

System Requirements and Additional Deployment Considerations

Netmail Store has been designed to work within standard TCP/IP networking environments. This is achieved through the use of widely supported network services and protocols. For a successful installation and deployment of Netmail Store, ensure that your system meets the requirements and recommendations outlined in [System Requirements for Netmail Store](#).

On this page:

- [Operating System](#)
- [Network](#)
- [About Memory Effects on Node Storage](#)
- [Setting Up Network Services](#)
 - [Setting Up NTP for Time Synchronization](#)
 - [Setting Up DHCP for IP Address Administration](#)
 - [Setting Up DNS for Name Resolution](#)
 - [Preparing for Named Objects](#)
 - [Setting Up SNMP for Monitoring](#)
 - [Setting Up Network Load Balancing](#)
- [Object Size Guidance](#)
- [Adaptive Power Conservation](#)
- [Proactive Power Conservation](#)
- [Local Area Replication Using Subclusters](#)
- [Primary vs. Secondary](#)

Operating System

The Netmail Store CSN has been developed and tested with 64-bit RHEL 5.5; other RHEL versions or Linux distributions are not currently supported. Subsequent installation instructions will assume a pre-installed RHEL Linux environment with either internet connectivity or an alternately configured RHEL yum repository for use in installing required third-party packages.

Network

The CSN requires access to both the external network as well as a dedicated internal network. The internal private network ensures the Netmail Store cluster traffic is protected from unauthorized access and also that the external network is isolated from both the PXE network boot server (including DHCP) and cluster multicast traffic. Allocation of the address space in the internal network is broken down as follows, depending on the network size selected during initial configuration (small or large):

Network Size	CSN	Third-Party	DHCP	Netmail Store Netboot
Small (/24)	x.y.z.0-16	x.y.z.17-32	x.y.z.33-48	x.y.z.49-254
Large (/16)	x.y.0.0-254	x.y.1.0-254	x.y.2.0-254	x.y.3.0-x.y.255.254

The CSN range provides IPs for the various services on the Primary and Secondary CSNs. The Third-Party range is provided for third-party applications that need to run on the internal network to interface with the Netmail Store cluster. The DHCP range provides an initial IP to Netmail Store nodes during their initial boot until permanent addresses can be assigned to each Netmail Store process by the CSN. Other applications using the CSN's DHCP server on the internal network will reduce the number of Netmail Store that can be booted at the same time. The Netboot range is used to provide the permanent IPs for all Netmail Store processes.

From a network configuration standpoint, the CSN will automatically allocate half of the detected NICs to the internal network and half of detected NICs to the external network. All NICs allocated to a network are bonded into a bond interface using Linux mode 6, or balance-alb, bonding. In configurations where there is both an onboard NIC card as well as an add-on card *and* the hardware supports detection of the difference between the two, the internal/external network allocation will be disbursed across both cards for redundancy. The CSN NIC assignments may not match the physical NIC ports.

Network Summary

The following table summarizes all required or optional network interfaces used by Netmail Store:

Interface	Type	Required/Optional
SCSP	UDP/80	Required

Web Console	TCP/90	Optional but Recommended
DHCP	UDP/67	Optional but Recommended
DNS	UDP/53	Optional
NTP	UDP/123	Required
Logging	UDP/514	Optional but Recommended
SNMP	UDP/161	Optional but Recommended
mDNS	UDP/5353	Optional
TFTP	UDP/69	Optional

About Memory Effects on Node Storage

The Netmail Store cluster is capable of holding the sum of the maximum stream counts from all nodes in the cluster. The number of individual streams that can be stored on a Netmail Store node depends both on its disk capacity and the amount of system RAM. The following table shows an estimate of the maximum possible number of streams, regardless of size, you can store on a node based on the amount of RAM in the node.

Amount of RAM	Maximum number of immutable unnamed streams	Maximum number of unnamed anchor streams or named streams
4GB	58 million	29 million
8GB	116 million	58 million
12GB	174 million	87 million

Setting Up Network Services

The following topics discuss how to set up network services for your Netmail Store cluster:

Setting Up NTP for Time Synchronization

Messaging Architects strongly recommends you use one or more trusted network time protocol (NTP) servers, whether they are dedicated hardware solutions on your internal network or publicly available NTP servers. This assures that all nodes' clocks are synchronized with each other. When booting Netmail Store from a CSN, the CSN provides NTP services for all Netmail Store nodes and no further configuration is required.

If one or more trusted NTP servers are available, configure Netmail Store to use them by setting the `timeSource` parameter in the node or cluster configuration files. The value of the `timeSource` parameter is a list of one or more NTP servers (either host names or IP addresses) separated by spaces. Examples follow:

- Specifying IP addresses:

```
timeSource = 10.20.40.21 10.20.50.31
```

- To specify host names, the node must be able to resolve host names using a DNS server. For example,

```
timeSource = ntp1.example.com ntp2.example.com
```

For more information about configuring nodes to use NTP, see [Appendix A - Node Configuration](#).

Note:

- NTP version 3 has design limitation that causes time to wrap in the year 2036. If a computer's BIOS clock is set beyond this wrap point, NTP cannot correct the time. Make sure that all BIOS clocks are set to a date prior to 2036 before booting Netmail Store. This is not an issue if you use an NTP server running NTP version 4.
- In the event nodes have no access to external or internal NTP servers, see the discussion of `timeSource = system` in [Appendix A - Node Configuration](#).

Setting Up DHCP for IP Address Administration

While Netmail Store nodes are not required to have static IP addresses to discover and communicate with each other, administrators might find it easier to manage and monitor a cluster where each node receives a predetermined IP address. To do this with DHCP, you must map the Ethernet MAC address of each node to a static IP address. In addition to a node's IP address, Messaging Architects recommends that the DHCP server provide a node with the network mask, the default gateway, and a DNS server address.

Setting Up DNS for Name Resolution

The domain name service (DNS) is used to resolve host names into IP addresses. While DNS is not required for Netmail Store nodes to communicate with each other, DNS can be very useful for client applications to reach the cluster. If you use named objects, DNS is one method you can use to enable access to objects over the Internet.

Although client applications can initiate first contact with any node in the Netmail Store cluster – even choosing to access the same node every time - Caringo recommends that the node of first contact be evenly distributed around the cluster. Basic options follow:

- Define multiple DNS entries ("A" or "CNAME" records) that each specify the IP address of the same Netmail Store node of first contact.
- Use multiple IP addresses for a DNS entry to create a DNS round-robin which provides client request balancing.

The following example shows the entries for three node IP addresses tied to one name. This is the configuration file format of the widely used ISC Bind Daemon.

```
CAStor 0 IN A 192.168.1.101
0 IN A 192.168.1.102
0 IN A 192.168.1.103
```

In the preceding example, it is important that the time to live (TTL) value for each of the records in the round-robin group is very small (0-2 seconds). This is necessary so that clients that cache the resolution results will quickly flush them. This allows for the distribution of the node of first contact and allows a client to quickly move on to another node if it tries to contact a failed node. Although it is recommended that applications implement more robust mechanisms like zeroconf for distributing the node of first contact and skipping failed nodes, an administrator can use DNS to assist with less sophisticated applications.

Preparing for Named Objects

For users to be able to access named objects over the Internet, you must enable incoming HTTP requests to resolve to the correct domain. (A cluster can contain many domains, each of which can contain many buckets, each of which can contain many named objects.) Cluster and domain names should both be IANA-compatible host names like cluster.example.com. For example, a client application can create an object with a name like the following:

```
cluster.example.com/marketing/photos/ads/object-naming.3gp
```

In this example, cluster.example.com is the domain name, marketing is the name of a bucket, and photos/ads/object-naming.3gp is the name of an object. You must set up your network so the host name in the HTTP request maps correctly to the object's domain name. (The cluster name is not important in this regard.)

To enable users to access the preceding object, you must set up one of the following:

- Set up your hosts file to map domain names to IP address(es) of the node of first contact.

For a Linux system, configure /etc/hosts

For a Windows system, configure %SystemRoot%\system32\drivers\etc\hosts

A sample hosts file follows:

```
192.168.1.111 cluster.example.com
192.168.1.112 vault.example.com
```

- Define multiple DNS entries (that is, "A" or "CNAME" records) that specify the IP address(es) of the Netmail Store node of first contact.

Specifying multiple IP addresses for a DNS entry creates a DNS round-robin which provides client request balancing.

Setting Up SNMP for Monitoring

Netmail Store exposes monitoring information and administrative controls through SNMP. An SNMP console provides an administrator a mechanism with which to monitor a Netmail Store cluster from a central location. The SNMP MIB definition file for Netmail Store is located as

follows:

- If you boot from a CSN, an aggregate MIB for the entire cluster is available in `/usr/share/snmp/mibs`.
- If you do not boot from a CSN, the MIB is located in the root directory of the Netmail Store software distribution.

Setting Up Network Load Balancing

Although the Netmail Store nodes interact with client applications using the HTTP communication protocol, the nodes are not simple web servers and they have operation behaviors that are different from traditional web servers. For these reasons, the placement of Netmail Store nodes behind an HTTP load balancer device is not a supported configuration.

During normal operations, a Netmail Store node routinely redirects a client to another node within the cluster. When this happens, the client must be able to make another HTTP request directly to the node to which they were redirected. Any mechanism that virtualizes the IP addresses of the Netmail Store nodes or tries to control the nodes to which a client connects will interfere with Netmail Store and will create communication errors.

Object Size Guidance

This section provides guidelines you can use to size storage volumes for large object sizes. The largest object a Netmail Store cluster can store is one-fifth the size of the largest volume in the cluster. If you attempt to store a larger object, Netmail Store logs an error and does not store the object.

To further tune your hardware planning, keep in mind that the Netmail Store health processor reserves defragmentation space on a volume equal to twice the size of the largest stream that has been stored on a volume. Therefore, you might end up having much lower storage capacity than you expect.

If possible, size your hardware so that the largest streams consume between 10 and 20 percent of available space on disk drives used in the storage cluster. If the largest stream consumes 10 to 20 percent of disk drive space, you get 60% utilization of available space. The percent utilization improves as you add more disk space.

For example, if the largest stream consumes between 5 and 10% of disk space, utilization improves to 80%. If the largest stream consumes only 1.25 to 2.5% of available disk space, utilization is 95%. If disk utilization is diminishing, you should consider upgrading the size of the disk drives in your cluster nodes.

Adaptive Power Conservation

As of the 5.1 release, Netmail Store includes an adaptive power conservation feature that supplements Netmail Store's naturally green characteristics to spin down disks and reduce CPU utilization after a configurable period of inactivity. This feature causes a node that has not serviced any incoming SCSP requests (both client and internode) in the last configurable sleepAfter seconds to change to an Idle status in the Admin Console and to pause its health processor.

A cluster that is constantly in use will likely not benefit significantly from the adaptive power feature but a cluster that has long periods of inactivity on nights and weekends can expect significant power savings utilizing this feature. Because only inactive nodes are affected, maximum available throughput is not affected, although additional latency is incurred on the first access to an node. The cluster automatically awakes one or more nodes to carry out requests when needed and eventually revives all nodes if needed.

For more information about configuring these settings, please refer to [Node Configuration](#).

Proactive Power Conservation

In addition to the power savings gained from the adaptive power conservation feature, administrators may need the ability to proactively reduce the power consumption peak and flatten the power spectrum consumption of the grid where power caps are required for either budgetary or compliance reasons. To support this use case, Netmail Store allows administrators to optionally set the power cap for the cluster to a percentage of the maximum potential power consumption via either the admin console Settings popup or SNMP. It is highly likely that the power cap mode will result in some performance degradation so administrators should be aware of the potential impact to throughput prior to setting the power cap to anything lower than 100%. Note this feature is only currently supported on select Dell hardware.

Note: If the power cap percentage is changed via SNMP or the admin console and the corresponding cluster settings UUID is not updated in Netmail Store configuration, the admin console and SNMP may get out of sync with the actual node state, as the power cap is preserved across reboots even if the cluster settings UUID is not persisted.

Local Area Replication Using Subclusters

Local area replication (LAR) allows an administrator to create logical separations within a Netmail Store cluster in order to define storage distribution strategies. These logical separations change the normal behavior of the cluster so that Netmail Store will attempt to create the greatest logical spread between a stream's replicas by moving them into separate subclusters. If booting from a CSN, administrators should contact their support representative for instructions on manually configuring subclusters as the user interface does not currently support configuration via the CSN Console.

Examples where LAR subclusters are useful:

- Splitting a cluster based on location (data cabinet, building wing)
- Grouping nodes based on common infrastructure (network, power)

An example of splitting a cluster based on location is a company that has data centers in separate wings of their building and wishes to have copies of stored content exist in both wings in case of a partial building loss. A loss could be events like a fire, flooding, or air conditioning problems. Similar to location-based separation, an administrator may wish to split a cluster based on common infrastructure. Examples are grouping nodes by shared network switches, a common Power Distribution Unit (PDU), or a common electrical circuit within a rack.

The network connections between LAR subclusters must have the same speed and latency characteristics as the connections between the nodes. Additionally, all nodes must be in the same broadcast domain such that they are able to send data directly to all other nodes in the cluster and receive the multicast traffic sent from anywhere in the cluster.

Warning: Administrators should avoid frequent changes to the subcluster assignments for nodes. When nodes move to new subclusters, there is the potential for a lot of cluster activity while content is redistributed to account for the new subcluster assignments.

When you retire a volume, you must make sure that sufficient space exists in the LAR subcluster that contains the retiring volumes if you want the separation to persist. Because Netmail Store must maintain the correct number of replicas in the subcluster, retiring a volume without sufficient space can be problematic. For example, Netmail Store might create all replicas on the other side of the subcluster, and might ultimately result in filling up that side of the subcluster.

Primary vs. Secondary

With two CSNs on an internal network, there must be one CSN designated as the primary. The primary CSN is responsible for responding to DHCP requests for the internal network and also listening for all communication on the well-known IP addresses for the internal and external network. When a secondary CSN joins the internal network, it registers with the primary CSN via a privileged SSH channel using the primary root password entered during the initial configuration process. This allows the primary to sync needed information to the secondary, specifically the Primary's Backup Manifest UUID. The secondary CSN provides redundancy for all services on the primary CSN in the event of a disaster and also provides scalability for incoming SCSP requests to the SCSP Proxy.