

BEST PRACTICES GUIDE:

Nimble Storage Best Practices for Microsoft Exchange 2010



Table of Contents

NIMBLE STORAGE OVERVIEW	3
EXCHANGE 2010 STORAGE REFERENCE ARCHITECTURE.....	3
Store Database and Transaction Log Files on Separate Volumes	3
Connecting Data Storage	3
Use Multi-Path I/O (MPIO).....	4
CAPACITY PLANNING	5
Mailbox Sizing.....	6
Recoverable Items (Dumpster) Sizing.....	6
Exchange Search	6
VOLUME PROVISIONING.....	7
Use Exchange Performance Policies.....	7
Protection Planning	7
Use Volume Collections	7
HIGH AVAILABILITY AND DISASTER RECOVERY	8
BETTER EXCHANGE BACKUPS	8
Use Nimble Windows Integration Toolkit (WIT)	8
Exchange Backup to Disk or Tape	9
Exchange Backup Using Nimble Storage.....	10
SUMMARY.....	10

Nimble Storage Overview

Nimble Storage iSCSI arrays provide a complete Exchange-aware storage solution that includes primary storage, intelligent caching, instant snapshot backup, and replication. Nimble provides iSCSI target volumes to servers that are highly optimized for Exchange by providing the following benefits:

- **Inline Compression:** Reduces storage footprint on physical disks by 50%–70%.
- **Thin Provisioning:** Efficiently stores actual data written rather than reserved space.
- **Snapshot Backups:** Instant point-in-time backup that eliminates the backup window.
- **Zero-Copy Cloning:** Preemptive de-duplication to eliminate storage footprint of repetitive data.
- **WAN-Optimized Replication:** Dramatically reduces bandwidth required for Disaster Recovery.

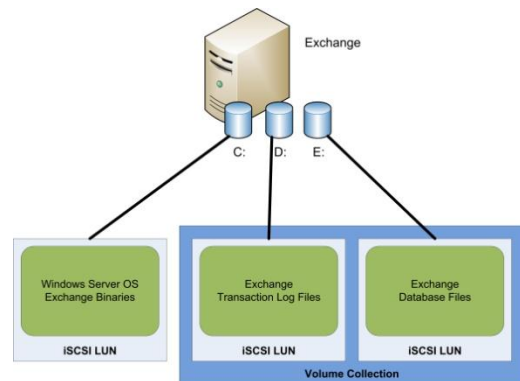
This guide aims to describe the best practices in designing your Exchange implementation on Nimble Storage to deliver optimal performance, protection, and availability.

Exchange 2010 Storage Reference Architecture

Store Database and Transaction Log Files on Separate Volumes

Nimble Storage arrays include high performance auto-tuning features that actively analyze I/O usage patterns and automatically locate hot blocks in high speed SSDs. Meanwhile, the unique, sequential layout characteristics of the CASL file system accelerate all writes. The data placement policies such as RAID-level do not require any manual configuration by the Exchange admin; settings such as block size and caching policy are automatically configured upon selection


of the application or data type from the available list of application profiles. Database and transaction logs require separation on different volumes because of their different tuning and performance characteristics. For example, transaction logs have heavy sequential write activities and need placement in a volume with an Exchange 2010 Log profile, whereas database files have more random read/write activities and need placement in a volume with an Exchange 2010 Database profile.



Connecting Data Storage

Physical and virtual machines should mount iSCSI database and transaction log volumes using the Microsoft iSCSI Initiator. Direct connection to storage volumes through the guest iSCSI initiator allows better storage optimization and enables the use of more granular data protection capabilities. Therefore, virtual servers should avoid mounting data storage to the hypervisor layer using virtual disks. Connecting the database and log volumes directly to the guest allows the Nimble Storage array


to coordinate application quiescence for application-consistent snapshots and replication, as well as provides better control over log truncation with verification schedules.



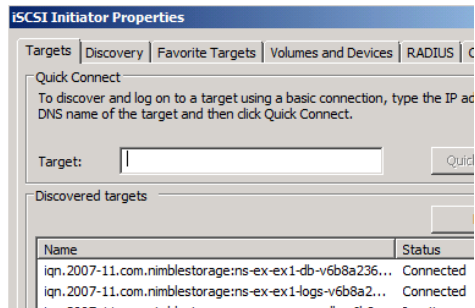
Use Multi-Path I/O (MPIO)

Nimble Storage arrays feature redundant hardware and multiple Ethernet ports that support multiple I/O functionality, increasing performance and providing resiliency against single failure points. Production volumes require connection using the Microsoft iSCSI Initiator (or another supported HBA) and MPIO capabilities provided by the Microsoft Windows Server operating system.

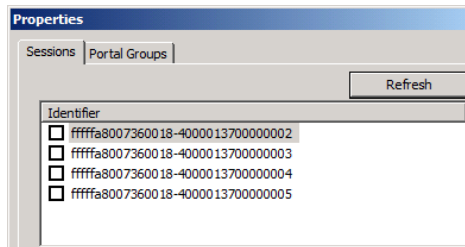
For example, the following Logical Disk Manager view of a guest virtual machine indicates the iSCSI configuration for an Exchange volume of 2 Terabytes.



The Microsoft iSCSI Initiator indicates the same volume connected, which is the first Target in the list below.



Selecting the iSCSI Initiator Target Properties shows four connections between the virtual machine and the Nimble Storage volume using MPIO.



For a detailed explanation on properly enabling the MPIO, refer to the Nimble Windows Integration Toolkit Guide.

Capacity Planning

Nimble Storage arrays can support thousands of mailboxes using a peered combination of solid-state disks and high-density SATA disks. Nimble arrays automatically tune in real-time to cache data in solid-state disks based on I/O performance patterns, such as random versus sequential I/O, as well as recency and frequency of usage. This approach differs from that of legacy storage architectures that add SSD as a tier on top of existing storage, which requires management or inefficient tuning that hinders immediate gains and treats SSD as a hard disk.

Nimble Storage arrays similarly perform real-time compression of application writes to maximize storage density. Depending on the data types and attachments prevalent in the organization, Exchange 2010 data tend to compress between 25%–40% in typical deployments. Therefore, Nimble manages most of the work and greatly simplifies the Exchange capacity sizing. Leveraging the performance and space-saving features of Nimble Storage arrays, we consider the ramifications for Exchange Mailbox Sizing, Dumpster Sizing, and Exchange Search overhead.

The rule of thumb when sizing an Exchange database is to consider the maximum mailbox quota and add the amount of deletion retention space stored in the Dumpster, multiplied by the number of mailboxes, increased by 10% for Search indexing, and then multiplied by 60% to allow for the Nimble compression reduction (1.0–0.40, assuming your data sets average 40% compression).

Database Size = (Max Quota + Retention) * # Mailboxes * 1.10 * 0.60


For detailed information when sizing Exchange database and log implementations, review the Microsoft Exchange sizing guidelines: <http://technet.microsoft.com/en-us/library/ee832796.aspx>

Mailbox Sizing

The mailbox size or quota is traditionally the most important factor in determining the scalability of your Exchange storage deployment and number of users hosted on Nimble Storage arrays. Exchange would typically exceed the storage capacity before the performance capability of a Nimble array with a high performance storage optimization. For example, based on JetStress performance testing, a single Nimble Storage CS-240G array could accommodate 20,000 mailboxes with 1 GB maximum quota per mailbox.

Recoverable Items (Dumpster) Sizing

Exchange 2010 stores deleted items for each mailbox in a Recoverable Items folder (AKA Dumpster), which is divided into three subfolders called Deletions, Versions, and Purges. By default, retention is 14 days for items and 120 days for calendar items. This native item-level recovery feature in Exchange requires additional storage space, but provides user-recoverable functionality. For business needs, the Nimble CASL architecture is ideal for the cost effective extension of this retention period because old deleted emails are purged out of the high performance SSDs and retained on low cost compressed SATA disks. While Nimble Storage arrays provide highly efficient snapshot backup features, you will typically use shorter snapshot retention cycles for Exchange with this Recoverable Items feature, rather than for other data. However, Nimble snapshots are useful in Exchange environments to provide database-level recovery, such as in cases where a virus attack physically corrupts data. Nimble snapshots allow immediate recovery from these data protection issues by either reverting the volume to a previous point-in-time or mounting the snapshot as a zero-copy clone and as a recovery database to correct and restore the item-level.



Exchange Search

By default, Exchange 2010 uses additional storage space to hold the database index, which consumes approximately 10% of the total database size. The index is stored on the same volume that the database uses, and this additional storage space is necessary to support search functionality.

Volume Provisioning

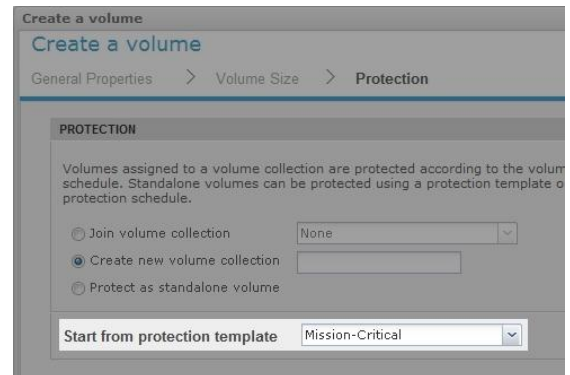
Nimble Storage provides several options when provisioning Exchange mailbox database volumes. A performance policy that matches the application write characteristics of the data stored on that volume should be specified. Use volume collections to group database volumes with their associated transaction volumes to maintain data consistency in snapshot backups, as well as to group data with similar protection schedules and off-site replication requirements.

Use Exchange Performance Policies

The Nimble Storage array includes performance profiles that pre-configure new volumes using optimized configuration settings specific for different usage scenarios. For example, the "Exchange 2010" performance policy uses optimum block sizes to provide the best performance for Exchange database volumes. However, Exchange transaction logs vary in performance, and thus in tuning, compared with their associated databases. The pre-configured "Exchange Logs" performance policy should therefore be used for volumes that store transaction logs.

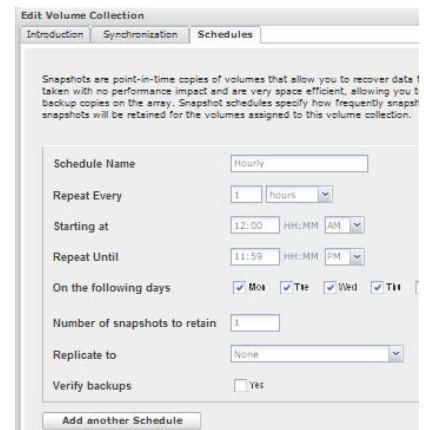
Protection Planning

Nimble Storage arrays provide Protection Templates that consist of pre-configured schedules for snapshots, replication, and retention policies. When creating a new volume collection, select a Protection Template that would insert a default schedule based on existing business rules. For example, you could create Protection Templates based on the criticality of the application data. Less critical applications, such as middleware servers, can use longer snapshot schedule intervals (4 hours) and shorter retention schedules (10 days). However, more critical applications, such as databases with frequently changing data usually require shorter snapshot schedule intervals (15 minutes or less) and longer retention schedules (90 days), thereby a suitable Protection Template with such capacity. Using Protection Templates reduces the amount of work required to create storage volumes and provides consistency for managing similar applications.



Use Volume Collections

A Volume Collection allows you to schedule the snapshot frequency, retention, and replication for a set of volumes to other Nimble Storage arrays. A volume collection can synchronize protection activities between separate yet related volumes, such as database transaction logs and file volumes, to ensure that



snapshots created for databases are application-consistent. The volume collections integrate with and trigger the Microsoft VSS to flush pending database writes in Exchange to disks as well as pause the write activity of the transaction log and database files into a momentarily quiesced state. This approach ensures the data integrity of the point-in-time snapshot backup. Using Nimble Protection Manager to coordinate quiesce provides the added benefit of automating the Exchange log truncation upon the completion of a successful snapshot.

The management of a volume collection allows quick changes to protection schedules for a group of related volumes. For example, you can create a volume collection for the most critical Exchange databases (such as Executive Management) as protection with snapshot backups and off-site replication every 5 minutes. General user mailbox databases could schedule snapshot backups every 30 minutes with off-site replication every 2 hours.

High Availability and Disaster Recovery

Nimble Storage arrays have a built in HA architecture with dual controllers and redundant hot-swappable components. Nimble arrays fully support Exchange 2010 Database Availability Groups (DAG) to create high availability and disaster recovery solutions. DAG uses application-level data replication that sufficiently protects mailbox servers for most environments. However, Client Access Servers (CAS) are unprotected and cannot be clustered with Microsoft Failover Clustering. CAS machines do not hold data but serve as part of the total Exchange solution. If a production failure occurs, then a new CAS could be built to continue email services, or consider hosting the CAS on Nimble Storage and leverage its native replication features to quickly recover these critical services.

Better Exchange Backups

The Exchange 2010 Recoverable Items folder allows an effective and easy item-level recovery. However, the Recoverable Items folder does not protect the integrity of the mailbox database against corruption caused by malware, administrator error, or CAS servers. Therefore, regular snapshot backups remain necessary for your Exchange 2010 mailbox and CAS servers to preserve an effective point-in-time backup. Fortunately, with this recovery feature, maintaining the full mailbox backup for extended periods of time may not be necessary. If you do experience logical corruption or need to recover a deleted mailbox, then Nimble snapshot backups makes these processes much easier. The following sections detail how Nimble Storage dramatically improves the traditional full backup by providing instantaneous full snapshot backups that would never exceed backup windows or inflict noticeable impact on production servers.

Use Nimble Windows Integration Toolkit (WIT)

The Windows Integration Toolkit (WIT) provides an interface between a Nimble Storage array and the native VSS interfaces of the Windows operating system. This interface places the application data into a consistent state that is safe for backup and recovery, ensuring the safe start of databases as well as

the readable and trustworthy data format for Exchange after an outage or user data loss. The Windows Integration Toolkit can be downloaded from the Nimble Storage support web site.

Computer servers continuously write system states and data changes into their storage devices. In taking snapshots, ensure that pending writes are flushed and the volume quiesced in a consistent state for continued proper application functionality in any later restorations to that snapshot. The two primary consistent states are called *crash-consistent* and *application-consistent*.

- Crash consistency generally refers to the ability of a system to crash and allow the operating system to boot as well as find its files and core functionality in a normal, readable state.
- Application consistency adds steps to ensure the safety of an application data. This is often called transactional consistency when referring to database applications such as Microsoft™ Exchange.

When you provision storage on a Nimble Storage array, you select one or more protection schedules that specify the intervals at which your data is protected to that certain point-in-time. During these protection schedules, the Nimble array triggers the WIT to coordinate a quiesce using VSS. After all writes are quiesced, the array performs the snapshot backup and the WIT performs an ESEUtil verification of the snapshot data. Once the snapshot backup and verification are complete, the WIT truncates the Exchange transaction logs to properly manage log growth.

The following timelines and graphs indicate the relative impacts of the different Exchange backup and restoration methodologies after a database failure. These demonstrate the recovery point and time associated with restoring a mailbox database back to full service. These timelines make certain assumptions that the hardware, operating system, and application software are unaffected by the outage. Similarly, the database failed because of an assumed logical failure or erroneous deletion of tables and files. The system load metric is a composite of the total production system involved in the backup process, which includes the CPU, memory, system bus, network, and storage subsystem. Production database processing must compete with the system load imposed by the backup process, which reduces the performance of processing transactions.


Exchange Backup to Disk or Tape

Traditional Backup to Disk/Tape

Max RPO = 1 Hour

Max RTO = 24 Hours

System/Storage Impact = Heavy



The Traditional Backup to Disk/Tape timeline indicates a typical daily backup schedule, which performs a full daily database backup followed by hourly transaction log backups. The system load is high during the full backup when the entire mailbox database and transaction logs are copied to a backup disk or tape device. Later, hourly spikes occur during the incremental backup of transaction logs.

Mailbox database recovery begins with the restoration of the most recent successful full backup, followed by that of each incremental backup. Restoration time is considerable because the entire mailbox database and all logs are copied from the backup media to the server. Thus, if backup takes 5 hours, then restoration would take at least the same amount of time.


Exchange Backup Using Nimble Storage

Nimble Storage Snapshot Backup

Max RPO = 15 Minutes

Max RTO = 1 Minute

System/Storage Impact = Minimal



Nimble Storage arrays are optimized to perform full mailbox database snapshot backups that minimize the impact on the production system. Nimble storage snapshots enable the quick recovery of a mailbox database by creating a clone of the snapshot and direct mounting for total mailbox database restoration. If recovery of a deleted mailbox is necessary, then you should review the detailed step-by-step instructions in the Nimble Storage Windows Integration Toolkit Guide that is available from Nimble's Technical Support web site.

Summary

Nimble Storage provides a complete, high-performance storage platform that significantly reduces the effort typically associated with the implementation of Exchange 2010. Nimble's auto-tuning and inline compression maximizes performance while optimizing the Exchange data footprint. Nimble's performance policies are pre-configured for the I/O characteristics of Exchange 2010 to remove estimations during storage provisioning. Exchange backup is also greatly simplified with Nimble protection schedules and snapshot backup. Nimble's converged storage architecture combines these features to provide a complete primary and backup storage solution for Exchange 2010.



Nimble Storage, Inc.

2740 Zanker Road, San Jose, CA 95134

Tel: 408-432-9600; 877-364-6253) | www.nimblestorage.com | info@nimblestorage.com

© 2012 Nimble Storage, Inc. All rights reserved. CASL is a trademark of Nimble Storage Inc. 0812