

Appendix A - Node Configuration

Every node in the Netmail Store cluster reads its configuration information from the *node.cfg* file stored on the USB memory device. This section describes the format and the fields that are contained within the configuration file.

On this page:

- [Format of node.cfg File](#)
- [Option Names and Descriptions](#)
- [Managing Netmail Store Administrators and Users](#)
 - [Defining Netmail Store Administrators and SNMP Administrators](#)
 - [Defining Netmail Store Operators](#)
 - [Securing the Administrator and Operator Passwords](#)
- [Managing Content Integrity Settings](#)
- [Managing Other Stream Replication Settings](#)
- [Volume Specifications](#)
- [Configuring Power Management Settings](#)
- [Managing Other Settings](#)
 - [Configuring external logging \(loghost\)](#)
 - [Operating Nodes Without Using NTP](#)

Format of *node.cfg* File

The *node.cfg* file is a text file that contains name/value pairs for setting configuration options. The format of the name/value pair lines is:

```
name = value
```

Any white-space before the name field, between the name field and the "=", and between the "=" and the value field is ignored.

All fields are case-sensitive.

Blank lines and lines beginning with "#" are ignored. If "#" characters appear on a line after the "=", they will be considered part of the option's value field – not as a comment.

Option Names and Descriptions

The following configuration options are available for administrators to control the operation of the Netmail Store cluster nodes.

Option Name	Default	Description
administrators	{'admin': 'ourpwdofchoicehere', 'snmp': 'ourpwdofchoicehere'}	List of usernames and passwords authorized for read/write access to SNMP and the admin console. By default, an 'admin' user with a default password of 'ourpwdofchoicehere' can perform write actions on the console and an 'snmp' user with a 'ourpwdofchoicehere' password can perform write actions via SNMP.
archiveMode	0	Default of 0 leaves node in a normal operating state. Set to 1 to set the node to archive mode, where the node will remain idle in low-power mode without participating in cluster activity until its capacity is needed.
autoRepOnWrite	0	Set to 1 to force a replicate-on-write for every stream. Set to 0 to allow the normal replication mechanism within the health processor.
autoValidateRead	0	Validate stream contents before returning successful read response.
cache.percentage	10	Percentage of IO buffer memory reserved for the content cache. The content cache speeds up access to frequently accessed content by storing it in geographically proximate locations.

		<p>The amount of IO buffer memory in a node displays as an announcement message when the node starts up. Look for the following in your announcements: MAIN ANNOUNCE: Memory allocation at startup.</p> <p>The default is 10. To disable the content cache, set it to 0. For performance reasons, avoid disabling the content cache unless advised to do so by your support representative.</p>
cache.maxCacheableSize	1048576	Largest size, in bytes, of an object that can be stored in the content cache. The default is 1.048.576 (1MB).
cipTTL	1	Allows configuration of multicast network traffic TTL (time to live). Using the default, the multicast traffic should remain local.
cluster	none	The name of the cluster for use in content distribution. Should not contain special characters like periods or quotes. To prevent confusion, both in viewing the cluster from the administration console and downstream in applications like Content Router that might analyze the source cluster, it is highly recommended that all nodes in the cluster are configured with the same cluster name.
clusterSettingsUUID	none	UUID of the anchor stream holding persistent settings on the cluster.
consolePort	90	Port on which web console listens.
consoleReportStyleURL	none	URL to report style sheet.
consoleStyleURL	none	URL to console style sheet.
defreps	2	The default number of stream instances that must be kept in the cluster if nothing is specified in the individual stream lifepoints. This is required to be 1 or higher and it must be between the values of minreps and maxreps, inclusive (see below).
diskAutomaticFormat	True	Set to False to prevent Netmail Store from automatically formatting a volume at start-up. This is useful in a Storage Attached Network (SAN) environment (especially when using AoE) where Logical Unit Numbering (LUN) masking is not used. In that type of environment, if vols = all, unintended disk formatting can occur. More information about the vols parameter can be found later in this table.
dnsDomain	none	DNS domain name for node.
dnsServers	none	DNS name servers separated by spaces.
domainHeaders	none	A comma-separated list of headers in which to search for the host in an SCSP request. By default, the setting is X-Forwarded-Host, Host
gateway	none	IP address of the default gateway in the subnet. This is only used when <i>ipaddress</i> is set in the configuration file.
group	225.0.10.100	This is the multicast address for the cluster. It must be a Class D IP address in the range 224.0.0.0 - 239.255.255.255. This address must be unique for each cluster. Administrators must exercise care when configuring multiple, distinct clusters to ensure the multicast groups do not overlap, as any node with the same multicast group will become part of a single cluster.

hpStartDelay	900	Number of seconds after a node starts before the health processor begins checking the cluster for the required number of replicas for content stored on this node.
icmpAcceptRedirects	true	Determines whether or not the node accepts routing information from ICMP redirect responses. Valid Values are true and false.
igmpVersion	2	The IGMP version used for host membership queries. Valid values: 1, 2, 3.
ioErrorVersion	200	Number of IOErrors after which a node is taken offline automatically during a retire.
ipaddress	none	Static IP address for a node. If this is not set, a node will use DHCP to obtain its address information.
licenseFileURL	file:///caringo/license.txt	Location and name of the Netmail Store license file.
loghost	none	IP address of an optional Syslog server.
loglevel	40	Controls the minimum level of messages sent to the loghost. Critical messages are logged at higher levels while minor messages are logged at lower ones. 60=announcement, 50=critical, 40=error, 30=warning, 20=info, 10=debug, 0=no logging.
logport	514	Sets the syslog port to use.
maxreps	16	The maximum number of stream instances that can be specified in the cluster. This is required to be 1 or higher and must equal or exceed the value of minreps (see below). Maxreps is a ceiling that cannot be exceeded by individual lifepoint specifications.
minreps	2	The minimum number of stream instances that must be kept in the cluster. This is required to be 1 or higher. This is a floor that cannot be reduced by individual lifepoint specifications.
networkMTU	1500	<p>Sets the maximum transmission unit (MTU), in bytes, that Netmail Store accepts. Set to a higher value to use jumbo frames.</p> <p>Before you change the default MTU value, make sure the node's network interfaces and all other network hardware support the selected MTU. If the hardware does not support the new MTU value, nodes can be prevented from replicating objects and might not be able to communicate with each other.</p>
netmask	none	IP network mask for a node. This is only used when <i>ipaddress</i> is set in the configuration file.
operators	{'snmp':'public'}	List of users authorized for read-only access to SNMP and the admin console. By default, all users are authorized on the admin console without authentication and read-only SNMP users are authorized with an 'snmp' username and 'public' password.
realmCacheStaleTimeout	60	Length of time, in seconds, before the realm cache is cleared. Netmail Store caches security realms and user lists for domains and buckets for the maximum length of time you specify with this parameter. Set this parameter to a shorter value if you expect realms and user lists to be updated frequently. The default and minimum is 60.
repMulticastFrequency	1	Frequency with which UUIDs are multicast to verify

		<p>replicas. By default this is set to 1% but can be changed to a higher approximate percentage (expressed as an integer) up to 100%. This parameter should be set to the same value for all nodes in a cluster.</p>
repPriority	2	<p>enables you to control how aggressively Netmail Store handles asynchronous replications of an object after it is initially written into the cluster.</p> <p>The default setting is 2, which makes the replication secondary to handling user requests. Setting the parameter to 1 gives the replication task equal priority with user requests.</p>
repQueryLowSizeProbKeep	0.0	<p>Probability of a rep query bid being cached at the low end of the stream size curve. Set to 0.0 to disable caching. Set to 0.96 to optimize write performance in large clusters. NOTE: this setting is not compatible with Content Router versions 1.2 or earlier.</p>
repQueryHighSizeProbKeep	0.0	<p>Probability of a rep query bid being cached at the high end of the stream size curve. Set to 0.0 to disable caching. Set to 0.5 to optimize write performance in large clusters. NOTE: this setting is not compatible with Content Router versions 1.2 or earlier.</p>
repMulticastFrequency	1	<p>Frequency with which UUIDs are multicast to verify replicas. By default this is set to 1% but can be changed to a higher approximate percentage (expressed as an integer) up to 100%. This parameter should be set to the same value for all nodes in a cluster.</p>
repthreshold	5	<p>Replication threshold. The value of this parameter is an integer between 0 and 100, representing the allowed difference between volume percent utilizations. Lower values ensure better load balancing and higher throughput. Higher values minimize data movement at the expense of lower maximum throughput.</p>
scspport	80	<p>This is the port on which client applications access cluster nodes with HTTP requests.</p>
sleepAfter	7200 (2 hours)	<p>Number of seconds of inactivity before an active node or volume becomes idle in power saving mode. Minimum supported value is 60 seconds.</p>
snmpCommunity	ourpwdofchoicehere	<p>This parameter was deprecated with the 2010.1 release in lieu of separate read only and read/write authentication using the administrators and operators list. Prior to 2010.1, the parameter controlled access to the SNMP agent for reading and writing values.</p>
snmpSysContact	none	<p>SNMPv2-MIB::sysContact value.</p>
snmpSysLocation	none	<p>SNMPv2-MIB::sysLocation value.</p>
snmpSysName	none	<p>SNMPv2-MIB::sysName value.</p>
spaceErrLevel	15	<p>Percent free cluster capacity to trigger error highlight on the admin console.</p>
spaceWarnLevel	30	<p>Percent free cluster capacity to trigger warning highlight on the admin console.</p>
subcluster	none	<p>Sets the LAR subcluster membership of this node. Should not contain special characters like quotes. Maximum of 16 characters.</p>

timeSource	none	IP address of an optional NTP time server. If this is not set, the node will use its internal clock and the other cluster nodes to synchronize time.
volMinimumGB	64	Minimum size of a device in gigabytes before it is considered for use when using "vols = all" to automatically assign storage volumes. Set to 0 to include all disk devices.
volPluginURL	none	Location of the volume identification tar file plugin if different from the default IO script.
vols	none	Specifies the volume storage devices that Netmail Store will use. This field is required for the operation of a node. For detailed instructions on how to use this field, see Volume Specifications .
volStandbyTimeout	360	Number of seconds of inactivity before an inactive volume changes to a standby state. Minimum value is 0; maximum value is 1200.
volumeRecoverySuspend	0	Set to 1 to disable volume recovery and recovery behavior. Set to 0 (default) to allow normal volume recovery and recovery behavior. This parameter should be set to the same value for all nodes in a cluster.
wakeAfter	28800 (8 hours)	Number of seconds of inactivity before an idle node becomes active again in power saving mode.

Managing Netmail Store Administrators and Users

By default, Netmail Store defines the following realms of users:

- CASstor administrators: Have write access to the Admin Console, meaning they can make changes to the cluster configuration settings. You can also specify administrative SNMP users who have read-write SNMP passwords (that is, community strings).
- CASstor operators: Have read-only access to the Admin Console and to SNMP, and they cannot change cluster configuration settings.

Each realm is specified by a Python dictionary in the following format:

```
{administrators | operators} = {'username':'password', 'username':'password'}
```

Defining Netmail Store Administrators and SNMP Administrators

This section discusses how to set up the administrators parameter to specify users with administrative access to the Admin Console and SNMP. To create the CASstor administrators realm, add user names and passwords as values to the administrators parameter. The following is an example from `/caringo/node.cfg`:

```
administrators = {'admin':'ourpwdofchoicehere', 'snmp':'ourpwdofchoicehere'}
```

In the example above, the CASstor administrators realm has two users: admin and snmp. Both users have the same password: ourpwdofchoicehere. For security reasons, it is strongly recommended that you change these users' passwords as soon as possible.

Note: The names `admin` and `snmp` are reserved and should not be changed. Changing or deleting these names results in errors and unpredictable performance. If you do not wish to use these names, define long, complex passwords for them.

Defining Netmail Store Operators

To create the CASstor operators realm, add user names and passwords as values to the operators parameter. As privileges to the Admin Console are not hierarchical, you must add your administrators users to the operators user list as well. SNMP uses a single snmp user for all access and validates the community string password from the administrators and operators list to determine if the user is allowed read-only or read-write access. The default read-only community password is public.

Securing the Administrator and Operator Passwords

Instead of cleartext passwords, the config file can also use a hexadecimal-encoded MD5 hash of the following string:

```
username:realm-name:password
```

where username and password must consist only of ASCII characters and realm-name can be either *CAStor administrator* or *CAStor operator*.

To create the MD5 hash, use a programming language or a utility like *md5sum* or Apache *htdigest*. For example, to update your node or cluster configuration file with a password hash you create using *htdigest*:

1. Create a hash of the user name, password, and realm.

```
htdigest -c castor_admins "CAStor administrator" Jim.Jones
```

2. When prompted, enter and confirm the user's password.
3. Open *castor_admins* in a text editor.

The hash is the last entry in the string `Jim.Jones:CAStor`

```
administrator:08b0468c1d957b7bac24463dd2191a2d
```

4. Update the administrators parameter in your node or cluster configuration file. The following is a *node.cfg* example:

```
administrators = {'admin':'ourpwdofchoicehere', 'Jim.Jones' : '08b0468c1d957b7bac24463dd2191a2d' }
```

5. Save your changes and exit the text editor.
6. Restart the cluster to use the new setting.

Note that the admin user must always be specified in the administrators parameter, although Messaging Architects recommends you change the user's password. Specifying console administrative permissions in the configuration file on one node allows access to and control of all other nodes in the cluster when the console is accessed from that node. All nodes should specify the same administrators list to avoid differing permissions by node.

Managing Content Integrity Settings

autoRepOnWrite

This option gives an administrator a mechanism to force a second replica of a stream to be written to another node prior to returning a success status to the client. If either write operation fails, the client will receive an error status. Normally, when a stream is first written, its replication constraints will be checked in the normal course of the health checking of all streams on the node. While the behavior should be acceptable for most installations, there may be business requirements to immediately duplicate the stream to another node as quickly as possible. The `autoRepOnWrite` option will only force one immediate copy of a stream to another node. If the stream has a constraint for additional copies beyond two, the additional copies will be made during the normal health checking process. Be careful when using the replicate on write feature because it will slow down write performance for all clients.

repPriority

This option allows an administrator to control how aggressively Netmail Store handles asynchronous replications of an object after it is initially written into the cluster. The default setting is "2" which makes the replication secondary to handling user requests. Setting the parameter to "1" gives the replication task equal priority with user requests. The `autoRepOnWrite` option can be used to make replication synchronous – the highest priority.

autoValidateRead

This option gives an administrator a method to require validation of all content reads before returning a successful read completion. Although this option can be specified on a per-read basis, setting the value to "1" in the configuration file forces all reads to use validation. During the read from the disk, the content hash is computed. If the hash is wrong, indicating logical disk corruption, the socket will be closed before the last block is transmitted, forcing an error to the client. Using this option creates additional CPU load on the node.

Managing Other Stream Replication Settings

minreps, maxreps, and defreps

These options give an administrator the ability to control how many replicas of streams can be created and maintained within a cluster. Since each individual stream can request a certain number of replicas be created, these parameters allow the administrator to override those choices to account for the characteristics of the cluster as a whole. The `minreps` option gives an administrator the ability to globally set the minimum number of instances that any stream is allowed to have, irrespective of its lifepoint specifications. This value is a floor and will override a lower value that may be set by an application developer. If an application requests more instances than the `minreps` value, the higher number requested by the application will be used. Similarly, the `maxreps` option gives an administrator the ability to limit the requested number of replicas so that it cannot exceed a certain value. For small clusters, `maxreps` could be set to the total number of nodes in the cluster, thus preventing needless error messages when the cluster is unable to fully replicate streams because there are too few nodes. The final option, `defreps`, allows the administrator to specify a default number of replicas that is used only if an individual stream does not specify number of reps in its lifepoints. As the names imply, `maxreps` should be at least as large as `minreps` and `defreps` should fall between `minreps` and `maxreps` (inclusive).

hpStartDelay

The `hpStartDelay` option allows an administrator to control the delay before the health processor begins checking the cluster for the required number of replicas for content stored on this node. This delay window occurs once following the startup of a node and provides a grace period for other nodes in the cluster to stabilize. This is useful in situations where an entire cluster must be shutdown and restarted.

The health processor helps maintain data integrity on the cluster by performing the following tasks:

- Verifies the correct number of replicas exist in the cluster and makes sure replicas are properly distributed on subclusters.
- Enforces lifepoint policies (enforces replica counts and delete policies (terminal lifepoints specifically) at different policy time intervals).
- Validates the object's MD5 on the disk and creates a new replica if it doesn't match. This can happen over time because of disk errors.
- Periodically evaluates if an object is optimally stored in its current location.
- Defragments storage space on an as needed basis.

Volume Specifications

vols

The `vols` option in the node or cluster configuration file specifies the volumes that Netmail Store uses. This specification includes the device names and optional flags for handling these volumes. Only one `vols` parameter is allowed in a configuration file.

Warning: *Netmail Store will erase any non-Netmail Store data on all the volumes it uses. It is recommended that Netmail Store be run only on nodes that are free of non-Netmail Store data.*

The formal syntax of the `vols` option is as follows:

```
vols = volume-specification
volume-specification ::= all-volumes | volume-list | ''
all-volumes ::= 'all' [ ':' policy ] [ space 'except' space device-list ]
device-list ::= device [ space device [ ... ] ]
device ::= Linux-device-path
space ::= space or Tab character and not the word space
policy ::= 'k'
```

The recommended setting is `vols = all`, which enables Netmail Store to use all volumes except the USB flash drive. These are examples of `vols` entries:

```
vols = all
vols = all:k
vols = all except /dev/hda
vols = /dev/sda /dev/sdb
vols = /dev/hda:k /dev/hdc:k
```

If the `vols` option is blank, the Netmail Store node will startup as a volumeless node that is able to process client requests. Some sites may wish to

do this in order to have fixed IP addresses for a few nodes that will be stable primary access nodes for all client operations.

device

The device component is either the keyword `all` or the Linux device path string for the drive. When the keyword `all` is used, do not include any other device path specifications. Messaging Architects strongly recommends you use `vols = all` and discourages you from using Linux device paths. Examples of a device path are:

- `/dev/hda`
- `/dev/sda`
- `/dev/sdb`

The older IDE disks, also known as Parallel ATA, EIDE, ATA-33, ATA-66, ATA-100, or ATA-133, use the "hd" device names. These drives are configured as master or slave devices on each IDE controller. Typically, the master devices will be `/dev/hda` and `/dev/hdc`, while the slave devices will be `/dev/hdb` and `/dev/hdd`. SCSI and Serial ATA (SATA) disks use the "sd" device names. The device letters are assigned sequentially in the order in which the drives are discovered starting at `/dev/sda`. If an invalid device name is specified in the device component, the node's log will indicate an error during the format operation. Additionally, the node will not include the storage space from any disks that are incorrectly specified in the `vols` parameter.

policy

The policy option allows an administrator to give Netmail Store instructions on the appropriate handling on a volume. Currently these handling features involve the formatting characteristics of the physical device. The format policy allows an administrator to override the default volume expiration behavior by specifying the `:k` (keep) policy. See [Volume Management](#) for an explanation of volume expiration.

Specifying Exceptions

Although it is recommended that you allow Netmail Store to use all disk volumes, you can optionally exclude certain volumes from being formatted and used by Netmail Store. For example:

```
vols = all except /dev/sda /dev/sdb
```

As shown in the example, `except` is used only with `vols = all`.

Configuring Power Management Settings

This section discusses how to configure power management settings for nodes and volumes. This feature, referred to as Darkive™, causes a node that has not serviced any incoming SCSP requests (both client and internode) in the last configurable time period to change to an Idle status in the Admin Console and to pause its health processor. Power management can be set for a node or for a volume. A node power management setting affects the entire node; that is, all volumes on the node. A volume power management setting affects only individual volumes in the node.

The term "Idle", when displayed in the Admin Console, has different meanings for a node and volume:

- **Idle node:** A node which has had no SCSP activity for a specified length of time. In an idle state, a node's health processor pauses so an idle node is more likely to have idle volumes.
- **Idle volume:** A volume which has had no I/O activity for a specified length of time.

sleepAfter

This option determines how long a period of inactivity should occur with no incoming SCSP requests before Netmail Store pauses the node's health processor and displays it as Idle in the Admin Console. If you select **Full Performance Mode** on the Settings page in the Admin Console, a node never displays as Idle. Setting this value to a long period prevents Netmail store nodes from becoming idle and taking advantage of power savings. Setting the value to a small number allows nodes to become idle after a reasonable period of inactivity (two hours by default). The smallest supported value is 60 seconds.

Note: The `volStandbyTimeout` parameter is independent of this setting. In other words, it is possible for a volume to be idle when its node is not idle. However, because an idle node's health process or is paused, it is more likely that an idle node has idle volumes than an active node.

wakeAfter

This option allows an administrator to determine how long a period of inactivity should occur with no incoming requests before Netmail Store wakes up an idle node or volume in order to allow the health processor to validate disk and content integrity and replicas. Setting this value to a long period may unnecessarily put content at risk as the health processor does not run on sleeping nodes. Setting the number to a small number will reduce power savings, although the volumes and nodes will go back to 'sleep' if additional client activity is not received.

volStandbyTimeout

This option determines how long after a volume's last I/O before the volume spins down. This setting is independent of sleepAfter. The default is 360 seconds (six minutes). Minimum value is 0; maximum value is 1200.

Managing Other Settings

consolePort

The consolePort option allows an administrator to specify the port on which the web administration console listens for requests. When Netmail Store is deployed into untrusted network environments, this port should be firewalled so that only administrators can access it. This setting must be the same on all nodes.

domainHeaders

To comply with [RFC 2616](#), Netmail Store currently requires you to specify a Host header with every SCSP request. This mandatory header was added to HTTP/1.1 to support web servers or server farms that host more than one domain. Your client might use an HTTP proxy that modifies the Host header, but the Netmail Store domain name matches the original Host header. In that case, an HTTP proxy copies the original Host header into another header, typically X-Forwarded-Host. Use the domainHeaders configuration parameter to specify a search order for headers in which to find the host of a request. By default, domainHeaders is set to the following:

```
X-Forwarded-Host, Host
```

To change the list of headers or the search order, add domainHeaders to the cluster or node configuration file.

Configuring external logging (loghost)

The loghost option allows Netmail Store log messages to be sent to a central Syslog server. The examples here show the specific Netmail Store items that need to be added to the UNIX configuration files for syslog or syslog-ng. The actual configuration files will likely contain additional information for logging messages from other hosts and other logging facilities.

To configure your Syslog server:

1. Set syslog or syslog-ng daemon options to enable logging from a remote host.

- syslog:

Edit `/etc/sysconfig/syslog` to the SYSLOGD_OPTIONS to the following:

```
SYSLOGD_OPTIONS="-r -m 0"
```

- syslog-ng:

An example is shown in the next step.

2. Configure other logging options.

- syslog:

This example shows a sample configuration with the standard syslog program. When editing the `syslog.conf` file, remember that `facility.level` must be followed by a Tab character to separate it from the destination specification.

```
# /etc/syslog.conf
local6.* /var/log/castor.log
```

- syslog-ng:

This example shows a sample syslog-ng configuration:

- Add the following to the filters section (If a filters section does not exist, add one after the source statement and before the destinations

section.):

```
#CASTor filter
filter f_castor {facility(local6);};

#CASTor destination
destination d_castor {file("/var/log/castor.log");};

#Caringo-specific additions
#
log {source(s_net); filter(f_castor); destination(d_castor);};
```

Note: The statement `(source(s_net);` is the default remote service source for `syslog-ng`. If you use a different remote source name, replace `(source(s_net);` with the name you used.

- If you changed your `syslog-ng` settings, enter the following commands to restart the service:

```
sudo /etc/init.d/syslog-ng stop
sudo /etc/init.d/syslog-ng start
```

3. Set Netmail Store logging options in the node or cluster configuration file.

- `loghost`: Fully qualified host name or IP address of your `syslog` or `syslog-ng` server.
- `logport`: Server's port.
- `loglevel`: Set to one of the following:

Log level	Meaning
0	no logging
50	critical messages and announcements
40	errors and all of the preceding
30	warnings and all of the preceding
20	info and all of the preceding
10	debug and all of the preceding (the most verbose log level)

timeSource

It is strongly recommended that you use Network Time Protocol (NTP) servers to synchronize the clocks of all nodes in the storage cluster. Netmail Store's integral processes, such as versioning, lifepoints, and updates, depend on all the clocks of all nodes in the cluster being synchronized with each other. Failure to provide synchronization can lead to unexpected results.

Netmail Store does not synchronize the clocks of client machines. As a best practice, you should also synchronize clients' clocks with NTP but that is not required. Use NTP servers you trust, whether they are dedicated hardware solutions your internal network or external, public servers.

To set a node to use NTP, use the `timeSource` parameter in the node or cluster configuration file as follows:

- `timeSource` blank or not set

Netmail Store uses NTP servers from the public NTP pool for a one-time synchronization at start-up, provided `dnsServers` is set to a valid value or the node uses DHCP that provides DNS servers.

Note: If nodes do not have network access to public NTP servers, they eventually time out waiting for a connection and automatically restart. As a result, it is important that `timeSource` either be set or that the nodes have network access to the NTP pool servers.

- `timeSource` set to one or more IP addresses or host names (recommended)
 - `timeSource` set to one or more IP addresses

If the node has no DNS server defined for it, enter a space-separated list of NTP server IP addresses. For example,

```
timeSource = 192.168.0.20 192.168.0.50 192.168.0.110
```

Note that `dnsServers` must be set to a valid value.

- `timeSource` set to one or more host names

Use this option if the node uses DHCP that provides a DNS server to resolve host names. For example, you might choose to use [NTP pool servers](#). [ntp.org](#) recommends you use pool servers that are close to your servers' time zone as discussed on the [ntp.org help page](#). For example, to use United States-based pool servers,

```
timeSource = 0.us.pool.ntp.org 1.us.pool.ntp.org 2.us.pool.ntp.org
```

Operating Nodes Without Using NTP

Messaging Architects strongly recommends against operating a cluster in a production environment without using NTP. However, in demonstration or development environments where there is no internal or external NTP server available, you can choose a minimum of one or a maximum of two nodes as the master clock and synchronize all other nodes' clocks to the clock of that node.

Warning: *Extra care should be used when operating without an external NTP time source. Specifically, administrators will need to assure that all new nodes have their BIOS clocks set relatively close to the correct GMT time before they join the cluster. All Netmail Store nodes set their clocks relative to GMT, not local time, and they do not change for daylight savings time.*

To operate a cluster without using NTP, set the following in the configuration file of any single node in the cluster:

```
timeSource = system
```

All other nodes in the cluster attempt to synchronize their clocks to this node.