

Installing Netmail Store Content Router

Netmail Store Content Router installs as a service on a standard server, using standard Linux installation and package management utilities as outlined in the sections below. The following steps will ensure that Netmail Store Content Router is properly installed and functions correctly after installation.

Note: If you are installing Netmail Store Content Router on a Cluster Services Node (CSN), the required infrastructure setup, installation and configuration are performed automatically as part of CSN configuration.

On this page:

- [Recommended Content Router Node Hardware Configuration](#)
 - [Hard Disk Selection](#)
- [Operating System](#)
- [Third Party Package Prerequisites](#)
- [Network](#)
 - [Required Communications](#)
 - [Bandwidth Requirements](#)
 - [Recommended Infrastructure](#)
 - [IP Address Configuration](#)
- [Installing Content Router Services](#)
- [Configuring Content Router Services and Rules](#)
- [Removing Content Router](#)
 - [Configuring Publisher](#)
 - [Configuring Replicator](#)

Recommended Content Router Node Hardware Configuration

Although Netmail Store Content Router will run on a variety of x86 hardware, the following are recommended characteristics of a Content Router node:

- **CPU:** 64-bit, Intel Xeon, AMD Athlon64, or equivalent
- **Disk:** Standard SATA, SAS , or SCSI with RAID configuration for redundancy

Approximately 200GB of disk space is needed for every 100 million unique objects in the associated Netmail Store cluster.

- **RAM required for Publisher Service:** Minimum of 4GB; 8GB recommended
- **RAM required for Replicator Service:** Minimum of 2GB; 4GB recommended
- **Network Interface:** Dual Gigabit Ethernet
- **CD Drive:** to install the operating system

Hard Disk Selection

The choice of the hard disks used in a Content Router node can have significant impact on a node's performance as well as the recovery characteristics in the event of a node or disk failure. A disk failure on a Content Router node will not cause data loss, though it may cause a delay in processing remote replication requests. Recovering from a hard disk failure on a Content Router node is as easy as replacing the disk drive and restarting the node.

To support clusters of 20 nodes or more, it is advisable to use a fast disk drive on the Content Router node to maximize throughput and filtering capacity for replication. If a hardware RAID controller is available on a Content Router node with two or more drives attached, it could be configured to use RAID 0, striping across all drives to form a single logical volume. In general, Content Router nodes with faster disk subsystems will be able to support more Netmail Store nodes.

Operating System

Netmail Store Content Router has been developed and tested with 64-bit Red Hat Enterprise Linux (RHEL) 5.5. Other RHEL versions or Linux distributions are not currently supported.

Third Party Package Prerequisites

Netmail Store Content Router has dependencies on Python 2.5, Python Setuptools 0.6c9, and Twisted 8.2.0, which are not available as standard RPM packages under RHEL 5.5 but are included in the Content Router installation package and will be unzipped into the same location as the Content Router RPM. If installing on a Cluster Services Node (CSN), these RPMs will already be installed. If you are installing without a CSN, you may need to manually install the packages using the commands shown below (in this order):

```
rpm -ivh caringo-3rdparty-Python-2.5.4-x86_64.rpm
rpm -ivh caringo-3rdparty-setuptools-0.6c9-x86_64.rpm
rpm -ivh caringo-3rdparty-Twisted-8.2.0-x86_64.rpm
```

Network

The Netmail Store system has been designed to work within standard TCP/IP networking environments. This is achieved through the use of widely supported network services and protocols.

Required Communications

A Content Router node must be able to initiate TCP connections with all nodes in a Netmail Store cluster through the designated access port. This is typically HTTP port 80, but it is configurable by the administrator. Internally, Content Router nodes must be able to communicate with each Netmail Store node via UDP on port 7000 and through IP multicasting. The multicast address must be unique for each cluster and is configurable by the administrator. Multicast communication is only necessary between the Content Router nodes and the Netmail Store storage nodes. All nodes in the Netmail Store cluster, including Content Router nodes, must reside within the same IP subnet.

Bandwidth Requirements

The expected transfer rate for two clusters communicating over a wide area network can be roughly calculated by dividing the minimum effective bandwidth available (e.g., 512 kbps) by the number of volumes in the target cluster. By default, Netmail Store requires a minimum transfer rate of 1000 bytes/sec before timing out connections. The transfer rate is not checked for the initial 20 minutes to provide some latitude for network variability. Administrators anticipating a transfer rate below 1000 bytes/sec between clusters based on the calculation above will need to modify the `minBytesPerSec` parameter in the Netmail Store `node.cfg` file to account for a slower rate and avoid timeouts. For example, to lower Netmail Store's expectations for transfer rate to 512 bytes/sec, the following parameter should be added to `node.cfg` to override the default:

```
minBytesPerSec = 512
```

Recommended Infrastructure

The following networking services are recommended for the management of a Netmail Store cluster that includes Content Router nodes:

- Syslog server for receiving critical alerts
- NTP time server to provide clock synchronization

Gigabit Ethernet is the recommended connection speed between Content Router nodes and Netmail Store cluster nodes. A Content Router node should use the same speed connection as the fastest Netmail Store node it communicates with to prevent bottlenecks.

IP Address Configuration

Content Router Services expect a predictable communication pattern and therefore configuration of one or more static IP addresses for each Content Router server is required. This can be accomplished through either mapping MAC addresses via DHCP or by physically assigning the IP addresses on the server. For mirrored environments, Publisher and Replicator each required their own IP address to ensure communications for the two services do not intermingle. Please reference the Red Hat documentation for information on configuring static IPs. The following is an example of the steps to configure additional alias interfaces in the `/etc/sysconfig/network-scripts` directory:

1. `cd /etc/sysconfig/network-scripts`
2. Copy an existing interface file that has the same characteristics as the new alias interface:

```
cp ifcfg-eth1 ifcfg-eth1:1
```

3. Edit the new file to create the new alias. The essential fields to update are: `DEVICE`, `IPADDR`, and `VLAN`. The following is an example of possible modifications to the file created above for a new `eth1:1` interface:

```
# Public ALIAS bonded interface
```

```
# the device value needs to be unique on the system, and should match the
# name specified in the file name.
DEVICE=eth1:1

# the static ip for the interface
IPADDR=192.168.99.110
NETMASK=255.255.0.0
GATEWAY=192.168.1.1
NETWORK=192.168.0.0
BROADCAST=192.168.255.255

BOOTPROTO=none
ONBOOT=yes
USERCTL=no

# set VLAN to yes for ALIAS network interfaces
VLAN=yes
```

4. Start the new interface:

```
ifup ifcfg-eth1:1
```

Note: If the Content Router services are installed on a CSN, creation of the IP addresses is performed automatically when the services are configured.

Installing Content Router Services

The Content Router distribution is available as a RHEL rpm package that is installed with a shell script. The package and its dependencies must be installed as a 'root' user with the following steps:

1. Copy the distributed zip file to your RHEL server and unzip it. Unzipping the distribution on a separate server and then transferring it to the Content Router server is not recommended as the file may become corrupted during the transfer. The following commands will unzip the file after it has been copied to the server:

```
$ sudo su -
# unzip contentrouter-install-2.2.0b1.zip -d contentrouter-install-2.2.0b1
```

2. Install Content Router by running the shell script from the directory location where the shell script was unzipped. For instance, to install the 2.2 Beta 1 version of the software from the contentrouter-install2.2.0b1 directory created above, you would run the following commands:

```
$ sudo su -
# cd contentrouter-install-2.2.0b1
# ./contentrouter-install.sh
```

Note: If the installation fails due to 'Missing Dependency', make sure you have installed the packages discussed in [Third Party Package Prerequisites](#).

If Content Router is installed on a CSN, the necessary software will be installed automatically as part of the overall software bundle.

Configuring Content Router Services and Rules

After installing Netmail Content Router, you must configure the installed services by updating the configuration files for each service, Replicator and Publisher, you plan to run as well as creating a filter rules file for Publisher. The configuration files do not exist on the system after installation and must be created prior to attempting either service. Failure to create and configure these files will result in a 'Publisher/Replicator not configured' error when you attempt to start the service.

Both services would be utilized for a mirrored configuration where the server is both publishing UUIDs from a local cluster and receiving incoming UUIDs to replicate from a remote cluster. For a one-way Disaster Recovery scenario, you would likely only need to configure one of the services on a single server. Both the Publisher and Replicator services can be configured from the Netmail Administration Console.

Removing Content Router

If you need to remove Netmail Store Content Router from a server, you should first ensure there are no subscribers actively querying the Publisher for streams. The following command will then remove the configuration files and software for Netmail Store Content Router:

```
# sudo /opt/caringo/contentrouter/contentrouter-uninstall.sh
```

Configuring Publisher

A full explanation of all configuration options can be found in the following table:

Option Name	Default	Mandatory	Description
ipaddress	none	yes	The Publisher's IP address.
log	/var/log/caringo/contentrouter-publisher.log	no	Local log file name when not using a syslog.
loghost	localhost	no	Syslog host name or IP address. Content Router services log to the local 5 syslog facility. The parameter should not be set if local file-based logging will be used instead of syslog.
logbackups	8	no	The number of older, rotated log files to keep when file-based logging is used.
logsize	10 * 1024 * 1024 bytes	no	The number of bytes allowed for all file-based log files. Each log file is allocated 1/n of the configured logsize, where 'n' is the configured number of logbackups. A value of 0 will prevent Netmail Store Content Router from rotating logs if an alternate mechanism like logrotate is preferred.
loglevel	40	no	The level of logging verbosity with the following supported values: 50=critical, 40=error, 30=warning, 20=info, 10=debug, 0=no log.
group	225.0.10.100	yes	The multicast group address for the Netmail Store cluster from which UUIDs are gathered. Must be a Class D IP address in the range 224.0.0.0 to 239.255.255.255
scspport	80	no	The SCSP connection port for the Netmail Store cluster listed in the group parameter.
scsphosts	none	yes, if cluster is not specified	A comma-separated list of Netmail Store node IP addresses used to validate version compatibility at boot time. Either scsphosts or cluster must be specified.
cluster	none	yes, if scsphosts is not specified	The name of the source Netmail Store cluster; used to dynamically locate a node and verify version and compatibility. Either scsphosts or cluster must be specified.
castorProxyIp	none	no	IP address for the proxy that provides external Content Router subscribers access to the Netmail Store cluster for replication.

castorProxyPort	none	no	Port for the proxy that provides external Content Router subscribers access to the Netmail Store cluster for replication.
rulesFile	/etc/caringo/contentrouter/rules.xml	no	The name and location of the XML rules file the Publisher uses to filter UUIDs.
consolePort	8090	no	The port for the Publisher web console.
publicationServerPort	80	no	The port Publisher uses to publish UUID data.
storageDir	/var/opt/caringo/contentrouter/publisher	no	<p>A unique, writable directory path for use in storing information about objects in the Netmail Store cluster, as well as additional data required to queue the object to subscribers, such as the Replicator. To prevent errors, the specified location must contain at least enough space defined by the following calculation. The number of objects referred to in the calculation is based on the number of objects processed by the Publisher in its lifetime, including updates to those objects.</p> <p>For every 100 million objects, Publisher requires at least the following amount of disk storage: 26 GB (for non-metadata object information) and an additional 70 MB (0.07 GB) multiplied by the average metadata size per object, in bytes (for compressed object metadata). For example, if Netmail Store objects have 200 bytes of metadata on average, the requirement per 100 million objects is: $26 + 14 = 40$ GB. Note: The size requirement is significantly less if you use only the PublishAll channel in the <i>rules.xml</i> file and do not create a Metadata enumerator. In that case, the requirement is fixed at 20 GB per 100 million objects.</p>
subscriberOfflineAfter	180	no	Time in seconds before the Publisher displays a Subscriber as offline in the Publisher console. Applies only to Subscribers that do not send a offlineAfter parameter at runtime.
subscriberErrOfflineAfter	3600	no	Time in seconds before the Publisher logs a critical error message if a Subscriber has not been heard from. Applies only to Subscribers that do not send an errOfflineAfter parameter at runtime.
subscriberTimeoutInterval	90000	no	Time in seconds before the Publisher terminates a Subscriber if it has not been heard from. Applies only to Subscribers that do not send a Timeout parameter at runtime. Minimum of 9000.
snmpCommunity	ourpwdofchoicehere	no	The password used to control access to restricted capabilities in the Publisher admin console.
consoleReportStyleURL	none	no	Provides the location of an override stylesheet for overriding style types in the admin console stylesheet.
consoleStyleURL	none	no	Provides the location of an override stylesheet for overriding style types in the

admin console stylesheet.

errorRetentionDays	4	no	The number of days the Publisher should retry failed attempts to read stream metadata before the event is dropped.
enumeratorDefaultMaxItems	5000	no	Determines the number of events that are retrieved by an enumerator from the Publisher per request. The configured parameter can be overridden by an individual enumerator using the maxItems query argument on a Start or Next request. Replicator throughput can be increased by increasing the value of this parameter.
publicationServerStrictArgsChecking	False	no	Determines whether Enumerator API query argument syntax is strictly enforced on requests to the Publisher. By default, invalid arguments and values are tolerated by dropping the value and sending a warning message to the log. This enables legacy applications to continue to function. If set to "True", publication server query argument checking is strictly enforced, with a 400 (error) returned on the request if an invalid argument or value is given on a request. This checking is helpful when developing new subscribers as it provides immediate feedback on error conditions.

Configuring Replicator

A full explanation of all configuration options can be found in the following table:

Option Name	Default	Mandatory	Description
ipaddress	none	yes	The Replicator's IP address.
log	/var/log/caringo/contentrouter-replicator.log	no	Local log file name when not using a syslog.
loghost	localhost	no	Syslog host name or IP address. Netmail Store Content Router services log to the local 5 syslog facility. The parameter should not be set if local file-based logging will be used instead of syslog.
loglevel	20	no	The level of logging verbosity with the following supported values: 50=critical, 40=error, 30=warning, 20=info, 10=debug, 0=no log.
logbackups	8	no	The number of older, rotated log files to keep when file-based logging is used.
logsize	10 * 1024 * 1024 bytes	no	The number of bytes allowed for all filebased log files. Each log file is allocated 1/n of the configured logsize, where 'n' is the configured number of logbackups. A value of 0 will prevent Netmail Store Content Router from rotating logs if an alternate mechanism like logrotate is preferred.
group	225.0.10.100	yes	The multicast group address for the Netmail Store cluster to which UUIs should be replicated. Must be a Class D IP address in the range 224.0.0.0 to 239.255.255.255.

scsphosts	none	yes, if cluster is not specified	A comma-separated list of Netmail Store node IP addresses used to validate version compatibility at boot time. Either scsphosts or cluster must be specified.
cluster	none	yes of scsphosts is not specified	The name of the target Netmail Store cluster; used to dynamically locate a node and verify version compatibility at boot. Either scsphosts or cluster must be specified.
castorProxy IP	none	no	IP address of a proxy configured in front of the target cluster.
castorProxyPort	none	no	Port of a proxy configured in front of the target cluster.
subscribeTo	none	yes	Specifies one or more Publishers to query for UUIDs. Syntax of <channelName> @<host>:<port>, <channelName> @<host>:<port>,etc (one Name, host, port group for each Publisher the Replicator subscribes to).
subscriptionCheckInterval	10	no	Time in seconds between checks for new UUIDs; values can be between 5 - 3600 seconds.
offlineAfter	120	no	Time in seconds before the Publisher displays Replicator as offline in the Publisher console. Minimum of 60 seconds.
errOfflineAfter	1800	no	Time in seconds before the Publisher logs a critical error message if a Subscriber has not been heard from. Minimum of 60 seconds.
timeoutInterval	90000	no	Time in seconds before the Publisher will terminate a Subscriber if it has not been heard from. Minimum of 9000 seconds.
consolePort	8088	no	Replicator responds to requests for its state on this port. There is not currently a separate console for Replicator.
storageDir	/var/opt/caringo/contentrouter/replicator	no	A unique, writeable directory in which to store Replicator data.
maxActiveEvents	20	no	The number of events Replicator should process at one time. Messaging Architects recommends using the following settings: <ul style="list-style-type: none"> • If there are 10 or fewer nodes in the target cluster, leave this parameter at its default of 20. • Set to 10 if there are more than 10 nodes in the target cluster. • Set to 5 if there are more than 20 nodes in the target cluster.
ignoreDeleteEvents	0	no	Determines whether or not a Replicator should process delete events. By default, delete events are processed (ignoreDeleteEvents=0)
