoft® Exchange Server 2010 is specifically designed to overcome these barriers. It enables you to give your users larger mailboxes at lower cost without sacrificing performance or reliability. Built-in high availability and disaster recovery, storage system improvements, and self-healing from disk faults let you use large, inexpensive disks in configurations that maximize data redundancy. And, Exchange Server 2010 lets you keep your users online during mailbox moves. With these changes, the benefits of large mailboxes are now within the grasp of all organizations.

Problems with small mailboxes

Historical email is a record of knowledge and decisions. When users have small email quotas, they are forced to choose not only which messages they need at the time, but which are going to be important in the future, and they must spend valuable time moving or deleting messages to stay under quota. When that organizational knowledge is deleted or stored in an inaccessible location, users are forced to recreate content that could have simply been reused—provided they even recall what they created in the first place.

Because of the new storage capabilities in Exchange Server 2010, NASDAQ OMX employees across the organization will have 10 GB mailboxes. Larger mailboxes mean that employees can spend less time managing the size of their mailboxes to meet mailbox size quotas and focus more on serving customers.
– NASDAQ OMX

Users may simply not have time to go through their entire mailbox to decide what is important or not, so they often archive older messages in .PST files. This creates a host of problems. Users lose “anywhere access” to the data, as .PST files can only be accessed from the machine on which they are stored and data in lost or corrupted .PST files is gone forever. These files also create problems for administrators and risk for the organization. Because .PST files exist outside the Exchange Server infrastructure, they are difficult to search for legal discovery purposes and are rarely backed up. While you can attempt to prohibit or delete .PST files or create separate storage systems for archived email, the best solution is simply to increase mailbox sizes.

By taking advantage of changes in disk technology and advances in the way Exchange accesses storage, increasing mailbox sizes is easily achieved in a cost effective manner. To realize these benefits, you need to first use the disk storage allocated for Exchange mailboxes more efficiently.

Use storage more efficiently

Storage disks are getting larger and less expensive with little or no increase in their performance capabilities. Storing all email data in one system and increasing mailbox sizes is the best way to efficiently use these large, inexpensive, yet relatively slow disks. This may seem counterintuitive to many IT professionals who are used to the idea of tiered storage: keeping current (hot) data on small, fast, expensive disks and storing archival (cold) data on large, less expensive, slow disks or other storage media such as tape backup. Putting all of your data on the same less expensive disk storage system is the first step to using storage more efficiently and reducing overall storage costs.

To understand why large mailboxes and a single storage strategy enable more efficient use of storage, it is important to consider the two different requirements of Exchange data storage: **disk capacity** and **disk performance**.

Disk performance is measured in IO operations per second (IOPS). Each disk can perform a maximum number of IOPS and each Exchange user consumes a certain number of IOPS. A simplified way to look at this is as a limit on the number of messages that a disk can allow you to send and receive. If the system tries to send/receive more messages than the disk can provide, then users will experience poor performance.

More intuitively, the storage capacity of a disk also impacts mailbox size. It is, however, the proper balance between disk capacity and disk performance which ensures optimal efficiency.

To illustrate this idea, let’s look at an example. Many Exchange systems today use relatively expensive, high performing 146 gigabyte (GB) disks which can perform 150 IO operations per second (IOPS). The easiest way to calculate the number of Exchange users which can be allocated to each disk and provide acceptable user performance is to divide the IOPS of the disk with the IOPS required per user.

Most Microsoft Exchange Server 2007 users require 0.3 IOPS of disk performance. The 150 IOPS provided by the 146GB disk can deliver acceptable performance to a maximum of 500 users. Additional users accessing the disk would exceed the number of IOPS that the disk could provide and would result in performance problems for all users. To use the full capacity of the 146GB disk, those 500 users could be allocated a mailbox size of 300MB each (Figure 1).

This allocation of 500 users to the disk, with the small 300MB mailbox size, provides an efficient balance of both performance and capacity when using these smaller sized disks.

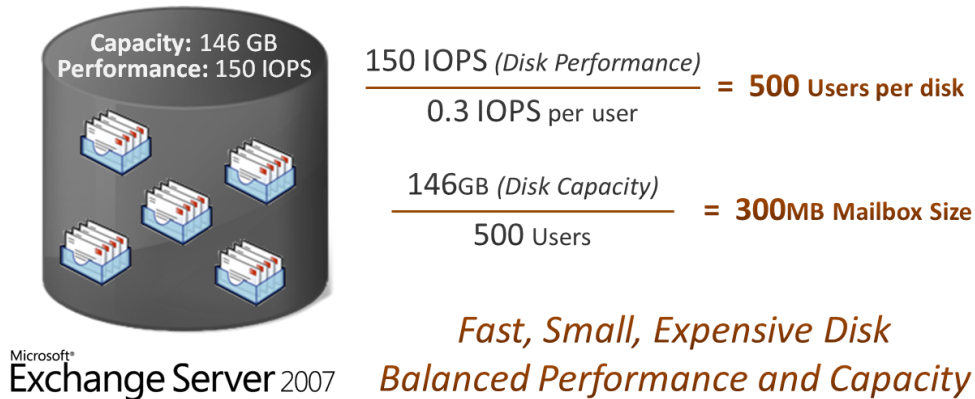


Figure 1: Exchange Server 2007 balancing disk performance and capacity on a 146GB disk

How IO changes make a difference

Just as disks have changed over time to become much less expensive while providing more capacity, the IOPS requirements for Exchange users have also changed. With each version of Exchange Server, the IOPS requirement per user has been dramatically reduced by changing the way Exchange stores data as well as changing how data is read from and written to the disk. For Exchange Server 2010, the IOPS per user is now 0.1—70 percent less than Exchange 2007. This results in better Exchange performance for users with less impact on disks.

One way to take advantage of these Exchange IO savings on the 146GB disk we’ve used in our example is by simply increasing the number of users accessing the disk. However this is not a practical solution. You could support 3 times the number of users per disk—1,500— then you encounter a storage capacity limitation which would reduce the mailbox size for each of those users to only 100MB (Figure 2). This smaller mailbox size results in more problems as users struggle to manage a very small mailbox quota.

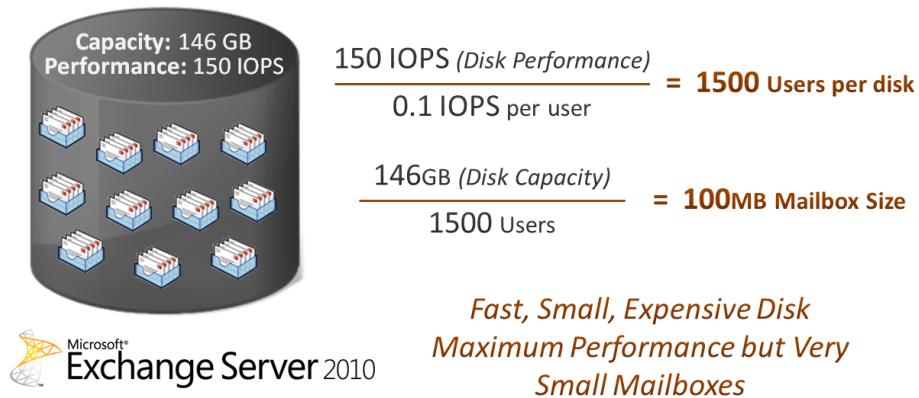


Figure 2: Impact of the Exchange Server 2010 IO changes on smaller disks

Instead of increasing the number of users accessing the disk, you could continue to have 500 users access the disk and retain the original 300MB mailbox size. However, this configuration only uses one third of the available disk performance. By using the full capacity of the disk, but without using the available performance of the disk, the configuration would not be using the disk storage in the most efficient way.

Why using larger disks is more efficient

Modern storage options (less expensive but larger disks) provide a cost-effective way to allocate large mailboxes, which better balances disk performance with disk capacity.

“We haven’t seen any real difference in performance between the SAN and SATA drives. We are getting the same performance from a cheaper disk. The disk capacity for the individual drive is eight times as large and 25 percent of the cost. So our storage cost is about 30 times less with the Exchange Server 2010, when comparing byte to byte.”
- Steve Derbyshire, IT Manager, NEC Philips

A 1 terabyte (TB) disk that can perform only 50 IO operations per second (IOPS) can provide substantially more capacity than the 146GB disk at only 25 percent of the cost.

Using Exchange Server 2010, where the IOPS per user has been reduced to 0.1, a configuration which maximizes the performance of this type of disk would support the same 500 users on the disk as the Exchange 2007 configuration for the 146GB disk (Figure 3).

As disks get larger, if the performance limit is reached on a disk before the capacity of the disk is used, this results in unusable disk storage capacity. In our example, if you configured 500 users on the 1TB disk with only a 300MB mailbox size, then

while the performance capacity of the disk would be exhausted, the disk would only have 150GB allocated for mailboxes and would only be 7.5 percent used, leaving valuable disk capacity unused.

By allocating 2 GB mailboxes for each user, you are able to use both the full performance and the full capacity of the disk (Figure 3) which results in a balanced storage configuration.

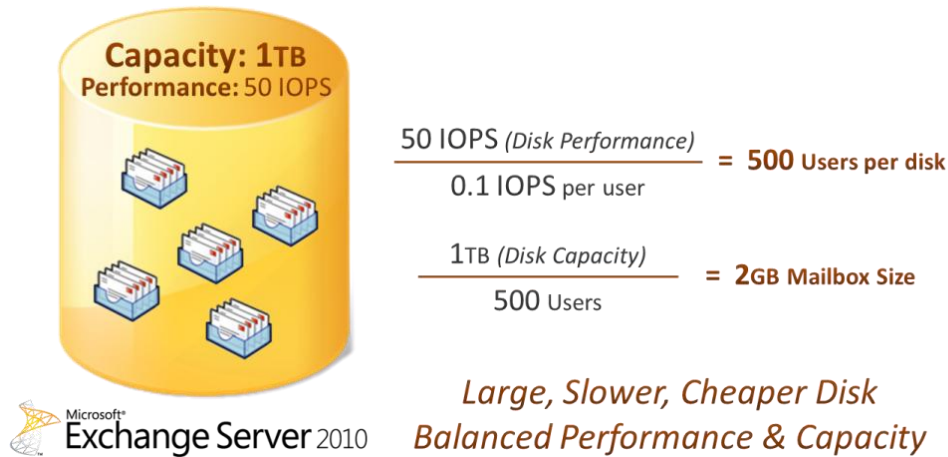


Figure 3: Exchange Server 2010 balancing disk performance and capacity on a 1TB disk

In both these examples, the 500 users are taking advantage of the maximum performance of the disk, resulting in good user performance whether on the faster 146GB disk or the slower 1TB disk. However with the larger, less expensive disk, the mailbox size was significantly increased, taking full advantage of the increased disk capacity while actually reducing disk costs.

By taking advantage of the increasing capacity in disk technology and combining it with the Exchange Server 2010 IO performance reductions, organizations can reconsider their Exchange storage options and provide large mailboxes for their users without breaking their budgets.

Built-in mailbox resiliency

Enabling you to efficiently use high-capacity, low-cost storage is a big step on the way to larger mailboxes. But you also need to make those mailboxes continuously available to users. How do you do that using less expensive disks that could be seen as less reliable? And what about backing up all that additional data? In many organizations, nightly backups represent a significant operational cost that grows with the amount of email stored. If you need to restore the system, doing so from tape or other traditional backups takes longer the more data you have.

Exchange Server 2010 addresses these issues with a built-in mailbox resiliency solution based on database replication to provide high availability, disaster recovery, and backup. The redundancy it provides can significantly reduce your reliance on traditional backups *and* provide the fault tolerance necessary to confidently use less expensive disks. In doing so, it opens new storage options that are less expensive but still provide the reliability that organizations require.

Database Availability Group



The mailbox resiliency solution built into Exchange Server 2010 is based on the Database Availability Group (DAG). A DAG is a set of up to 16 mailbox servers that communicate to manage failures that affect individual databases. Administrators can incrementally add copies of each database depending on the organization's availability requirements and storage limits. Database copies can be configured within the same location to provide high availability, and can also be configured in remote locations for disaster recovery purposes. Exchange Server 2010 automatically switches between these database copies to maintain availability, which means that a database-level disruption, such as a disk failure, no longer affects all users on a server. The ability to switch between database copies has reduced the failover time to around 30 seconds. The system is designed to be cost-effective, scalable, and easy to use.

With high availability built in, you can use less expensive disks to reduce storage costs while actually *improving* the availability of your Exchange system.

Using replication for recovery

The built-in Mailbox Resiliency solution enables a system in which nightly backups for recovery purposes are no longer necessary. When data changes, such as when an email message arrives or is deleted, the change made to the active database is replicated to the passive database copies. If a database fails, a passive copy is activated quickly and automatically. If a server fails, all the databases on that server fail

“With multiple copies of our databases and faster failover times, we can rely on the Exchange Server 2010 high-availability architecture to help us recover quickly from failures, instead of having to restore from backups. If one server goes down, there will be two more to support our mailboxes, so we expect no downtime.” - Breda Hudej, Lead System Engineer, DURS

over to replicas on other servers—leading to a much more efficient recovery than loading a “point in time” backup of the lost data from a separate storage system, which could take days. Also, with a recovery from traditional backups, when a major failure occurs you lose all the data that has come into the system since it was last backed up. With replication based recovery, there is little or no data loss.

One issue that sometimes arises with replication is that if a user or administrator makes a change or deletes something accidentally, the update is replicated to the other database copies. There is no inherent way to “go back in time” as there is with backed up data. Exchange Server 2010 addresses this issue

in a number of ways. Previously, when a user accidentally deleted an email and could not recover it themselves, the administrator would have to restore an entire database to retrieve the deleted item. Exchange Server 2010 introduces Single Item Recovery, which enables administrators to easily recover single messages. Exchange Server 2010 also provides the ability to create a “lagged” database copy within a Database Availability Group which delays updates from being made to a database copy. Lagged

copies are designed to prevent loss of data in the case of logical corruption (a relatively rare occurrence) by providing the ability to recover the database to a specific point in time.

More robust storage systems at lower cost

The key in creating a resilient messaging system is to eliminate single points of failure. A single storage resource, for example, can be susceptible to disaster simply because all the data is stored in one place. With database replication as your redundancy strategy, you can achieve high levels of availability because components of the system can be completely independent of one another.

A common way to reduce single points of failure at the disk level is to use a Redundant Array of Independent Disks (RAID) configuration in which additional disks are used for redundancy to allow recovery in case of failure. Deploying multiple database copies reduces the need for redundancy at the disk level. The advances made by Exchange Server 2010 in database redundancy and resiliency to storage problems combine to make RAID-less disk configurations feasible for Exchange deployments.

Also known as a “just a bunch of disks” or JBOD, this RAID-less configuration has the ability to dramatically reduce your storage costs by reducing the number of disks required. Each database and its log stream are stored on a single disk and replication is used to provide high availability. In the event of a disk failure, a copy of the database is activated, minimizing downtime for users. The number of database copies can be configured to provide whatever level of fault tolerance you require. The combination of JBOD and less expensive disks can result in dramatic cost savings for your organization.

Self-healing

Physical disk errors are known to occur occasionally on low-cost disks (for example, bad blocks). As the amount of data stored by Exchange Server grows, the exposure to disk faults grows too. Since low-cost storage options are a key part of the large mailbox strategy in Exchange Server 2010, the system is designed to deal with these types of disk errors efficiently without having users incurring any downtime. When Exchange detects an error while reading or writing to a database, it fetches a correct copy of the data from a replica of database and repairs the error. When the error is repaired, the disk also re-maps the bad data block, effectively isolating the damaged part of the drive (typically only a miniscule fraction of the disk’s total capacity) so it does not cause future errors. The healthy portion of the disk remains fully functional.

Online mailbox moves

We have seen how Exchange Server 2010 enables the efficient use of large, inexpensive disks to create a resilient, highly available system that reduces reliance on time-consuming backups. This leaves one final management issue to be resolved for large mailbox deployments. In nearly every Exchange Server deployment, administrators will need to move mailboxes between databases or between servers at some point. This could be as users move to different locations or to balance workloads across servers or disks. In many organizations, productivity is lost if users are offline for a matter of minutes, let alone the

hour or more it can take to move a large mailbox. Exchange Server 2010 enables users to stay online during mailbox moves regardless of the size of the mailbox. The user's mailbox remains functional as the data is copied to the new location. When the move is completed, the user is switched over to the new database with no downtime.

Large mailbox performance

While the management issues of larger mailboxes have been addressed by Exchange Server 2010, users still need to have acceptable performance. Larger mailboxes tend to contain more items. To ensure

Elabs will improve customer satisfaction by using Exchange Server 2010 to offer an increase in mailbox size to 5 GB or larger. This larger mailbox size ensures that employees can access the information they need to do their jobs in a quick and efficient manner.
– Elabs

users had acceptable performance with Exchange Server 2007, it was recommended that folders contain no more than 20,000 items each. Thanks to database and indexing improvements, Exchange Server 2010 provides good user performance with folders containing up to 100,000 items.

For users that want to reduce the number of active email items in their primary mailbox, Exchange Server 2010 introduces integrated email archiving and retention policies. This includes a new Personal Archive—a specialized mailbox associated with a user's primary mailbox. When using the archive, all email data is

still stored on the Exchange Server. The combination of the user's primary and archive mailboxes can be considered together as the user's large mailbox.

Conclusion

Larger mailboxes are better for users, IT administrators, and organizations. They use today's storage systems more efficiently. They help users stay productive by giving them better access to the valuable organizational knowledge contained in historical email. Users spend less time managing their email to stay under quota. By keeping as much data as possible inside the Exchange Server 2010 infrastructure, large mailboxes help reduce the proliferation of .PST files and associated backup, compliance, and risk problems. With cost-effective storage options, built-in high availability, an increased number of items possible per folder, built-in email archiving, and online mailbox moves, Exchange Server 2010 enables you to deliver these large mailboxes to your organization whether it is a small business or a global enterprise.

At the same time, you can reduce costs and dramatically improve the availability of your Exchange Server deployment. By integrating everything into one system, you make the most efficient use of today's much less expensive storage options, eliminate the need to maintain separate systems for hot and cold data, and reduce your reliance on costly traditional backups and time-consuming system restores. The net result is a more efficient and reliable business messaging infrastructure that preserves access to the crucial organizational knowledge contained in email messages.