Microsoft® Exchange and Data ONTAP® 7.0

Flexible storage for Microsoft Exchange.

John Phillips | November 2006 | TR-3350

Abstract

Early versions of Microsoft Exchange supported a single mailbox database (information store) per server. The number of users on each server and the size of the databases were purposely limited so that backup and restore operations could complete in a reasonable amount of time.

Today, Microsoft Exchange allows for multiple database instances (known as storage groups) in which multiple databases may be defined. Distributing users among multiple databases eases operational backup and restore procedures and increases data availability. This provides for greater server scalability and more users on a smaller number of mailbox servers.

The features of Microsoft Exchange Server offer great advantages only if the underlying storage can support the availability and performance requirements under all operational conditions. NetApp iSCSI and Fibre Channel (SAN) solutions for Microsoft Exchange 2000, 2003, and 2007 deliver superior performance and advanced features such as online Snapshot™ copies, disk-to-disk mirroring over a LAN or WAN, and the ability to reconfigure storage with no downtime. Intuitive management interfaces facilitate storage management and dynamic provisioning with a fraction of the complexity associated with lesser systems. All NetApp solutions for Microsoft Exchange are listed in the Microsoft Windows® Server Catalog.

Data ONTAP 7.0 offers new functionality that further simplifies Exchange data management while providing administrators the capability to organize available spindles to enhance performance.

Overview of NetApp Solutions for Microsoft Exchange

Network Appliance™ SnapManager® for Microsoft Exchange (SME) software is available for Microsoft Exchange Server 2000, 2003, and 2007. NetApp SnapManager for Exchange is tightly integrated with Microsoft Exchange, which allows Exchange to leverage NetApp storage technology. SnapManager for Exchange is somewhat analogous to backup software. When using an Exchange 2003 Server under Windows 2003, SnapManager for Exchange is integrated into the Windows 2003 operating system volume shadow copy services (VSS) Snapshot framework. SME is a VSS (Snapshot copy) requestor. In this configuration, NetApp storage systems are VSS (Snapshot copy) hardware providers.

For more information on NetApp solutions for Microsoft Exchange and the SnapManager family of products, visit *www.netapp.com/solutions/sc/exchange.html*.

Microsoft Exchange Database Fundamentals

Microsoft Exchange databases are referred to as *information store* (IS) databases. User mail data (mailboxes) resides in a mailbox information store database. Public folders and collaborative data are stored in public information store databases. (From this point forward, there will be references to mailbox store and public store database[s], where applicable.) Exchange *storage groups* are ESE (Extensible Storage Engine) database management instances in which multiple information store databases may be defined.

Microsoft Exchange Server supports multiple storage groups and multiple databases per storage group. Each storage group instance contains the following:

- From one to five mailbox or public information stores (.edb/MAPI message format)
- One set of transaction logs (shared by all databases within a storage group)
- One .stm (streaming file) for each database

Storage Requirements for Exchange

Microsoft Exchange Servers require responsiveness from the underlying storage in support of user activity, backup and restore operations, database maintenance, etc. The requirements for all of these actions can be distilled into two basic categories:

- Capacity
- I/O performance (measured in data flowing in and out of the file system per second, or IOPS)

Common examples that require capacity:

- The number of users times the mailbox quota size (in megabytes) provides a starting measurement for database size (for mailbox stores)
- Policy changes that result in an increase in mailbox quota size (it can be safely assumed that mailboxes will become full)
- Growth in the number of users, resulting in larger or additional mailbox stores
- Third-party Exchange add-on applications such as unified messaging servers, fax servers, and workflow applications that require additional storage to operate

Common examples that generate I/O load:

- User-related loads such as users sending/retrieving e-mail, attachments, calendaring events, working on collaborative data in public folders, etc.
- Server-specific events such as online database maintenance, indexing, etc.

Calculating Capacity vs. I/O Requirements

Capacity requirements are easier to calculate than performance requirements. This is largely because administrators can exert control over the *amount* of data via user mailbox quotas, etc. Calculating the I/O requirements for Microsoft Exchange Servers is far more complicated. For example:

- Users' actions and usage patterns are not entirely predictable
- Peak user activity loads, such as those on Monday mornings, after lunch, etc., are heavier than "steady-state" user loads, which occur during nonpeak business hours, evenings, etc.
- Backup, restore, and data replication/mirroring events require various amounts of disk bandwidth
- I/O performance is highly dependent on spinning hard disks, with mechanical latencies

Calculating I/O Requirements

There are a large number of factors to consider when calculating the I/O requirements and thus the number of disks to support an Exchange workload. NetApp has significant expertise in successfully scaling to support tens of thousands of Exchange users. In addition to years of experience, NetApp has developed internal software tools that employ sophisticated algorithms to calculate complete and accurate system configurations. The software tools

that programmatically support the sizing methods are available 24x7 to all Network Appliance personnel around the world involved in sizing Exchange for NetApp customers.

When sizing storage for Microsoft Exchange, the following (common) factors are taken into consideration:

- The number of users
- Per-user mailbox quotas
- User concurrency
- IOPS per mailbox
- Single instance (message) ratios
- Capacity and I/O related to Exchange Server online database maintenance and deleted items cache
- The number of storage groups and databases per storage group

In addition, the following factors specific to NetApp are entered into the equation:

- Snapshot and metadata IOPS for disk-to-disk operations
- Disk bandwidth requirements for advanced replication and mirroring operations
- High-availability requirements
- Application-specific workload characteristics
- Expected IOPS for 10,000 RPM and 15,000 RPM disks (specific to the Exchange workload)
- Volume sizing and layout relative to Exchange storage group LUNs
- Snapshot operations per day and the number of Snapshot copies kept online at any given time (maximum 255 per volume)

Summary of the method used to calculate I/O and disk requirements:

- Exchange IOPS = (number of users * IOPS/mailbox * % user concurrency * database maintenance multiplier)
- Divide the IOPS into reads and writes (Exchange is most often 70/30 reads/writes)
- Add modifiers for user profiles, various mailbox quota ranges, etc.
- Multiply Exchange load to find the “storage IOPS” related to Snapshot copies, mirroring, etc., operations
- Divide the IOPS against 10,000 RPM and 15,000 RPM disks to determine the correct number of disks
- Calculate the transaction log space, data required by SME, and separate transaction log and database data
- Calculate space and volume requirements for each (right-sized/usable) supported disk size

Operational Events That Impact Sizing

The issues and details covered thus far describe the basis for how NetApp sizes Microsoft Exchange deployments. There are additional factors when service-level agreements (SLAs) and policies result in operational events for any system. Those events, not merely I/O and capacity, can also factor into volume sizing.

For example, SME Snapshot copies and database verifications both utilize volume-level management mechanisms. There are cases in which it makes more sense from an operational point of view to have more volumes. On the surface, creating a larger number of volumes may seem like a valid solution. The challenge is that by dividing the total number of disks across many logical volumes, the aggregate I/O potential of the disks is also divided. While it’s possible to acquire and add more disks to the volumes to guarantee adequate I/O performance, there may be excess capacity left stranded in each volume that can’t be used by databases, etc., in other volumes.

Data ONTAP 7G Performance Benefits for Microsoft Exchange

Data ONTAP 7G provides new functionality that allows the creation of logical volumes for managing data without the need to assign physical disks to the volumes. Instead, the logical volumes all derive performance benefits from a larger pool of physical disks called an *aggregate*.

This results in a number of additional benefits for Microsoft Exchange environments:

- A large number of volumes may be created, all with independent Snapshot copy schedules, mirroring events, etc.
- All volumes can be managed independently while receiving the maximum I/O benefit of a much larger pool of disks

Examples Using Traditional vs. Flexible Volumes

Diagram 1: Using Data ONTAP 7G with Traditional Volumes

The volume structure below illustrates a scenario using traditional volumes that yields the greatest operational flexibility. Using traditional volumes, each disk is associated with a specific volume and is for use for that volume only.

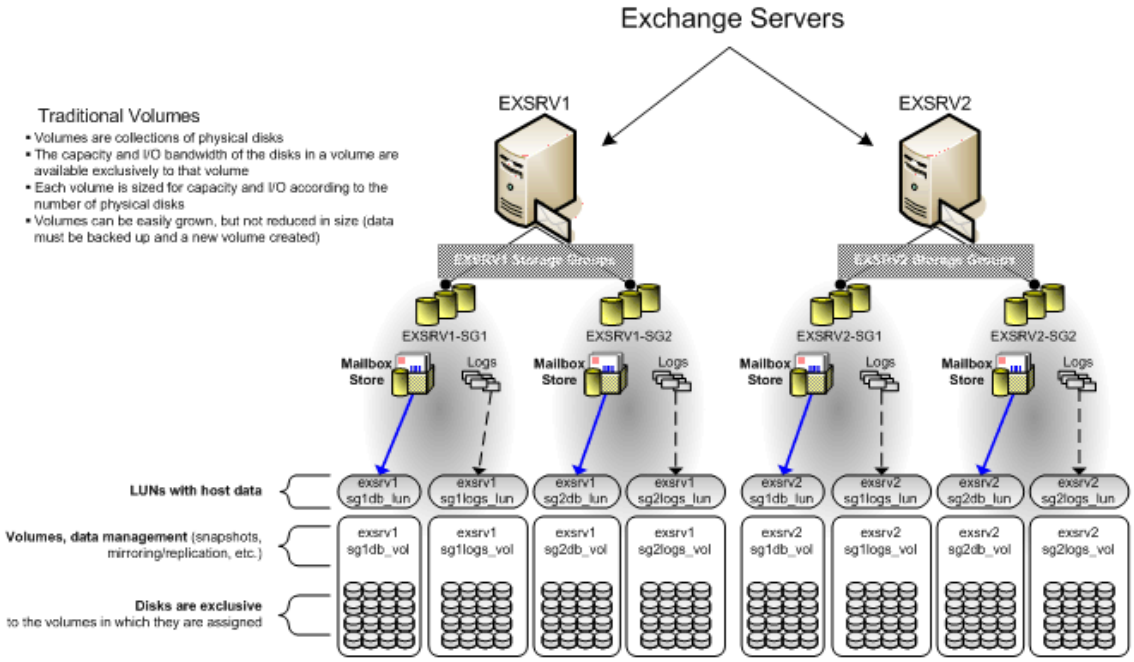

Diagram 2: Using Data ONTAP 7G with Flexible Volumes

This is a simplified example of how the same volume structure that existed previously could be used in a configuration with flexible volumes. This configuration yields the same operational flexibility for SME, but there are other significant differences and advantages.

- Using flexible volumes, the capacity and performance bandwidth of a large collection of fast disks can be made available to all volumes. Even very small volumes have the benefit of a very large number of disks. Volumes can be better leveraged for managing data.

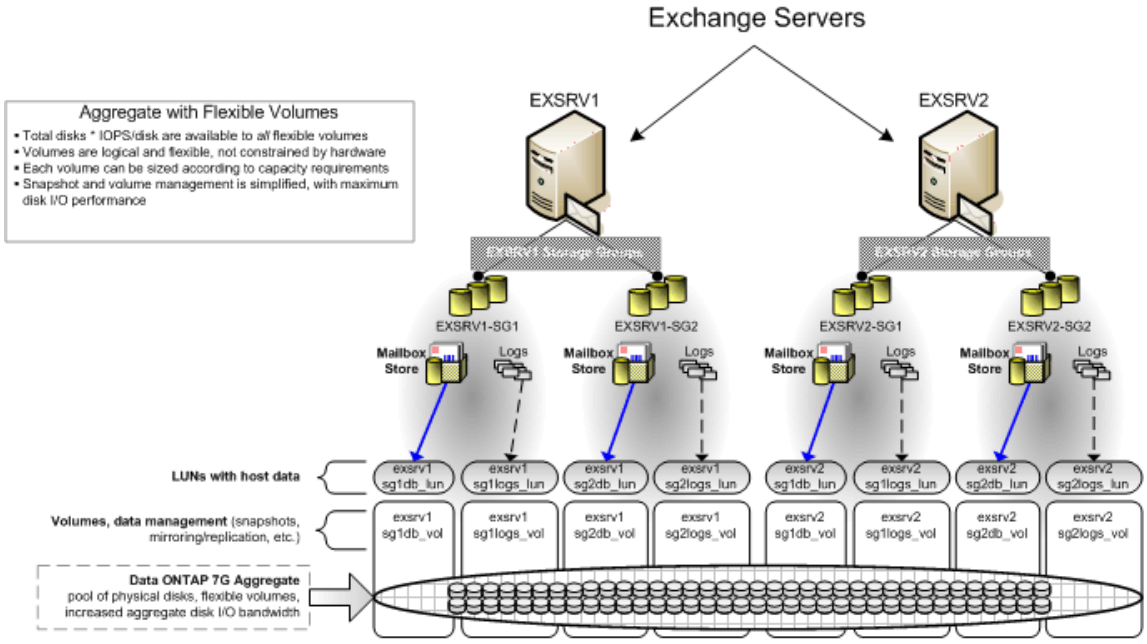

Summary

Data ONTAP 7G provides flexible options that allow Exchange administrators to efficiently utilize the available disks attached to NetApp platforms. The benefits are easily realized for all Exchange Servers storing data on one or more NetApp systems. Traditional volumes or flexible volumes can be used, depending on the operational requirements and performance needs of small to very large enterprise deployments.