

Oracle® Communications Messaging Server

Installation and Configuration Guide for Cassandra Message
Store

Release 8.0.2

E79615-01

October 2017

Copyright © 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	vii
Audience	vii
Related Documents	vii
Documentation Accessibility	viii
 1 Cassandra Message Store Installation and Configuration Overview	
Overview of Cassandra Message Store Installed Software	1-1
Overview of the Message Store Installation Procedure	1-1
Messaging Server Installation Options	1-2
Directory Placeholders Used in This Guide	1-2
 2 Planning Your Cassandra Message Store Installation	
About Cassandra and DataStax	2-1
About the Cassandra Message Store	2-1
About Solr Indexing and Search	2-2
Planning Your Cassandra Message Store Installation	2-2
System Deployment Planning	2-3
Planning for High Availability	2-3
Planning for Disaster Recovery	2-3
Using Load Balancing	2-3
Planning Backup Strategies	2-3
About Installing a Secure System	2-4
 3 Developing a Cassandra Message Store Architecture	
Understanding the Cassandra Message Store Architecture	3-1
Overview of Cassandra Message Store Logical Architecture	3-2
Deploying the Cassandra Message Store on Data Centers	3-3
Planning for a Highly Available Cassandra Message Store Deployment	3-6
 4 Cassandra Message Store System Requirements	
Software Requirements	4-1
Supported Operating Systems	4-1
Required Software	4-1
Hardware Requirements	4-2

Time Synchronization Requirements	4-2
Information Requirements	4-3
Connection Information	4-3
Data Center and Replication Factor Requirements	4-3
Store Affinity Group Requirements	4-3

5 Cassandra Message Store Pre-Installation Tasks

Summary of General Pre-Installation Tasks	5-1
Installing Java	5-1
Installing Python	5-2
Installing DataStax Enterprise	5-2
Downloading the DataStax Enterprise Software	5-2
Installing the DataStax Enterprise Software	5-2
Setting Up the Cassandra Cluster	5-3
Changing Initial Cassandra Settings	5-3
Changing Initial Tuning Settings	5-3
Linux Tuning Settings	5-3
dse File	5-4
dse.yaml File	5-4
cassandra.yaml File	5-4
cassandra-env.sh File	5-5
jvm.options File	5-5
cassandra-rackdc.properties File	5-6
Changing Initial Solr Tunings	5-6

6 Installing Cassandra Message Store

About Messaging Server Components	6-1
About Unified Configuration	6-2
Before Installing Messaging Server	6-2
Using a Configuration Management Tool to Simplify Installation and Configuration	6-2
Installing and Configuring Messaging Server	6-3
Downloading the Messaging Server Software	6-3
Installing and Configuring the Messaging Server Software	6-3
Installing Messaging-CAS Software on DSE/Solr Nodes	6-3
Installing Message Access Tier Hosts	6-4
Installing MTA Tier Hosts	6-5
Installing MMP Hosts	6-6
Using a Non-MMP IMAP Proxy	6-7
Configuring a Remote Indexed Search Converter	6-7
Enabling SSL	6-7
Starting the Cassandra Message Store Across Tiers	6-7
Restarting a Cassandra Cluster	6-7

7 Cassandra Message Store Post-Installation Tasks

Verifying the Installation	7-1
Scaling Your Cassandra Message Store Deployment Horizontally	7-1

8 Migrating to Cassandra Message Store

Overview of Migrating to Cassandra Message Store	8-1
--	-----

9 Uninstalling the Cassandra Message Store

Uninstalling Cassandra Database	9-1
---------------------------------------	-----

Uninstalling Message Server Software	9-1
--	-----

A config-fit Reference

Overview of the config-fit Script	A-1
---	-----

Syntax.....	A-1
-------------	-----

Preface

This guide provides instructions for installing and configuring Oracle Communications Messaging Server Cassandra message store. Where appropriate, this guide refers to *Messaging Server Installation and Configuration Guide* for concepts and tasks pertaining to the overall installation of Messaging Server software.

Audience

This document is intended for system administrators or software technicians who install and configure Messaging Server. This guide assumes you are familiar with the following topics:

- Messaging protocols, such as IMAP and SMTP
- Oracle Directory Server Enterprise Edition and LDAP
- System administration and networking
- DataStax Enterprise/Cassandra

Related Documents

For more information, see the following documents in the Messaging Server documentation set:

- *Messaging Server Installation and Configuration Guide*: Provides instructions for installing and configuring Messaging Server.
- *Messaging Server System Administrator's Guide*: Provides instructions for administering Messaging Server.
- *Messaging Server Reference*: Provides additional information for using and configuring Messaging Server.
- *Messaging Server Release Notes*: Describes the fixes, known issues, troubleshooting tips, and required third-party products and licensing.
- *Messaging Server Security Guide*: Provides guidelines and recommendations for setting up Messaging Server in a secure configuration.
- *Messaging Server MTA Developer's Reference*: Describes the Messaging Server MTA SDK.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Cassandra Message Store Installation and Configuration Overview

This chapter provides an overview of the Oracle Communications Messaging Server Cassandra message store installation and configuration process.

Overview of Cassandra Message Store Installed Software

During the Cassandra message store installation process, you install and configure the following software:

- On Cassandra/Solr nodes:
 - Oracle Java SE Runtime Environment (JRE) or OpenJDK
If you are using SSL and need to use the Java **keytool** utility, you must install Java SE Development Kit (JDK).
 - Python
 - DataStax Enterprise (Cassandra database with Solr integration)
 - Messaging-CAS software (which contains the Field Input Transformer (FIT) component)
- On Messaging Server message access tier, MTA, and MMP hosts:
 - Messaging Server software
 - Optional: Java is required on message access tier hosts if they also run the Indexed Search Converter (ISC).

See *Messaging Server Installation and Configuration Guide* for information about additional Messaging Server components that you must install and configure, and other software dependencies.

Note: You cannot run both Cassandra message store and classic message store in the same Messaging Server deployment. That is, you cannot mix Cassandra and classic message stores in the same deployment.

Overview of the Message Store Installation Procedure

The installation procedure follows these steps:

1. Plan your installation. When planning your installation, do the following:

- Determine the scale of your implementation, for example, a small development system, or a large production system.
 - Determine how many physical machines you need, and which software components to install on each machine.
 - Plan the system topology, for example, how the system components connect to each other over the network.
2. Review system requirements. System requirements include:
 - Hardware requirements, such as disk space.
 - System software requirements, such as operating system (OS) versions and OS patch requirements.
 - Information requirements, such as IP addresses and host names.
 3. Install and configure DataStax Enterprise.
 4. Install and configure Messaging Server and configure the new message store.
 5. Perform post-installation configuration tasks.
 6. Verify the installation.
 7. Continue with the rest of your Messaging Server installation.

Messaging Server Installation Options

You install Messaging Server software by running the installer in either interactive or silent mode. When you run the installer in silent mode, you are running a non-interactive session. The installation inputs are taken from the following sources:

- A silent installation file
- Command-line arguments
- Default settings

You can use silent mode to install multiple instances of the same software component and configuration without having to manually run an interactive installation for each instance.

For more information, see the discussion on running a non-interactive session in *Messaging Server Installation and Configuration Guide*.

Directory Placeholders Used in This Guide

[Table 1–1](#) lists the placeholders that are used in this guide:

Table 1–1 Messaging Server Directory Placeholders

Placeholder	Directory
<i>MessagingServer_home</i>	Specifies the installation location for the Messaging Server software. The default is /opt/sun/comms/messaging64 .
<i>InstallRoot</i>	Specifies the installation location for other Communications Suite software. The default is /opt/sun/comms .
<i>ConfigRoot</i>	Specifies the location of the configuration files. The default is /var/opt/sun/comms/messaging64/config .

Planning Your Cassandra Message Store Installation

This chapter describes how to plan for installing the Oracle Communications Messaging Server Cassandra message store.

About Cassandra and DataStax

Apache Cassandra is a free and open-source NoSQL database designed for:

- Large amounts of data
- Use on commodity Linux servers
- Continuous read/write availability
- High availability with no single point of failure
- Multiple data centers to ensure constant uptime

DataStax is a company that provides commercial support for an enterprise, production-certified edition of Cassandra known as DataStax Enterprise (DSE). You must use DSE as the database for the Cassandra message store.

DataStax also integrates the Apache Solr indexing engine into DSE. Cassandra message store requires DSE Solr integration to provide fast IMAP SEARCH functionality.

In addition, DSE provides a set of tools for Cassandra, including monitoring and analytics, to help you manage your deployment.

For more information, see the topic about Apache Cassandra at:

<http://docs.datastax.com/en/cassandra/3.0/cassandra/cassandraAbout.html>

About the Cassandra Message Store

In Messaging Server, the Cassandra message store consists of *keyspaces*. A keyspace, in a NoSQL data store, is a namespace that defines data replication on nodes. (It resembles the schema concept in Relational database management systems.) The Messaging Server Cassandra message store consists of the following keyspaces:

1. **ms_msg**: Contains the email message content. It can be very large.
2. **ms_mbox**: Contains user and mailbox metadata. It is relatively small and has a lot of mutations.
3. **ms_index**: Contains Cassandra tables for Solr indexing.

4. **ms_cache**: Contains cache tables for converted message content. The data in these tables is for Solr use only, and enables the Field Input Transformer (FIT) to use cached data instead of converting large message content for each update operation.

After Messaging Server has been installed and configured, the Cassandra message store keyspaces are created on one or more data centers once the Messaging Server stored processes start for the first time. When deploying the Cassandra message store, the recommendation is to use multiple Cassandra clusters. See ["Understanding the Cassandra Message Store Architecture"](#) for more information.

About Solr Indexing and Search

DSE is the database platform hosting the Cassandra message store and Solr is the indexing and search engine. DataStax bundles Solr with Cassandra in DSE. Solr replaces the need to install Oracle Communications Indexing and Search Service as the indexing and search engine for Messaging Server.

The following two components are Oracle extensions to Solr, which is open source software. These two extensions are packaged and with Messaging Server:

- **Field Input Transformer (FIT)**: Deployed to DSE as a Solr plugin, the FIT generates dynamic Solr fields for a message's header and body parts to be indexed in Solr.
- **Indexed Search Converter (ISC)**: Deployed as a standalone web application, the ISC processes the entire message body—headers, attachments, and plain text body—and generates plain text content for each of the message parts. It stores this converted content in a persistent Cassandra cache table, which is later consumed by the FIT to be indexed in Solr. The cache enables the ISC to avoid processing the same contents and attachments multiple times (for example, when there are multiple recipients, email forwards, and so on).

You must deploy the ISC on the message access tier hosts, and the FIT on each DSE/Solr node. When configuring the FIT on the DSE/Solr nodes, you specify the access tier host(s) running the ISC to the **fit.ischosts** option.

Email messages have two paths to the Cassandra/Solr message store:

1. Messages that arrive through LMTP are processed by ISC in the message access layer. The ISC pre-converts the header, body and attachments to plain text and stores them in cache. This is also referred to as "pre-conversion." The FIT then consumes this data and indexes it in Solr.
2. Messages that arrive in the store through the IMAP **APPEND** command or the **imsrestore** command are sent by FIT to be processed by ISC on the message access tier host(s) (determined by the **fit.ischosts** option set on the DSE/Solr node). In this scenario, the FIT is "triggered" twice:
 - a. In the first trigger, all email fields, except the body and header, are immediately indexed. The body conversion and indexing happen asynchronously.
 - b. The second trigger happens after the message is converted and cached.

Planning Your Cassandra Message Store Installation

See the topic on planning your Messaging Server installation in *Messaging Server Installation and Configuration Guide* for general planning information. You must consider the following items before installing Cassandra message store:

- **Cassandra software:** You must purchase DSE as the database for the Cassandra message store separately from DataStax. DSE is not included in the Messaging Server distribution.
- **Operating system:** You must install Cassandra nodes on Linux. Solaris OS is not supported for Cassandra nodes.
- **Messaging Server version:** Cassandra nodes and message access tier hosts must run Messaging Server 8.0.2. Front-end hosts must run at least Messaging Server 8.0.1. Additionally, you must use Messaging Server Unified Configuration to administer message access tier hosts. Legacy configuration is not supported on these hosts.
- **Store protocol:** You must use Local Mail Transfer Protocol (LMTP) on the message access layer to deliver messages into the Cassandra message store.
- **Notifications:** You must use Event Notification Service (ENS) for notifications on the message access tier hosts. (Support for Java Message Queue (JMQ) is deprecated.)
- **Store delivery:** You must configure store affinity groups by using the **proxy:mailhost.storehostlist** Unified Configuration option for robust Cassandra message store delivery.
- **Message store storage:** Use of solid-state drives (SSDs) is recommended for nodes running DSE/Cassandra. For actual storage requirements, see "[Hardware Requirements](#)".

System Deployment Planning

This section contains system-level planning topics you must consider before installing Cassandra message store:

Planning for High Availability

Cassandra database is designed to be highly available with no single point of failure. Thus, you do not need to install additional cluster software to create a high availability solution for the message store itself.

Planning for Disaster Recovery

When planning for disaster recovery, consider deploying Cassandra across multiple data centers in different geographic locations. Cassandra enables replicating and synchronizing data across multiple data centers. You can use different data centers as a live backup, or as a failover site.

Using Load Balancing

Load balancers balance network connections uniformly or by algorithm across multiple servers. You cannot use load balancers on the message store or directory masters. You can use them for connections to MMPs, Oracle Convergence, MTAs, directory consumers, and with the MTA's use of anti-virus and anti-spam products.

Planning Backup Strategies

Cassandra uses data replication to store copies (replicas) of data on multiple nodes to ensure that data is not lost if the system crashes, hardware fails, and so on. Nevertheless, backing up and restoring data is an important administrative task for

your Messaging Server deployment. To back up the Cassandra database, use the snapshot feature, which can back up all keyspaces, a single keyspace, or an entire cluster. For more information, see the Cassandra backup and restore documentation at:

<https://docs.datastax.com/en/cassandra/3.0/cassandra/operations/opsBackupRestore.html?hl=backup>

Also, to back up DataStax OpCenter (DataStax's management and monitoring solution), see the OpCenter Backup Service documentation at:

https://docs.datastax.com/en/opscenter/6.0/opsc/online_help/services/opscBackupService.html

About Installing a Secure System

You can secure your Messaging Server infrastructure by first determining your firewall or DMZ architecture. You can also configure Secure Sockets Layer (SSL) on Messaging Server front-end hosts. You must also protect Messaging Server individual components.

For Cassandra message store, keep the following security items in mind:

- Messaging Server does not use SSL or authentication for the Cassandra Query Language (CQL) connections from the message access tier (the IMAP/LMTP-server/store-CLI-tools) to Cassandra. Instead, keeping those hosts together on a physically secure network that blocks external access to CQL is recommended.
- To secure data at rest on Cassandra nodes, refer to Datastax for more information.
- You can optionally secure the connection from the LMTP server to the Indexing Search Convertor (ISC) with SSL.
- ENS can be secured by using SSL and authentication support that was introduced in Messaging Server 8.0.1.

For more information, see the discussion on secure installation and configuration of Messaging Server in the *Messaging Server Security Guide*.

Developing a Cassandra Message Store Architecture

This chapter provides information about how to design your Oracle Communications Messaging Server Cassandra message store architecture, and information about how Messaging Server components are distributed across hardware and software resources. For more information about general concepts, see the "Developing a Messaging Server Architecture" chapter in *Messaging Server Installation and Configuration Guide*.

Understanding the Cassandra Message Store Architecture

In general, you never deploy the Cassandra message store on a single node. Instead, you deploy the Cassandra message store on multiple nodes in a *data center*. In Cassandra, the term *data center* is a grouping of nodes, configured together for replication purposes. (A *cluster* contains one or more data centers, which can span physical locations.)

A Cassandra node can be installed on a Docker container, a virtual machine (VM), or a physical host. Using containers provides more flexibility with your deployment than using physical hosts. For example, if a physical host contains extra bandwidth, you can increase throughput by installing more Cassandra instances within containers.

In a Messaging Server deployment, containers are recommended for installing Cassandra. However, if a physical host can support only a single Cassandra instance, do not use a container.

Each node can read and write messages that are sent to the data center. One node is chosen to be *coordinator* for each read or write operation, and forwards the read or write operation to appropriate *replicas* for the operation's partition key. Cassandra uses token aware routing, so the coordinator for an operation usually holds one copy of the relevant data.

Cassandra uses replicas on multiple nodes to ensure data reliability and fault tolerance. A *replication factor* (on a per-keyspace basis) determines the number of replicas. A replication factor of 1 means that there is only one copy of each row on one node. A replication factor of 2 means two copies of each row, where each copy is on a different node. When you install the Cassandra message store, Messaging Server automatically sets the replication factor on the Cassandra keyspaces. You can control the replication factor through Messaging Server configuration options. For more information, see the replication factor options in [Table 4-3, "Data Center and Replication Configuration Options"](#).

Internode communications between nodes occurs through a peer-to-peer protocol called *gossip*. Through gossip, nodes periodically exchange state information about

themselves and about other nodes they know about. To facilitate gossip communications and bootstrap nodes in a data center, you designate some nodes as *seed nodes*. The seed node acts as a point of contact for other nodes to find each other and learn the topology of the data center. In addition, it is recommended to configure *snitches*, which determine which data centers and racks the nodes belong to. It is recommended to configure the DataStax option **GossipingPropertyFileSnitch** as **endpoint_snitch**. Snitches inform Cassandra about the network topology so that requests are routed efficiently and enable Cassandra to distribute replicas by grouping machines into data centers and racks. You also use Message Server data center and replication factor options, to configure racks and data centers. For more information, see Table 4–3, "Data Center and Replication Configuration Options".

For more information about Cassandra architecture, see DataStax Enterprise Reference Architecture at:

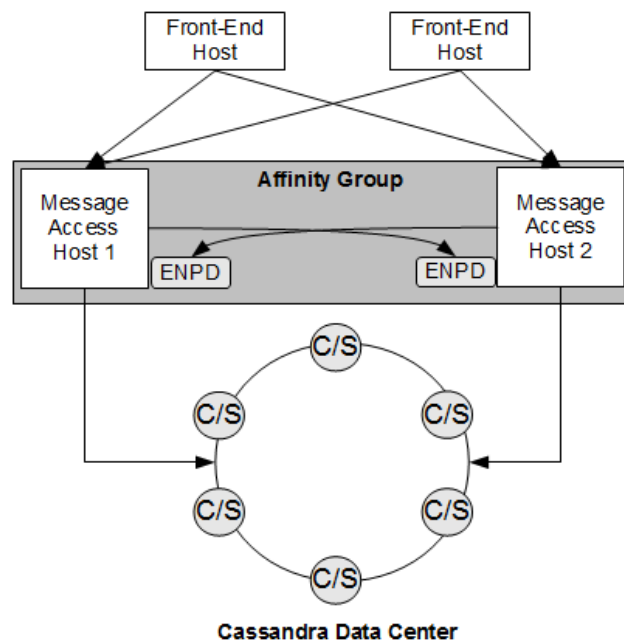
<http://www.datastax.com/wp-content/uploads/2014/01/WP-DataStax-Enterprise-Reference-Architecture.pdf>

Overview of Cassandra Message Store Logical Architecture

In general terms, the message store holds and stores user mail. Sometimes it is referred to as a "back end." The message store also refers to the Message Access Components including the IMAP server and the POP server. Starting with Messaging Server 8.0.2, the Message Access Components have been separated out from the message store into their own logical layer, referred to as the message access layer.

The original Messaging Server message store is based on storing mailbox caches, mailbox indexes and messages in the file system, and use of Oracle Berkeley Database (BDB) for other metadata. The Cassandra message store, based on Cassandra database, changes the approach to how you architect the store components of your Messaging Server deployment. At a high-level, Figure 3–1 shows the Cassandra message store logical architecture.

Figure 3–1 *Cassandra Message Store Logical Architecture*



The preceding figure shows the following components of the Cassandra message store logical architecture:

- **Front-end hosts:** The front-end hosts, which can be Message Transfer Agent (MTA), Messaging Multiplexor (MMP), LMTP client, or Webmail (**mshttpd**) hosts, communicate with the message access hosts.
- **Message access hosts:** The message access hosts run the **imapd** process and LMTP server. The message access hosts also run the Event Notification Service (ENS) **enpd** process for horizontal scalability in the form of a store affinity group. Store affinity groups enable routing to an alternate LMTP server in the list if the first one listed is unavailable. Store notifications relevant to the store affinity group are published to all ENPD servers in the store affinity group. Other connections to a store affinity group (POP, IMAP, LMTP server) can be handled by any host in the store affinity group.

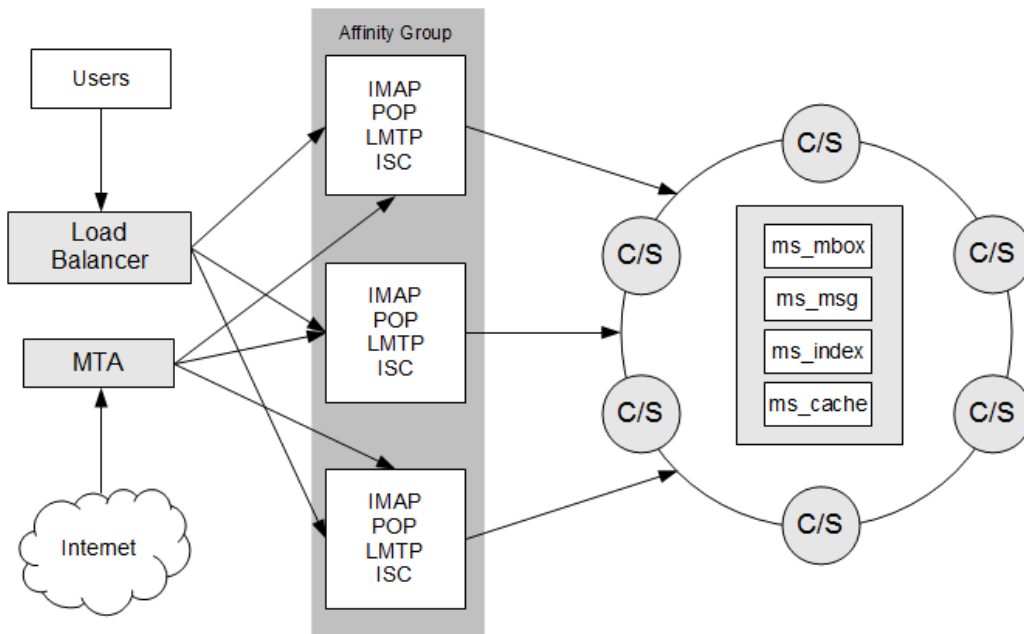
Note: Only two message access hosts are required in a store affinity group. This differs from classic store automatic failover, which requires three hosts in a store affinity group.

The message access hosts also run the Indexed Search Converter (ISC) process, which pre-converts messages before delivery to the message store.

- **Cassandra data center:** The data center is a collection of Cassandra and Solr hosts, which store and index the email data. The Solr core consists of the **ms_index** index/schema and Field Input Transformer (FIT), a plugin that generates dynamic Solr fields to index the message header and body contents.

Deploying the Cassandra Message Store on Data Centers

Figure 3–2 shows the logical architecture of a Cassandra message store on a single data center.

Figure 3–2 Message Store Deployed Upon a Single Cassandra Data Center

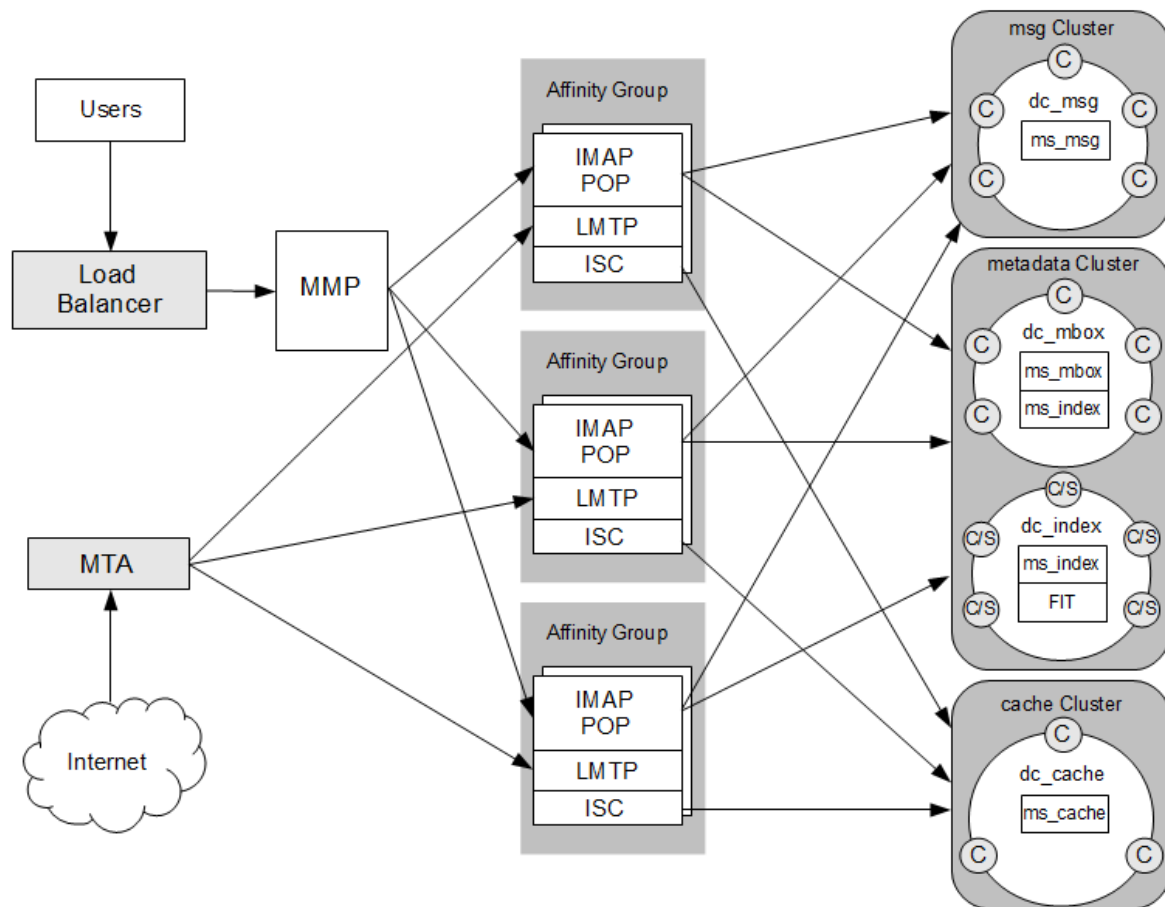
This figure shows a Messaging Server Cassandra message store deployment that consists of a load balancer; a Message Transfer Agent (MTA) host for incoming traffic; multiple Messaging Server message access hosts in a single affinity group; and a single Cassandra data center. Each Cassandra/Solr node (indicated by a "C/S" in the figure) in the data center contains the four keyspaces that make up the message store and indexing and search component: **ms_mbox**, **ms_msg**, **ms_index**, and **ms_cache**.

Note: See [Figure 3–1, "Cassandra Message Store Logical Architecture"](#) for more details on the logical architecture of the message access tier and ENS, which is not depicted in this figure.

When you perform the initial configuration of Messaging Server, or when the service startup of the **stored** process happens when in Cassandra message store mode, by default, it creates all four keyspaces on a Cassandra/Solr node on one data center called **DC1**.

In general, use a single data center for a proof of concept deployment, and not for production purposes. While you can use a single store affinity group deployment, the scale limits of a single-affinity deployment are unknown. In the event ENS notifications become a scale bottleneck, the recommendation is to add a new store affinity group. When there is more than one store affinity group, you must have an MMP between the load balancer and the IMAP/LMTP/ISC tier.

[Figure 3–3](#) shows the logical architecture of a Cassandra message store on multiple data centers and clusters.

Figure 3–3 Message Store Deployed Upon Multiple Cassandra Clusters

This figure shows a Messaging Server Cassandra message store deployment that consists of a load balancer; a Message Transfer Agent (MTA) host for incoming traffic; a Messaging Multiplexor (MMP) host to act as the connection point to the multiple back-end message stores; multiple Messaging Server front-end hosts, configured as three store affinity groups (and running ENS), for message access and delivery handling; and three Cassandra clusters. The clusters contain at least one Cassandra data center, indicated by "**dc_name**." Each Cassandra node is indicated by "C" in the figure, and each Cassandra/Solr node is indicated by "C/S." The **dc_msg** data center contains the **ms_msg** keyspace. The **dc_mbox** data center contains the **ms_mbox** and **ms_index** keyspaces. The **dc_index** data center contains the **ms_index** keyspace and Solr components. The **dc_cache** data center contains the **ms_cache** keyspace. Because the cache data can be regenerated, only a single **ms_cache** keyspace is needed; disaster recovery is unnecessary. Also, the **ms_cache** keyspace is only needed by the Indexed Search Converter (ISC). The Cassandra replication factor ensures there are *N* copies of data distributed across the nodes in the data center.

Note: See Figure 3–1, "Cassandra Message Store Logical Architecture" for more details on the logical architecture of the message access tier and ENS, which is not depicted in this figure.

Planning for a Highly Available Cassandra Message Store Deployment

Unlike previous versions of Messaging Server, where you need to install and manage clustering software to obtain a highly available message store, the Cassandra message store is itself designed to provide continuous uptime. Cassandra relies the concept of installing multiple nodes with no single point of failure. Because of its peer-to-peer distributed system, where data is handled by all nodes in a cluster, Cassandra is able to address the problem of failures. Thus, the need to use additional clustering software is rendered unnecessary (nor is it supported).

At a high-level, Cassandra operates in the following way to ensure data consistency across all nodes in a cluster:

1. Uses a peer-to-peer protocol called gossip for nodes to exchange state information.
2. Ensures data durability through use of a sequentially written commit log on each node.
3. Indexes and writes data to an in-memory structure, called a memtable, which is similar to a write-back cache.
4. When the in-memory structure becomes full, Cassandra writes the data to disk. Cassandra automatically partitions and replicates the writes throughout the cluster.

For more information on Cassandra architecture, see the DataStax documentation at:

<http://docs.datastax.com/en/cassandra/3.0/cassandra/architecture/archTOC.html>

High availability is achieved on the Messaging Server front-end tier (Message Transfer Agent (MTA), Messaging Multiplexor (MMP), and **mshttpd** daemon) through the use of load balancers, or MX records (MTA-submit). High availability is achieved on the message access tier (IMAP, POP, LMTP server, ENPD, and ISC) with store affinity groups. High availability on the back-end tier (Cassandra/DSE, FIT and ISC) is achieved through Cassandra/DSE itself.

You must configure store affinity groups consistently on all front-end tier and message access tier hosts. Store affinity groups enable the MMP to switch to an alternate IMAP server that can perform read/write IMAP operations for the user if one of the message access tier hosts goes down. Classic message store only supports failover (with read-only on backup IMAP servers), while Cassandra store supports read/write for all IMAP servers in the affinity group.

Store notifications are broadcast to all **enpd** servers in the store affinity group (which then relays them to all subscribed IMAP servers in the affinity group). Without store affinity groups, all store notifications would have to be broadcast to all *enpd* servers in the deployment, which would eventually pose a scalability problem.

For more information on configuring store affinity groups, see "[Installing and Configuring the Messaging Server Software](#)".

Cassandra Message Store System Requirements

This chapter describes the software, hardware, and operating system requirements for installing Oracle Communications Messaging Server Cassandra message store hosts. See *Messaging Server Installation and Configuration Guide* for the software, hardware, and operating system requirements for Messaging Server front-end and message access tier hosts.

Software Requirements

This section describes the software requirements for installing Cassandra message store.

Supported Operating Systems

Table 4–1 lists operating systems that support Cassandra message store. We recommend that you run the latest software update with the latest recommended patch set.

Table 4–1 Supported Operating Systems

Operating System	CPU	Required Patches
Oracle Linux 7.2 Red Hat Linux 7.2	x64	See the Oracle Linux and Red Hat Enterprise Linux documentation for patch information.

Required Software

Table 4–2 lists software required for installing and running Cassandra message store. See *Messaging Server Installation and Configuration Guide* for other requirements related to Messaging Server itself.

Table 4–2 Software Requirements

Product	Version	Notes
Oracle Java SE Runtime Environment 8 (JRE) or OpenJDK 8	1.8.0_40 minimum	Required to install DataStax Enterprise on Cassandra nodes. Required on message access tier hosts if they also run the Indexed Search Converter (ISC). If you are using SSL and need to use the Java keytool utility, you must install Java SE Development Kit (JDK).

Table 4–2 (Cont.) Software Requirements

Product	Version	Notes
Python	2.7.x	Required to install DataStax Enterprise on Cassandra hosts.
DataStax Enterprise Max	5.1.0	You must obtain a license from DataStax.

Hardware Requirements

Cassandra node storage requirements are:

- Message data center nodes: 5 TB when using SSD (utilization should be less than 70 percent)
- Solr index data center nodes: 300 GB (utilization should be less than 60 percent)

Additionally, see the table on hardware requirements for installing Cassandra in the DataStax documentation at:

https://docs.datastax.com/en/datastax_enterprise/5.0/datastax_enterprise/install/installRHELdse.html?hl=hardware,requirements

The number and configuration of the systems that you employ for your Messaging Server installation depends on the scale and the type of deployment you have planned.

Cassandra nodes that are configured with the **ms_msg** keyspace require the most disk capacity to store the message content. Cassandra nodes configured with the **ms_cache** keyspace collectively should have enough disk capacity for 14 days of converted data stored in the conversion cache.

Note: The sizing estimates in this section assume proper application configuration and tuning, in a manner consistent with leading practices of Oracle Communications consulting and performance engineering. This information is provided for informational purposes only and is not intended to be, nor shall it be construed as a commitment to deliver Oracle programs or services. This document shall not form the basis for any type of binding representation by Oracle and shall not be construed as containing express or implied warranties of any kind. You understand that information contained in this document will not be a part of any agreement for Oracle programs and services. Business parameters and operating environments vary substantially from customer to customer and as such not all factors, which may impact sizing, have been accounted for in this documentation.

Time Synchronization Requirements

Configure Network Time Protocol (NTP) to ensure that time is synchronized across your deployment.

Cassandra nodes must synchronize their clocks between all hosts accessing Cassandra. All the Cassandra hosts and all the IMAP/LMTP-server tier hosts must have fully synchronized clocks or the deployment becomes unsupportable (both Cassandra-level errors and application errors related to IMAP QRESYNC and CONDSTORE [RFC 7162] will result).

Information Requirements

This section describes the information needed before installing and configuring Cassandra message store.

Connection Information

You need the IP addresses of the Cassandra nodes that act as **cassconnectpoints**, and the IP addresses of the Solr nodes that act as **solrconnectpoints**, when installing and configuring the message access tier nodes.

Data Center and Replication Factor Requirements

Table 4–3 shows the Messaging Server options you use to configure racks and data centers when installing and configuring message access tier hosts.

Table 4–3 Data Center and Replication Configuration Options

Keyspace	Data Center Option	Replication Factor Option
ms_cache	store.cascachedc	store.cascacherf
ms_index	store.cassolrdc	store.cassolrrf
ms_mbox	store.casmetadc	store.casmetarf
ms_msg	store.casmsgdc	store.casmsgrf

Store Affinity Group Requirements

You must understand the store affinity group layout, and how each LDAP **mailHost** value translates to two or more access tier hosts. For more information, see the topic on the **storehostlist** option in *Messaging Server Reference*.

Cassandra Message Store Pre-Installation Tasks

This chapter provides information on the pre-installation tasks you must complete on Cassandra nodes before you can install Messaging Server software.

Summary of General Pre-Installation Tasks

The following list summarizes the general pre-installation tasks you must complete before installing any Messaging Server component.

- Create a UNIX system user and group for Messaging Server, and set permissions for the directories and files owned by that user.
- Check that DNS is running and configured properly for the Messaging Server host.
- Check the number of file descriptors for Linux, and if this number is less than 16384, you need to increase the value.
- Install Oracle Directory Server Enterprise Edition, if your site does not currently have Directory Server deployed.

See the chapter titled "Messaging Server Pre-Installation Tasks" in *Messaging Server Installation and Configuration Guide* for detailed information.

The following list summarizes the pre-installation tasks you must complete on Cassandra nodes:

- [Installing Java](#)
- [Installing Python](#)
- [Installing DataStax Enterprise](#)

Installing Java

To install Java, see "Installing Oracle JDK on RHEL-based Systems" on the DataStax web site at:

http://docs.datastax.com/en/dse/5.1/dse-admin/datastax_enterprise/install/installJdkRHEL.html

Note: The `JAVA_HOME/bin` directory must be in the `PATH` environment variable.

Installing Python

To install Python, see the Python documentation at:

<https://docs.python.org/2/installing/>

Be sure to use the version of Python that is supported by the version of DataStax Enterprise Max that you are installing.

Installing DataStax Enterprise

The tasks to install DataStax Enterprise are:

- [Downloading the DataStax Enterprise Software](#)
- [Installing the DataStax Enterprise Software](#)
- [Setting Up the Cassandra Cluster](#)
- [Changing Initial Cassandra Settings](#)
- [Changing Initial Tuning Settings](#)
- [Changing Initial Solr Tunings](#)

Downloading the DataStax Enterprise Software

To download the DataStax Enterprise Max software:

1. Register with DataStax and download the DSE software from the DataStax download site, located at:
<https://academy.datastax.com/downloads>
2. Copy the installer file to your Cassandra message store hosts.

Installing the DataStax Enterprise Software

To install DataStax Enterprise software:

1. On each Cassandra/Solr node, you configure a **datastax.repo** file, install the DataStax Enterprise packages, start the DSE software, and verify that DSE is running.

For more information, see the DataStax Enterprise installation documentation at:

http://docs.datastax.com/en/dse/5.1/dse-admin/datastax_enterprise/install/installTOC.html

2. On all Solr nodes, enable Solr by setting the following option in the **/etc/default/dse** file:

```
SOLR_ENABLED=1
```

3. Ensure that for Oracle Linux 6.x and later, the 32-bit versions of the **glibc** libraries are installed.

For more information, see the DataStax Enterprise documentation at:

https://docs.datastax.com/en/datastax_enterprise/4.8/datastax_enterprise/install/installDseInstallGlibc.html

4. Optionally, install OpsCenter, a visual management and monitoring solution for DataStax Enterprise. For more information, see the OpsCenter installation documentation at:

https://docs.datastax.com/en/latest-opscenter/opsc/online_help/opscOverview_c.html

Setting Up the Cassandra Cluster

To set up the Cassandra cluster, see the following DataStax documentation:

- For a single data center, see:

http://docs.datastax.com/en/dse/5.1/dse-admin/datastax_enterprise/production/singleDCperWorkloadType.html

- For multiple data centers, see:

http://docs.datastax.com/en/dse/5.1/dse-admin/datastax_enterprise/production/multiDCperWorkloadType.html

When setting up multiple data centers, the Messaging Server recommendation, which minimizes the overhead in replicating and repairing DataStax keyspaces across all data centers, is to configure four data centers in three clusters with keyspaces arranged as shown in [Table 5–1](#).

Table 5–1 Recommended Multiple Data Centers and Clusters Configuration

Data Center Name and Node Types	Keyspaces	Cluster Configuration
DC_MSG, Cassandra nodes	ms_msg	Cluster Content
DC_META, Cassandra nodes	ms_mbox, ms_index	Combined with DC_INDEX into Cluster Metadata
DC_INDEX, Cassandra/Solr nodes	ms_index	Combined with DC_META into Cluster Metadata
DC_CACHE, Cassandra nodes	ms_cache	Cluster Cache

Cluster settings, such as the cluster name and seed nodes, are defined in the **cassandra.yaml** file. See the following section for more information.

To support more concurrent index updates, the ratio of **DC_META** nodes to **DC_INDEX** nodes should be at least 1 to 2.

Changing Initial Cassandra Settings

On each Cassandra node, optimize the DataStax Enterprise installation by following the recommendations in the DataStax documentation at:

https://docs.datastax.com/en/dse/5.1/dse-admin/datastax_enterprise/config/configRecommendedSettings.html

Changing Initial Tuning Settings

On each Cassandra node, change the configuration files described in this section so that the node operates correctly in the Cassandra message store deployment.

Linux Tuning Settings

To optimize Cassandra on Linux, see the DataStax recommendations at:

https://docs.datastax.com/en/landing_page/doc/landing_page/recommendedSettingsLinux.html

dse File

For **DC_INDEX** nodes, which run Solr, make the following change to the **/etc/default/dse** file:

```
SOLR_ENABLED=1
```

dse.yaml File

For **DC_INDEX** (Solr) nodes, make the following changes to the **/etc/dse/cassandra/dse.yaml** file to improve performance:

```
max_solr_concurrency_per_core: 6
back_pressure_threshold_per_core: 5000
cql_slow_log_option:
    enabled: false
```

cassandra.yaml File

Make the changes in this section to the **/etc/dse/cassandra/cassandra.yaml** file.

For all nodes, to enable separate clusters for better performance, specify **cluster_name**.

Make the following changes to the **num_tokens** setting:

- **DC_MSG, DC_META, and DC_CACHE** nodes:

```
num_tokens: 256
```

- **DC_INDEX** (Solr) nodes:

```
num_tokens: 16
allocate_tokens_for_local_replication_factor
```

On **DC_INDEX** (Solr) nodes, make the following change to the **allocate_tokens_for_local_replication_factor** setting:

```
allocate_tokens_for_local_replication_factor replication_factor
```

where *replication_factor* is derived from the **store.cassandra_rf** configuration option, and by default has a value of **2**.

Note: This recommendation is for the DataStax Enterprise 5.10. release.

To improve performance, locate data on SSD drives:

- **data_file_directories:**

```
/var/lib/cassandra/data
```

- **commitlog_directory:**

```
/var/lib/cassandra/commitlog
```

- **saved_caches_directory:**

```
/var/lib/cassandra/saved_caches
```

- **hints_directory:**

```
/var/lib/cassandra/hints
```

To support large mailbox and message, increase the **commitlog** size:

```
commitlog_segment_size_in_mb: 256
```

To specify seed nodes, you must use two nodes from each data center in the cluster, preferably located on different racks, so that each cluster has different seeds, for example:

- **DC_MSG** cluster:

```
seeds: "192.0.2.12,192.0.2.24"
```

- **DC_META/DC_INDEX** cluster:

```
seeds: "192.0.2.1,192.0.2.2,192.0.2.10,192.0.2.3"
```

- **DC_CACHE** cluster:

```
seeds: "192.0.2.14,192.0.2.7"
```

For **DC_INDEX** nodes, make the following changes to improve performance:

```
memtable_heap_space_in_mb: 2048
```

For all nodes, make the following change to improve performance:

```
memtable_flush_writers: 8
```

For all nodes, to specify **listen_address**, **rpc_address**, and so on, make the following changes:

```
listen_address: 10.128.128.12
```

```
rpc_address: 10.128.128.12
```

cassandra-env.sh File

For all nodes, to specify the location of the heap dump, make the following change to the **/etc/dse/cassandra/cassandra-env.sh** file:

```
export CASSANDRA_HEAPDUMP_DIR=/scratch/heapdump
```

jvm.options File

For **DC_MSGDC_META**, **DC_INDEX**, and **DC_CACHE** nodes, make the following heap size changes:

```
-Xms32G
```

```
-Xmx32G
```

For **DC_META** nodes, to improve performance, make the following heap size changes to improve performance:

```
-Xms16G
```

```
-Xmx16G
```

For all nodes, make the following changes to improve performance:

```
-XX:InitiatingHeapOccupancyPercent=70
```

```
-XX:ParallelGCThreads=12
```

```
-XX:ConcGCThreads=12
```

For all nodes, print garbage collection measurements, which are useful for monitoring system performance:

```
-XX:+PrintGCDetails
```

```
-XX:+PrintGCDateStamps
```

```
-XX:+PrintHeapAtGC
-XX:+PrintTenuringDistribution
-XX:+PrintGCApplicationStoppedTime
-XX:+PrintPromotionFailure
-XX:PrintFLSStatistics=1
-Xloggc:/var/log/cassandra/gc.log
-XX:+UseGCLogFileRotation
-XX:NumberOfGCLogFiles=10
-XX:GCLogFileSize=10M
```

cassandra-rackdc.properties File

For all nodes, make the following changes to the **cassandra-rackdc.properties** file.

1. Configure the endpoint snitch:

```
endpoint_snitch: GossipingPropertyFileSnitch
```

2. Set the data center and rack names as appropriate:

```
dc=mydc
rack=myrac
```

For example, for a node in **DC_CACHE** in a physical rack in one location, set **dc=DC_CACHE** and **rack=RAC1**. And, for another node in **DC_CACHE** in a physical rack in another location, set **dc=DC_CACHE** and **rack=RAC2**.

Note: Data center and rack names are case sensitive.

Changing Initial Solr Tunings

On each Cassandra/Solr node, you might need to make changes to the **solrconfig.xml** file, which is the configuration file with the most tuning parameters affecting Solr itself. For more information about Solr tuning parameters, see the DataStax documentation at:

http://docs.datastax.com/en/dse/5.1/dse-admin/datastax_enterprise/search/performanceTuningTOC.html

Note: DSE Search has a *live indexing* feature to increase indexing throughput, which is turned off by default. Enabling this feature causes sporadic search failures under load. This is a known DSE bug (DSP-12600) as of DSE 5.0.4.

For more information about the live indexing feature, see the DataStax documentation at:

http://docs.datastax.com/en/datastax_enterprise/4.8/datastax_enterprise/srch/tuningIndexing.html

To make changes to the **solrconfig.xml** file:

1. Use the **dsetool read_resource** *keyspace.table* **name=resfilename** command to read the xml.

where:

- *keyspace.table* is *ks-preindex* (*ks-pre* is the prefix configured by the **store.caskeyspaceprefix** option; the default is **ms_**)

- *resfilename* is **solrconfig.xml**
2. Edit the xml.
 3. Use the **dsetool write_resource** *keyspace.table* **name=resfilename** **file=path_to_file_to_upload** command to write changes to the xml.
where:
 - **file=path_to_file_to_upload** is the name and path of the resource file to upload
 4. Use the **dsetool reload_core** *keyspace.table* command to reload the Solr core.

Example:

```
dsetool read_resource ms_index.msgindex name=solrconfig.xml > /tmp/solrconfig.xml
vi /tmp/solrconfig.xml
##### Make required edits to the file #####
dsetool write_resource ms_index.msgindex name=solrconfig.xml
file=/tmp/solrconfig.xml
dsetool reload_core ms_index.msgindex
```

Installing Cassandra Message Store

This chapter describes how to install and configure Oracle Communications Messaging Server software on Cassandra message store nodes, message access tier hosts, and front-end tier hosts. For more information on the Messaging Server install command, see the discussion on the **commpkg install** command in *Messaging Server Installation and Configuration Guide*.

About Messaging Server Components

When you install the Cassandra message store, you install and configure one or more of the following components:

- **Cassandra and Solr Node.** Consists of the Cassandra database and Solr indexing and search system. The Cassandra keyspaces and schema are created whenever the Messaging Server **stored** process starts on any message access node that points to this Cassandra data center.
- **Message Access Tier.** Consists of a set of components that store, retrieve, and manipulate messages for mail clients. When using Cassandra message store, you configure at least IMAP and the LMTP server components on this tier. In addition, LMTP pre-conversion is set on by default (`iscimmediate` option is enabled). That is, when LMTP is involved, the message access tier's Indexed Search Converter (ISC) converts the messages before delivering them to the store. (The result is that email messages become searchable almost immediately upon delivery into the message store.) The message access tier also must be affinity-group aware for notification publishing.
- **Front-End Tier.** Consists of hosts that connect to the message access tier affinity group based on a user's LDAP **mailHost** attribute. For Cassandra message store, this includes the Message Transfer Agent (MTA), Messaging Multiplexor (MMP), LMTP Client, and **mshttpd** daemon.

Note: Messaging Server front-end tier components—Message Transfer Agent and Messaging Multiplexor (MMP) hosts, LMTP clients, and Webmail Server (**mshttpd**) hosts—must be running at least Messaging Server 8.0.1.

The Cassandra message store changes the semantics of the **mailHost** attribute in LDAP. Instead of referring to a physical host, in Messaging Server 8.0.2 "mail host" refers to a *store affinity group*. You configure the store affinity group by using the **proxy:mailhost.storehostlist msconfig** option. You must configure store affinity groups with the same **storehostlist** on all front-end and access-layer hosts in the deployment. Store affinity groups enable the MTA to route to an alternate LMTP server in the list if

the first one listed is unavailable. Affinity groups are used to distribute ENS delivery events so that IMAP IDLE works without the need to broadcast all delivery events to all IMAP servers.

Note: The store affinity group concept is also used by the Messaging Server 8.0.1 automatic failover (without cluster) feature. That is why it is supported for Messaging Server 8.0.1 front-end hosts to communicate with Cassandra message access-tier hosts.

About Unified Configuration

You must use Unified Configuration to configure and administer Messaging Server on the message access tier hosts.

Unified Configuration is an improved, streamlined process to configure and administer Messaging Server. Unlike in legacy configurations, Unified Configuration uses validation to verify configuration accuracy, and employs a single tool to configure the entire Messaging Server configuration (with a few exceptions). For more information, see the discussion on Unified Configuration in *Messaging Server System Administrator's Guide*.

Before Installing Messaging Server

You must have already completed the following steps on all Messaging Server components:

1. Ensure that DNS is running and configured properly.
2. Ensure you have sufficient file descriptors on Linux.
3. Make sure you do not configure conflicting port numbers on a host when various components run on a single machine.
4. Preparing Directory Server by running the **comm_dssetup.pl** script against it.

For more information on completing these tasks, see "Messaging Server Pre-Installation Tasks" in *Messaging Server Installation and Configuration Guide*.

The following steps must be completed for Cassandra nodes:

1. Install Java.
You also need to install Java on the message access tier hosts if they run the Indexed Search Convertor (ISC), in addition to the normal processes (**imap**, **enpd**, and **lmtps**).
2. Install Python.
3. Install DataStax Enterprise software.
4. Start DataStax Enterprise as a transactional (Cassandra) node.

See [Chapter 5, "Cassandra Message Store Pre-Installation Tasks"](#) for more information.

Using a Configuration Management Tool to Simplify Installation and Configuration

Currently, Messaging Server does not supply its own tool for configuration management. However, you can use a configuration management tool, such as Chef or Puppet, to simplify the installation and management of Messaging Server across

multiple servers. For more information, see the topic on using a configuration management tool with Messaging Server in *Messaging Server Installation and Configuration Guide*.

Installing and Configuring Messaging Server

The tasks to install the front-end and message access tiers, and the Cassandra message store are as follows:

- [Downloading the Messaging Server Software](#)
- [Installing and Configuring the Messaging Server Software](#)

Downloading the Messaging Server Software

1. Download the media pack for Oracle Communications Messaging Server software and for Messaging-CAS software (the Comms FIT component) from the Oracle software delivery website, located at:
<http://edelivery.oracle.com/>
2. Copy the Messaging Server ZIP file to a temporary directory on your Messaging Server hosts and extract the files.
3. Copy the Comms FIT ZIP file to a temporary directory on your Cassandra/Solr hosts and extract the files.

Installing and Configuring the Messaging Server Software

You must install and configure the Cassandra message store, and associated components, in the following order, to ensure proper operation:

- [Installing Messaging-CAS Software on DSE/Solr Nodes](#)
- [Installing Message Access Tier Hosts](#)
- [Installing MTA Tier Hosts](#)
- [Installing MMP Hosts](#)

Note: You must use ENS for notifications with the Cassandra message store. The default when you run the **configure** command is for ENS to be enabled, so no additional steps are necessary to enable ENS. However, manual store affinity group configuration is required to have ENS tolerate host failures.

Caution: When completing the steps in this section, do not start Messaging Server until instructed to do so. Otherwise, errors might result.

Installing Messaging-CAS Software on DSE/Solr Nodes

This task assumes that you have already installed and started Cassandra, as explained in ["Installing the DataStax Enterprise Software"](#). The Messaging-CAS software contains the FIT component, which you must install on every Cassandra/Solr node.

To install and configure Messaging-CAS software on the DSE/Solr nodes:

1. Log in as or become the superuser (**root**).
2. Use the **commpkg install** command to install Messaging-CAS software.
See "commpkg Reference" in *Messaging Server Installation and Configuration Guide* for more information about the **commpkg install** command.
3. Run the **config-fit** script to create the FIT configuration.

See "[config-fit Reference](#)" for more information about the **config-fit** script.

```
cd MessagingServer_home/cas/bin
config-fit
```

Be sure to use the same FIT port throughout the deployment. (The default is 8070.)

Note: The UNIX owner of the FIT directories and files must be the same as the owner of the *DSE_HOME* location. Permissions for the configuration files must be similar to those used by the **configure** script (644 for the **config.xml** file, and 600 for the **xpass.xml** file).

Installing Message Access Tier Hosts

The message access tier hosts must run at least Messaging Server 8.0.2 and be configured for Unified Configuration.

To install and configure Messaging Server software on the message access tier hosts:

1. Log in as or become the superuser (**root**).
2. Use the **commpkg install** command to install Messaging Server software.
See "commpkg Reference" in *Messaging Server Installation and Configuration Guide* for more information about the **commpkg install** command.
3. Create the Cassandra configuration:

```
cd MessagingServer_home/bin
configure
```

Note: By default, the **configure** command also performs the necessary configuration to set up LMTP pre-conversion of messages. Round robin load balancing does not exist for hosts specified as ISC servers to perform LMTP pre-conversion. Instead, the first available server in list is used.

4. Configure required Cassandra message store options by using the **CasAccessTier.rcp** recipe file.

This recipe enables the message store, sets the message store type to Cassandra, sets replication factors and data centers, sets connection points, sets up the LMTP server, sets up store affinity groups, and disables unwanted Messaging Server processes.

- a. Change to the *MessagingServer_home/lib/recipes* directory.
- b. Edit the **CasAccessTier.rcp** recipe for your deployment to configure the access tier hosts. Relevant options that the **CasAccessTier.rcp** file sets are:

```
store.enable = 1
store.purge.enable = 1
```

```

store.serviceadmingroupdn = Service_Administrators_Group_DN
store.dbtype = cassandra
store.casconnectpoints = IP_address_CSnode1 IP_address_CSnode2 ...
store.casmetarf = 3
store.casmsgsf = 3
store.cascacherf = 1
store.cassolrdc = DC2
store.cassolrrf = 3
store.cascachedc = DC3
store.casmetadc = DC1
store.solrconnectpoints = IP_addressSL_node IP_address_SL_node ...

```

Note: The **store.enable**, **store.purge.enable**, and **store.serviceadmingroupdn** options should be set automatically by the **configure** command.

- c. Run your edited recipe file:

```
msconfig run CasAccessTier.rcp
```

5. If Cassandra cluster authentication is required, use the **msconfig** command to set the **store.casusername** and **store.caspassword** options.
6. (Optional) Configure how the ISC is invoked.

By default, the ISC is configured with the **iscimmediate** option, for LMTP pre-conversion. See *Messaging Server Reference* for more information on the **iscdeferred** option to delay ISC operations until after delivery to the message store.

Installing MTA Tier Hosts

The MTA tier hosts must run Messaging Server 8.0.1 or greater and be configured for Unified Configuration (for store affinity group support).

To install and configure Messaging Server software on the MTA tier (Submission/Relay) hosts:

1. Log in as or become the superuser (**root**).
2. Use the **commpkg install** command to install Messaging Server software.
See "commpkg Reference" in *Messaging Server Installation and Configuration Guide* for more information about the **commpkg install** command.
3. Create the configuration:

```

cd MessagingServer_home/bin
configure --noisc

```

Running the **configure** command automatically sets up the LMTP channel correctly on the MTA tier hosts. Also, running the **configure --noisc** command installs the front-end services without requiring that Java be present.

4. Configure the MTA by copying the **CasFrontEnd.rcp** recipe from the *MessagingServer_home/lib/recipes* directory to the *ConfigRoot/recipes* directory, then editing the file.

When configuring an MTA host, make sure to:

- Disable the message store

- Disable the LMTP server
- Disable **mshttpd**

For the store affinity group:

- Specify a list of fully qualified server names that have access to the same message store data. The first host name in the list should be the default master for the message store while the other hosts in the list should be configured to act as failover hosts for the master. When a **mailHost** attribute is found in LDAP, it is first resolved through this configuration. The store affinity group settings (the **storehostlist** option values) must be the same on all hosts in the deployment. It is not permitted for the same server name to appear in more than one **storehostlist**. For the Cassandra message store, you can reorder the **storehostlist** on different servers as long as the content is the same.

Note: The MTA ignores the **storehostlist** option unless the **affinitylist** option is set.

5. Run your edited recipe file:

```
msconfig run CasFrontEnd.rcp
```

Installing MMP Hosts

The MMP tier is required if there is more than one store affinity group. The classic message store and Cassandra message store must be in different store affinity groups if both are present in the deployment for migration purposes.

The MMP tier hosts must run Messaging Server 8.0.1 or greater and be configured for Unified Configuration (for store affinity group support).

To install and configure Messaging Server software on the Cassandra MMP tier hosts:

1. Log in as or become the superuser (**root**).
2. Use the **commpkg install** command to install Messaging Server software.
See "commpkg Reference" in *Messaging Server Installation and Configuration Guide* for more information about the **commpkg install** command.
3. Create the configuration:

```
cd MessagingServer_home/bin  
configure --noisc
```

Running the **configure --noisc** command installs the front-end services without requiring that Java be present.

4. Configure the MMP by copying the **CasFrontEnd.rcp** recipe from the *MessagingServer_home/lib/recipes* directory to the *ConfigRoot/recipes* directory, then editing the file.

When configuring an MMP host, make sure to:

- Disable the message store
- Disable the LMTP server
- Disable **mshttpd**
- Enable the MMP (**mmp.enable** option)
- Uncomment the **delete_** options

For the store affinity group:

- Specify a list of fully qualified server names that have access to the same message store data. The first host name in the list should be the default master for the message store while the other hosts in the list should be configured to act as failover hosts for the master. When a **mailHost** attribute is found in LDAP, it is first resolved through this configuration. The store affinity group settings (the **storehostlist** option values) must be the same on all hosts in the deployment. It is not permitted for the same server name to appear in more than one **storehostlist**. For the Cassandra message store, you can reorder the **storehostlist** on different servers as long as the content is the same.

5. Run your edited recipe file:

```
msconfig run CasFrontEnd.rcp
```

Using a Non-MMP IMAP Proxy

If you want to use a different IMAP proxy, that proxy must support IMAP LOGIN referrals (RFC 2221) and for robustness needs to support a facility to handle the store affinity groups.

Configuring a Remote Indexed Search Converter

To configure a remote Indexed Search Converter (ISC), where the ISC is not running on the same host as LMTP, run the following command:

```
msconfig set isc_client.ischosts ISC_host_name1 ISC_host_name2 ...
```

Enabling SSL

You can configure the ISC and FIT components to communicate using Secure Sockets Layer (SSL). To enable SSL, you set the appropriate ISC and FIT configuration options, and generate a certificate to import into the certificate database. For detailed steps, see the topic on configuring the ISC component for SSL in *Messaging Server Security Guide*.

Starting the Cassandra Message Store Across Tiers

After you have installed Messaging Server software and configured the Cassandra message store (across all tiers), start the tiers in the following order:

1. DSE non-Solr nodes: DSE should already be started (by running the **service dse start** command).
2. Solr nodes: DSE should already be started. When starting Messaging Server, only the **isc** process should start.
3. Message access hosts: When starting Messaging Server, the appropriate message access services (**imapd**, **stored**, **isc**, and so on) should start.
4. Front-end hosts: When starting Messaging Server, the appropriate front-end services (MMP, MTA, **mshttp**, and so on) should start.

Restarting a Cassandra Cluster

A restarted Solr process on a rebooted Solr node takes a long time to join the Cassandra cluster. Therefore, when restarting a cluster, you must follow specific steps to avoid read failures in the FIT logs and DSE taking a long time to start on the Solr

nodes. See the topic on restarting a Cassandra message store cluster in *Messaging Server System Administrator's*.

Cassandra Message Store Post-Installation Tasks

This chapter describes the post-installation tasks that you must complete to finish the Cassandra message store installation, and how to add and remove additional Cassandra nodes.

Verifying the Installation

To verify the Cassandra message store installation was successful:

1. Verify the Solr core was created by logging in to the Solr administration console at:

`http://host:8983/solr`

If the administration console cannot be reached, check the **system.log** file on the Cassandra host for errors. By default, this file is located at **/var/log/cassandra/system.log**.

2. Verify the Indexed Search Converter (ISC) process launched correctly by accessing the statistics servlet at:

`http://host:8070/converter/statistics`

If the ISC cannot be reached, check the ISC log file for errors. By default, this file is located at **/opt/sun/comms/messaging64/log/isc.log.***.

Scaling Your Cassandra Message Store Deployment Horizontally

To scale your Cassandra message store deployment horizontally by adding more nodes to an existing deployment, see *Messaging Server System Administrator's Guide*.

For adding a node, and adding a data center, also see the Cassandra documentation at:

<https://docs.datastax.com/en/cassandra/2.1/cassandra/operations/opsAddingRemovingNodeTOC.html>

Migrating to Cassandra Message Store

This chapter describes how to migrate from a Messaging Server 8.0.1 or prior "classic" message store to Messaging Server 8.0.2 and Cassandra message store. It assumes that you have chosen a target deployment, and have developed an architectural design and deployment plan.

Overview of Migrating to Cassandra Message Store

The high-level steps to migrate to Cassandra Message store include:

1. Upgrading current front-end Message Transfer Agent (MTA) hosts, Messaging Multiplexor (MMP) hosts, and Webmail Server (**mshttpd**) hosts to at least Messaging Server 8.0.1.

You must upgrade these hosts to at least Messaging Server 8.0.1 to support store affinity groups.

The Cassandra message store changes the semantics of the **mailHost** attribute in LDAP. Instead of referring to a physical host, in Messaging Server 8.0.2 "mail host" refers to a *store affinity group*. You configure the store affinity group with the **proxy:mailhost.storehostlist msconfig** option. You must configure store affinity groups on front-end MTAs for robust Cassandra message store delivery. Store affinity groups enable the MTA to route to an alternate LMTP server in the list if the first one listed is unavailable. Affinity groups are used to distribute ENS delivery events so that IMAP IDLE works without the need to broadcast all delivery events to all IMAP servers.

2. If not already using LMTP, enabling LMTP server on the classic message store host. When you enable LMTP server, you must restart the dispatcher on each classic store.

For more information, see ["Installing Message Access Tier Hosts"](#).

3. If not already using LMTP for delivery, switch front-end MTAs to use LMTP for delivery. When you switch front-end MTAs to use LMTP, you must run the **imsimta restart smtp** command on those MTAs whose delivery options you changed.

For more information, see ["Installing MTA Tier Hosts"](#).

4. Installing and configuring the Cassandra message store.

For more information, see ["Cassandra Message Store Pre-Installation Tasks"](#) and ["Installing and Configuring the Messaging Server Software"](#).

5. Using the **rehostuser** command to migrate users from classic store to Cassandra store.

For more information, see the **rehostuser** topic in *Messaging Server System Administrator's Guide*.

For example, the following **rehostuser** command migrates user **joe** to affinity group **ngstore** using message access server **ma1**.

```
rehostuser -u joe -d ngstore -R ma1
```

The **-d target** for the **rehostuser** command becomes the user's **mailHost** attribute. For the end users to have access to their email after migration, that *target* must either match one of the host names in the **storehostlist** option for that affinity group, or must match the name of the affinity group. Use of the **-R** option is only required if the **-d target** is not a resolvable host name.

Note: The **storehostlist** option for an affinity group must be set to the same value on all front-end MTA, MMP, and **mshttpd** hosts, and all message access-tier hosts. Use fully qualified domain names when referring to hosts, both with the **rehostuser** command and in the **storehostlist** option.

Uninstalling the Cassandra Message Store

This chapter describes how to uninstall Cassandra message store.

Uninstalling Cassandra Database

To uninstall the Cassandra database, see "Uninstalling DataStax Enterprise" at:

http://docs.datastax.com/en/dse/5.1/dse-admin/datastax_enterprise/install/dseUninstall.html

Uninstalling Message Server Software

To uninstall the Messaging Server software, see the "Uninstalling Messaging Server" chapter in *Messaging Server Installation and Configuration Guide*.

config-fit Reference

This appendix provides information about the **config-fit** script.

Overview of the config-fit Script

The **config-fit** script configures the Field Input Transformer (FIT) on a Cassandra/Solr node. The **config-fit** script is provided as part of the Comms FIT software, which you download and install separately from the Messaging Server software. By default, the **config-fit** script is located in the `/opt/sun/comms/cas/bin` directory.

Syntax

config-fit [options]

Table A-1 describes the **config-fit** options.

Table A-1 *config-fit Options*

Option	Description
-r DSE_HOME, --dsehome DSE_HOME	Specifies the DSE HOME directory. The default is <code>/usr/share/dse</code> .
-h ISC_Hosts, --ischost ISC_Hosts	Specifies a space-separated list of hosts for the ISC server. For example: -h <code>isc1.example.com isc2.example.com</code>
-c Cache_Nodes, --cacheconnectpoints Cache_Nodes	Specifies a space-separated list of IP addresses/hosts for the cache nodes. Specify at least two for high availability.
-u username, --username username	Specifies the user name to be used by the FIT to communicate with the ISC server.
-w PASSWDFILE, --authpasswordfile PASSWDFILE	Specifies the file that contains the password for the auth user.
-d LOGDIR, --logdir LOGDIR	Specifies the location of the FIT log files.
-e LOGLEVEL, --statsjloglevel LOGLEVEL	Specifies the FIT statistics log level. The default is INFO .
-j LOGLEVEL, --jloglevel LOGLEVEL	Specifies the FIT log level. The default is INFO .
-i INTERVAL, --statsinterval INTERVAL	Specifies the frequency (in minutes) of statistics printed to the logs. The default is 10 .
-q TIMEOUT, --querytimeout TIMEOUT	Specifies the FIT Cassandra query timeout in milliseconds. The default is 2000 .

Table A–1 (Cont.) config-fit Options

Option	Description
-S <i>FLAG</i> , --sslusessl <i>FLAG</i>	Specifies the flag to enable SSL for ISC server connections. Use 1 to enable SSL, 0 to disable SSL. The default is 0 .
-v <i>PORT</i> , --serverport <i>PORT</i>	Specifies the HTTP port for the ISC process. The default is 8070 . Use the same port throughout the deployment.
--state <i>filename</i>	Uses the silent state <i>filename</i> for configuration.
--saveState <i>filename</i>	Specifies to save a state file. The location of the state file is \$DSE_HOME/fit/setup , which by default is the /usr/share/dse directory.
--help , -?	Displays the help.

Table A–2 describes the configuration options for the **config-fit** script. You can also use these options in a state file for silent installation.

Table A–2 config-fit Configuration Options

Option	Description
fit.ischosts	Specifies a list of host names that run the ISC server. Round robin load balancing is used for hosts specified as ISC servers.
fit.authpassword	Specifies the password of the user name required by FIT to communicate with the ISC server.
fit.authusername	Specifies the user name required by FIT to communicate with the ISC server. The default is admin .
store.caskeyspaceprefix	Specifies the prefix of the Cassandra keyspaces. The default is ms_ .
store.cacheconnectpoints	Specifies a space-separated list of Cassandra Cluster connect points (IP addresses) for cache keyspace access.
fit.logdir	Specifies the FIT log directory. The default is the Cassandra log path.
fit.jloglevel	Specifies the log level for the FIT component. The default is INFO .
fit.stats.jloglevel	Specifies the log level for the FIT statistics log. The default is INFO .
fit.stats.interval	Specifies the collection interval for the FIT statistics log. The default is 10.
fit.querytimeout	Specifies the query timeout level for the FIT component. The default is 2000 .
fit.sslusessl	Specifies to use SSL. The default is 0 .
fit.server_port	Specifies the port used by the FIT component. The default is 8070 . Use the same FIT port throughout the deployment.