

Netmail Detach Configuration Guide

Netmail Detach helps organizations minimize the size of their mail system by removing attachments from their messages and storing them only once in Netmail Store, through single instance storage (SIS). Stripped attachments are then replaced by HTTP links that point to the attachments in Netmail Store. If an organization already has a storage system in place (e.g., SAN or NAS), it is possible to configure Netmail Store Emulator. The Netmail Store Emulator behaves like Netmail Store but runs on the file system. It gives Netmail HTTP access to your SAN, allowing it to fetch attachments from your storage system, when required.

Netmail can remove attachments at three points:

- [At the gateway](#) (for incoming messages from the Internet)
- [From messages already in the Exchange mail server](#)
- [From new messages being sent internally](#) (before they ever reach the Exchange server)

When messages are forwarded internally within the organization, attachments are not re-attached to their respective messages. Recipients need only to click the HTTP links to access their attachments. This applies regardless of whether the original messages came directly from outside the organization or were stripped of their attachments from within the Exchange system.

For outbound messages (sent to external recipients), attachments are re-attached at the gateway before the messages are sent out. It is possible to forgo the re-attachment at the gateway for external recipients, but doing so requires opening the access to Netmail Store to the Internet. While it is safe, most organizations choose to re-attach the files.

Note: Attachment removal at the gateway is applicable to both GroupWise and Exchange environments. However, attachment removal from within the mail system applies only to Exchange environments, as GroupWise still provides SIS for attachments.

Netmail Detach System Requirements

Netmail Detach makes use of Netmail Archive, Netmail Secure, and Netmail Store. They can be deployed together or as separate pieces, depending on your organization's needs and email environment. As such, ensure that your system meets the requirements needed for each of the components your organization is using.

▼ [Netmail Archive](#)

Archive Server (Master or Worker Nodes)

The Archive Server is responsible for connecting to the back-end mail server (GroupWise or Exchange), ingesting data, converting it to XML, and storing the data.

Hardware:

- Processor: 4 total cores minimum
- Memory: 8 GB RAM (4 GB for worker nodes)
- Network: GB Ethernet
- Storage: Up to 4 different spaces can be set up for storage. Smaller organizations might choose to store everything together, while larger organizations may opt to for more granularity.
 - Storage #1: Operating System
 - 250 GB for OS, applications, temp and swap space. VM is delivered with this volume, thin-provisioned.
 - Storage #2: Archives
 - Adequate local or attached storage for archived messages is required.
 - Messages are stored in XML format, accounting for approximately 20% of the total mailbox size.
 - Archives have a read/write profile of 5/95, an expected performance of 500 IOPS for both random read and write, with a recommend block size of 4K.
 - Archives do not change once written, so they can be stored on read/write media or a WORM device.
 - Storage #3: Attachments
 - Adequate local or attached storage for archived attachments is required.
 - Attachments are stored in their native format, accounting for approximately 80% of the total mailbox size.
 - Attachments do not change once written, so they can be stored on read/write media or a WORM device.
 - Storage #4: Audit Files
 - Adequate local or attached storage for audit files is required.
 - Size of audit files requires approximately 15% the amount of archived data (e.g., 1 TB of archives would mean 150 GB of audit files).
 - Audit files are very small in size. Recommended block size of storage device is 2K.
 - Audit files track all operation on archive data and need to be updated, so they cannot be written to a WORM device.

Operating System (pre-installed on VM):

- Microsoft Windows Server 2012 R2 Standard

Software (pre-installed on VM):

- Internet Explorer 11, Firefox 25
- Microsoft Visual C++ Redistributables
- PowerShell 2.0 and 4.0 (required for Exchange deployments only)
- .NET Framework 3.5 and 4.5
- Java JRE 8u45 32-bit
- PostgreSQL 9.3 for logging (separate DB can be used)
- GroupWise 2014 client for OAPI support (required for GroupWise deployments only)
- Strawberry Perl 5.18.1.1
- OpenOffice 3.4.1

Virtual Environment:

VMware

- ESX 3.5 or greater
- vSphere 5.1 or greater
- Vmotion supported
- VMFS (RDM not required, but suggested for better performance)
- Disk Caching on disk array

Hyper-V

- Hyper-V Server 2008 R2 SP1 and higher (Hyper-V 2012 R2 recommended)
- Disk Caching on disk array

▼ Netmail Secure

Secure Server

Each Netmail Secure node that is deployed must be allocated the resources below.

Hardware:

- Processor: 2 cores (or processors)
- Memory: 4 GB of RAM
- Storage: 100 GB on local disk (in the VM)

Software (pre-installed on VM):

- Strawberry Perl 5.18.1.1
- Java JRE 8u45 32
- Open Office 3.4.1

Virtual Environment:

VMware

- ESX 3.5 or greater
- vSphere 5.1 or greater
- Vmotion supported
- Thin provisioning supported; thick is suggested
- VMFS (RDM not required, but suggested for better performance)
- Disk Caching on disk array

Hyper-V

- Hyper-V Server 2008 R2 SP1 and higher (Hyper-V 2012 R2 recommended)
- Thin provisioning supported; thick is suggested
- Disk Caching on disk array

▼ [Netmail Store](#)

Minimum Requirements for Netmail Store

Netmail Store installs and runs on enterprise-class x86 commodity hardware. At least three Netmail Store nodes are required in a cluster to ensure adequate resiliency to failures.

Netmail Store nodes are designed to run using lights-out management (or out-of-band management) and do not require a keyboard, monitor, and mouse (KVM) to operate.

The following table lists the minimum (and recommended) hardware requirements for a Netmail Store cluster.

Hardware:

- Node: x86 with Pentium-class CPUs; x86 with Intel Xeon or AMD Athlon64 (or equivalent) CPUs recommended
- Number of nodes: 3+ (to ensure adequate recovery targets in the event of a failure)
- Switch: Nodes must connect to switches configured compatible with L2 multicasting (IGMP snooping must be turned off)
- Log server: Recommended for accepting tracing from the nodes (included in the CSN server, below)
- Node boot capability: USB flash drive or network drive
- Network interfaces: Minimum 100Mbps Ethernet NIC with one RJ-45 port, recommended 1Gbps or more
- Hard Drives: One hard drive minimum; more recommended (RAID configuration not supported)
- Memory: 2 GB RAM minimum; 4 GB recommended
- NTP server: Required to synchronizing clocks across nodes (included in the CSN server, below)

Notes:

- *All nodes in a cluster must use the same speed network interface to remain compatible. Variable network speeds for different nodes are not a supported configuration. For example, a node with a 100Mbps network interface card must not be in the same cluster with a node that has a 1Gbps card.*
- *Although Netmail Store runs on a variety of x86 hardware, the figures above list the recommended base characteristics for a Netmail Store cluster.*
- *Netmail Store cluster performance improves exponentially by adding additional systems with more intensive CPU hardware and additional memory. The number of recommended systems varies depending on your storage and performance requirements.*

▼ [Expand for more detailed information about hardware requirements](#)

Memory Sizing Requirements

Review the following sections for factors that will influence how you size memory and erasure coding, as well as how you configure Netmail Store.

How RAM Affects Storage

The storage cluster is capable of holding the sum of the maximum object counts from all nodes in the cluster. The number of individual objects that can be stored on a Netmail Store node depends both on its drive capacity and the amount of its system RAM.

The following table shows estimates of the maximum possible number of replicated objects (regardless of size) that you can store on a node, based on the amount of RAM in the node, with the default 2 replicas being store. Each replica takes one slot in the in-memory index maintained on the node.

Amount of RAM	Maximum number of immutable unnamed streams	Maximum number of unnamed streams or named streams
4BG	33 million	16 million
8GB	66 million	33 million
12GB	132 million	66 million

How the Overlay Index Affects RAM

Larger clusters (those above 32 nodes by default) need additional RAM resources to take advantage of the Overlay Index.

To store the same number of `reps=2` object counts above and utilize the Overlay Index, increase RAM as follows:

- Immutable unnamed objects: 50% additional RAM
- Aliased or named objects: 25% additional RAM

Smaller clusters and larger clusters where the Overlay Index is disabled do not need this additional RAM.

How Erasure Coding Affects RAM

The number of erasure-coded objects that can be stored on a node per GB of RAM is dependent on the size of the object and the configured encoding. The erasure-coding manifest takes two index slots per object, regardless of the type of object (named, unnamed immutable, or unnamed anchor). Each erasure-coded segment in an erasure set takes one index slot. Larger objects can have multiple erasure sets, so you would have multiple sets of segments.

In $k:p$ encoding (integers for the data (k) and parity (p) segment counts), there are $p+1$ manifests (up to the `ec.maxManifests` maximum). For $5:2$ encoding, that would mean 3 manifests.

For example, with the default segment size of 200 MB and a configured encoding of $5:2$,

- **1-GB object:** $(5+2)$ slots for segments and $(2+1)$ for manifests = **10 index slots**
- **3-GB object:** 3 sets of segments @ 10 slots each = **30 index slots**

Additional RAM: Larger clusters (above 32 nodes by default) need additional RAM resources to take advantage of the Overlay Index. For erasure-coded objects, allocate 10% additional RAM to enable the Overlay Index.

In summary, Erasure coding users about half the space of replication, but it requires more RAM.

How to Configure for Small Objects

Netmail Store allows you to store objects up to a maximum of 4 TB. However, if you store mostly *small* files, configure your storage cluster accordingly.

By default, Netmail Store allocates a small amount of disk space to store, write, and delete the disk's file change logs (journals). In typical deployments, this default amount is plenty because the remainder of the disk will be filled by objects before the log space is consumed.

However, for installations writing mostly small objects (1 MB and under), the file log space can fill up before the disk space. If your cluster usage focuses on small objects, be sure to increase the configurable amount of log space allocated on the disk before you boot Netmail Store on the node for the first time.

The parameters used to change this allocation differ depending on the software version in use.

By default, Netmail Store is configured to allocate a small amount of disk space to store write and delete journals.

Supporting High-Performance Clusters

For the demands of high-performance clusters, Netmail Store benefits from fast CPUs and processor technologies, such as large caches, 64-bit computing, and fast Front Side Bus (FSB) architectures.

To design a storage cluster for peak performance, maximize these variables:

- Add nodes to increase cluster throughput – like adding lanes to a highway
- Fast or 64-bit CPU with large L1 and L2 caches

Important: If the cluster node CPU supports hyper-threading, be sure to disable this feature within the BIOS setup to prevent single-CPU degradation in Netmail Store.

- Fast RAM BUS (front-side BUS) configuration
- Multiple, independent, fast disk channels, such as:
 - Hard disks with large, on-board buffer caches and Native Command Queuing (NCQ) capability
 - Gigabit (or faster) network topology between all Netmail Store cluster nodes
 - Gigabit (or faster) network topology between the client nodes and the Netmail Store cluster

Balancing Resources

For best performance, try to balance resources across your nodes as evenly as possible. For example, in a cluster of nodes with 7 GB of RAM, adding several new nodes with 70 GB of RAM could overwhelm those nodes and have a negative impact on the cluster.

Because Netmail Store is highly scalable, creating a large cluster and spreading the user request load across multiple storage nodes significantly improves data throughput, and this improvement increases as you add nodes to the cluster.

Tip: Using multiple replicas when storing objects in the cluster is an excellent way to get the most out of Netmail Store, because each copy provides redundancy and improves performance.

Selecting Hard Drives

Selecting appropriate hard drives for the Netmail Store nodes improves both performance and recovery, in the event of a node or disk failure. When selecting drives, these are the key criteria, which are detailed below:

- **Drive type:** enterprise-level drives rated for continuous, 24x7 duty cycles and having time-constrained error recovery logic
- **Drive performance:** buffer cache size and bus type/speed
- **Drive capacity:** trade off of high capacity versus recovery time
- **Disk controller compatibility:** matching what you have in your existing nodes

Drive Type

The critical factor is whether the hard drive is designed for the demands of a cluster. Enterprise-level hard drives are rated for 24x7 continuous-duty cycles and have time-constrained error recovery logic that is suitable for server deployments where error recovery is handled at a higher level than its on-board controller.

In contrast, consumer-level hard drives are rated for desktop use only; they have limited-duty cycles and incorporate error recovery logic that can pause all I/O operations for minutes at a time. These extended error recovery periods and non-continuous duty cycles are not suitable or supported for Netmail Store deployments.

Reliability

The reliability of hard drives from the same manufacturer will vary, because the drive models target different intended use and duty cycles:

- Consumer models targeted for the home user typically assume that the drive will not be used continuously. As a result, these drives do not include the more advanced vibration and misalignment detection and handling features.
- Enterprise models targeted for server applications tend to be rated for continuous use (24x7) and include predictable error recovery times, as well as more sophisticated vibration compensation and misalignment detection.

Drive Performance

You can optimize the performance and data throughput of the storage sub-system in a node by selecting drives with these characteristics:

- **Large buffer cache** — Larger, on-board caches improve disk performance
- **Independent disk controller channels** — Reduces storage bus contention
- **High disk RPM** — Faster-spinning disks improve performance
- **Fast storage bus speed** — Faster data transfer rates between storage components, a feature incorporated in these types:
 - SATA-300
 - Serial Attached SCSI (SAS)
 - Fibre Channel hard drives

Use of independent disk controllers is often driven by the storage bus type in the computer system and hard drives.

- **PATA** — Older ATA-100 and ATA-133 (or Parallel Advanced Technology Attachment [PATA]) storage buses allow two devices on the same controller/cable. As a result, bus contention occurs when both devices are in active use. Motherboards with PATA buses typically only have two controllers. If more than two drives are used, some bus sharing must occur.
- **SATA** — Unlike PATA controllers, Serial ATA (SATA) controllers and disks include only one device on each bus to overcome the previous bus contention problems. Motherboards with SATA controllers typically have four or more controllers. Recent improvements in Serial ATA controllers and hard drives (commonly called SATA-300) have doubled the bus speed of the original SATA devices.

Drive Capacity and Recovery

You can improve the failure and recovery characteristics of a node when a drive fails by selecting drives with server-class features yet that are not the highest capacity.

- **Higher capacity means slower replication.** When choosing the drive capacity in a node, consider the trade-off between the benefits of high-capacity drives versus the time required to replicate the contents of a failed drive. Larger drives take longer to replicate than smaller ones, and that delay increases the business exposure when a drive fails.
- **Delayed errors mean erroneous recovery.** Unlike consumer-oriented devices, for which it is acceptable for a drive to spend several minutes attempting to retry and recover from a read/write failure, redundant storage designs such as Netmail Store need the device to emit an error quickly so the operating system can initiate recovery. If the drive in a node requires a long delay before returning an error, the entire node may appear to be down, causing the cluster to initiate recovery actions for all drives in the node — not just the failed drive.
- **Short command timeouts mean less impact.** The short command timeout value inherent in most enterprise-class drives allows recovery efforts to occur while other system drives continue to support system drive access requests by Netmail Store.

Drive Controller Compatibility

The best practice is to check with Netmail before investing in new equipment, both for initial deployment and for future expansion of your cluster. Netmail can help you avoid problems not only with drive controller options but also with network card choices.

- **Evaluate controller compatibility** before each purchasing decision.
- **Buy controller-compatible hardware.** As a rule, the more types of controllers in a cluster, the more restrictions you face on how volumes can be moved. Study these restrictions, and keep this information with your hardware inventory.
- **Avoid RAID controllers.** RAID controllers are problematic, for two reasons:
 - Incompatibilities in RAID volume formatting
 - Inability of many to hot plug

Mixing Hardware

Netmail Store greatly simplifies hardware maintenance by making drives independent of their chassis and their drive slots. As long as your drive controllers are compatible, you are free to move drives as you need.

Netmail Store supports a variety of hardware, and clusters can blend hardware as older equipment fails or is decommissioned and replaced. The largest issue with mixing hardware is incompatibility among the drive controllers.

Track types of drive controllers

When you administer the cluster, monitor your hardware inventory with special attention to the drive controllers. Some RAID controllers, for example, reserve part of the drive for controller-specific information (DDF). Once a volume is formatted for use by Netmail Store, it must be used with a chassis having that specific controller and controller configuration.

To save time and data movement, many maintenance tasks involve physically relocating volumes between chassis. Use the inventory of your drive controller types to easily spot when movement of formatted volumes is prohibited due to drive controller incompatibility.

Disable volume autoformatting

For additional safety in a cluster with incompatible controllers, set this option:

```
disk.automaticFormat=False
```

This configuration setting prevents volume reformatting if you accidentally move a volume between incompatible controllers.

With automatic drive formatting disabled, you will need to format your new volumes outside the cluster, which you can do using a spare chassis running Netmail Store with the desired controller.

Test compatibility outside cluster

To determine controller compatibility safely, test outside of your production cluster, do the following:

1. Set up two spare chassis, each with the controller being compared.
2. In the first chassis, format a new volume in Netmail Store.
3. Move the volume to the second chassis and watch the log for error messages during mount or for any attempt to reformat the volume.
4. Retire the volume in the second chassis and move it back to the first.
5. Again, watch for errors or attempts to reformat the volume.
6. If all goes well, erase the drive using `dd` and repeat the procedure in reverse, where the volume is formatted on the second chassis.

If no problems occur during this test, you can confidently swap volumes between these chassis within your cluster. If this test runs into trouble, do not swap volumes between these controllers.