Esgyn Corporation

EsgynDB Installation Guide for Release R2.5.0

Published: December 2018

Edition: EsgynDB Release 2.5.0

Contents

Legal Notice

The information contained herein is subject to change without notice. This documentation is distributed on an “AS IS” basis, without warranties or conditions of any kind, either express or implied. Nothing herein should be construed as constituting an additional warranty. Esgyn Corporation shall not be liable for technical or editorial errors or omissions contained herein.

NOTICE REGARDING OPEN SOURCE SOFTWARE: Project Trafodion is licensed under the Apache License, Version 2.0 (the "License"); you may not use software from Project Trafodion except in compliance with the License. You may obtain a copy of the License at http://www.apache.org/licenses/LICENSE-2.0. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Acknowledgements

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation. Java® and MySQL® are registered trademarks of Oracle and/or its affiliates. Bosun is a trademark of Stack Exchange Inc. Apache®, Hadoop®, HBase®, Hive®, openTSDB®, Sqoop®, and Trafodion® are trademarks of the Apache Software Foundation. Esgyn and EsgynDB are trademarks of Esgyn Corporation.

1. About This Document

This EsgynDB Installation Guide describes how to install and configure Release R2.5.0 of the EsgynDB core product (based on Trafodion) and required components on a Hadoop cluster.

It does not describe how to install the OS or Hadoop distribution, which are prerequisites for the EsgynDB installation. Refer to the vendor documentation for help in installation and configuration of those components.

2. Intended Audience

This guide is intended for EsgynDB and Hadoop system administrators.

3. Pre-requisite

Starting with EsgynDB R2.2, installation requires a license key provided by Esgyn Corporation. Make sure you have one before starting the installation process.

4. Preparing Your PC

If you are using a Windows PC, additional PC software may be needed for the installation process. It is recommended that you preinstall the PC software before continuing with the EsgynDB installation.

Helpful PC software:

You may use functional equivalents of the following

- putty and puttygen (download from the PuTTY website)
- VNC client (download from http://www.realvnc.com)
- Firefox or Chrome browser
- SFTP client to transfer files from your PC to the server: WinSCP or FileZilla

5. Validate Your Cluster Environment

Before installing EsgynDB, validate your cluster environment.

5.1 Cluster Requirements

The following configuration settings have been tested and are known to work with the EsgynDB installation:

Hardware platform	**x86-64**
Operating systems	CentOS 6.5 ~ 6.9, 7.2 ~ 7.4 (64 bit) Red Hat 6.5 ~ 6.9, 7.2 ~ 7.4 (64-bit)
Hadoop distributions	Cloudera CDH 5.9 ~ 5.13 Hortonworks HDP 2.4 ~ 2.6.3
User IDs	A user ID with passwordless sudo access for command-line installer. See Obtaining Sudo Access and Passwordless SSH.
Cluster size	A cluster consisting of 1 to n nodes. There is currently no upper limit. Two nodes is the minimum, but recommendation is to use at least 4 nodes.
Disk space	Minimum of 20 GB. See Checking the Disk Space.
Memory	Minimal 1 GB * number of connectivity servers (MXOSRVR processes) configured per node in the cluster.

5.2 Check the Disk Space

Before installing the Hadoop distribution, ensure a minimum of 20 GB is available to support the database. The default installation location for Cloudera CDH is **"/var/lib/cloudera-scm-server-db"**. To check the size available to **/var**, start a putty session, or a VNC terminal window on the node in your cluster where Cloudera will be installed. To execute the following commands, you will need either **root** or **sudo** access.

Confirm there is a minimum of 20 GB available to **/var**.

```
$ cd "/var"
$ df -hP
```

If there is insufficient space available in **/var**, a possible solution is to provide a soft link to another drive for your Cloudera database. Locate a drive that does have sufficient space.

```
$ cd <new drive>  (e.g. cd /DATA)
$ mkdir cloudera-scm-server-db
$ chmod 777 cloudera-scm-server-db
$ cd /var/lib
$ ln -s <new drive>/cloudera-scm-server-db .
```

If **/var** is a subdirectory in your cluster's root filesystem, the Cloudera database should have sufficient space available.

If the Cloudera distribution has already been installed and it is showing red for log directories, it means that Cloudera was installed using **/var/lib** and may not have a large enough **/var** file system. In this case, there is an unsupported script that can be used to move the directories. Please see the **clouderaMoveDB.sh** script in the **installer/tools** directory, which is created when the installer tar.gz file is untarred. Execute **clouderaMoveDB.sh** without syntax to display help information.

6. Install a Supported Hadoop Distribution

EsgynDB R2.5.x is compatible with the Cloudera and Hortonworks distributions.

Distribution	Version	HBase Version	Installation
Cloudera Distribution Including Apache Hadoop (CDH)	CDH 5.9 ~ 5.13	1.2	Refer to installation instructions on the Cloudera site for the specific version you plan to install.
Hortonworks Data Platform (HDP)	HDP 2.4 ~ 2.6.3	1.1	Refer to installation instructions on the Hortonworks site for the specific version you plan to install.

6.1 Mandatory Hadoop Services and Settings

IMPORTANT: Before installing a Hadoop distribution, please review this list of mandatory services and settings:

- HDFS
- Yarn/MapReduce
- ZooKeeper
- HBase
- Hive
- Embedded Databases

Please make sure that those services and settings are selected during installation.

IMPORTANT: The EsgynDB command-line installer needs to run from one of the nodes that will be a part of the EsgynDB cluster. Off-platform installation is not currently supported. All EsgynDB nodes must have HBase installed.

7. Prepare for Install

7.1 Obtain sudo access and passwordless SSH (for command-line install option)

The EsgynDB installation requires a user ID with these attributes:

- sudo access
- passwordless ssh to all nodes on the cluster

> Note: You may need to request permission from your cluster management team to obtain this type of access.

The following example shows how to set up your user ID to have "passwordless ssh" abilities.

```
$ echo -e 'y\n' | ssh-keygen -t rsa -N "" -f $HOME/.ssh/id_rsa
$ cat $HOME/.ssh/id_rsa.pub >> $HOME/.ssh/authorized_keys
$ echo localhost $(cat /etc/ssh/ssh_host_rsa_key.pub) >> $HOME/.ssh/known_hosts
$ echo "NoHostAuthenticationForLocalhost=yes" >> $HOME/.ssh/config
$ chmod 600 $HOME/.ssh/config
```

After running these commands, you will need to copy the contents of the public key file, `$HOME/.ssh/id_rsa.pub`, and append those contents to each additional node's `$HOME/.ssh/authorized_keys` file. You will also need to copy your private `$HOME/.ssh/id_rsa` file from the current node to the other node's `$HOME/.ssh` directory, and secure it as private to yourself (chmod 700).

7.2 Configure an LDAP Identity Store

If you plan to enable authentication in EsgynDB, you will need to have a configured LDAP identity store available. The EsgynDB installer will prompt you to set up an authentication configuration file that points to an LDAP server (or servers), which will enable security (that is, authentication and authorization). Refer to the section on Enabling Security Features for more details.

7.3 Integrate with Kerberos

If Kerberos is enabled on your system, then it needs to be enabled in EsgynDB. The EsgynDB installer will prompt you to set up Kerberos attributes needed to install Kerberos principals and keytabs. Refer to the section on Enabling Security Features for more details.

7.4 User IDs and Passwords

This table lists the user IDs and passwords that you will use during installation.

NOTE: You will use two user IDs: a user ID with sudo access and **trafodion**. The **trafodion** id is auto created by the installer if it does not exist.

Logon	User ID	Password	Description
Cloudera Manager Web GUI logon	`admin`(default)	`admin`(default)	After installing Cloudera, you will be instructed to log on to the Cloudera Manager Web GUI. Use the default user ID and password. If you already had Cloudera installed, please use your previously defined user ID and password.
Apache Ambari Web GUI logon	`admin`(default)	`admin`(default)	After installing Hortonworks' HDP, you will be instructed to log on to the Apache Ambari Web GUI. Use the default user ID and password. If you already had HDP installed, please use your previously defined user ID and password.
User ID with sudo access	<sudo-username>	<password>	In the installation steps, you may be instructed to use "sudo" or "sudo userid" access. You will be using your user ID, which has been enabled with "sudo" access and passwordless ssh to all nodes of the cluster.

EsgynDB logon	`trafodion`	`traf123` (default)	This user ID is **automatically created** for you by the EsgynDB installer when EsgynDB is installed. Do not create this user ID manually.

7.5 Required Software

EsgynDB requires supplementary software to be installed on the cluster before it is installed.

A basic pre-requisite is JDK 1.8.

In addition, a set of Linux tools is required that are not typically packaged as part of the core Linux distribution. The installer will attempt to automatically get these packages over the Internet, but if the cluster's access to the Internet is disabled, you will need to manually download the packages and make them available for installation.

OS	Packages	
CentOS Linux 6.5 ~ 6.9, 7.2 ~ 7.4	pdsh	apr
	log4cxx	apr-util
Red Hat Linux 6.5 ~ 6.9, 7.2 ~ 7.4	sqlite	protobuf
	expect	lzo
	perl-DBD-SQLite	lzop
	xerces-c	unzip
	perl-Params-Validate	gcc-c++
	perl-Time-HiRes	unixODBC
	gzip	unixODBC-devel
	gnuplot	libiodbc
	lsof	libiodbc-devel
	keepalived	openldap-clients
	libcgroup	snappy

7.6 Information Gathering

The installer will prompt for additional information over the course of the installation process. Before you start installation, make sure you have collected the data

Information	Default	Notes
Location of the working directory where the installer will untar files	None	Specify a location
Location of the directory on the node where the installer package was untarred	None	You will need to specify the location
License key	None	You will need a license key from Esgyn Corporation
Number of concurrent client sessions per node	8	This number specifies the concurrent sessions per node to be supported. Each session could require up to 1GB of physical memory. The number can be changed post-installation. For more information, refer to the DCS Installation Guide
Upgrade of software on an existing EsgynDB installation or a new install	None	The installer will take additional actions if the installation is on a new cluster
Trafodion user id and password	User ID: `trafodion` Password: `*****`	Recommend not to change the User ID. The python command line installer will prompt for a password. The other installers

		(integrated into Cloudera Manager and Ambari) will create a service account with no password.
List of nodes in the cluster	None	All nodes should be accessible by passwordless ssh as the sudo or root user.
Prefix of the home directory path of the trafodion user ID	/home	If the home directory of the trafodion user ID is "/opt/home/trafodion", specify the prefix as "/opt/home"
Location of the JDK	None	Fully qualified path of the JDK. For example: "/usr/java/jdk1.8.0_112-cloudera"
Location of the EPEL rpm	None	If your machine does not have external Internet access then you must install the EPEL repo manually.
Location of the EsgynDB package tar file	None	Specify the full pathname of the EsgynDB package
URL for the Hadoop distribution	None	Specify in the form: <IP-address>:<port> or <node name>:<port> Example: "vm-1.yourcompany.local:7180"
Hadoop distribution details 1. Admin UI user ID, password	Distribution-dependent Cluster 1	

2. Cluster name 3. HDFS user ID 4. HBase user ID, group 5. HBase service name	hdfs hbase, hbase hbase	
EsgynDB installation directory	None	Specify the full path for the directory. This will allow you to maintain multiple versions of the software if desired.
DCS HA (High Availability)	Disabled	You will need the Floating IP address and the interface~~, and the backup nodes for DCS master.~~ Recommend you specify a list of DCS Master nodes (more than one). If DCS HA is enabled, the Floating IP will be used. If not, the clients can use the multiple IP feature.
Security	Disabled	If security is to be enabled, ensure LDAP is configured and the name of the LDAP configuration file is available.
Kerberos	Hadoop option	If Kerberos is enabled on your system, EsgynDB needs to create Kerberos principals and keytabs for the `trafodion` ID.

8. Install EsgynDB

8.1 Command Line Installer

The EsgynDB command line installation tool is called `db_install.py` It is distributed as an independent package.

EsgynDB must be installed on all nodes that host an HBase RegionServer (that is, where a supported Hadoop distribution is installed).

1. You will execute the python installer from one of the nodes in the EsgynDB cluster, as user <sudo-username>. Create the `esgyndb_downloads` directory (if one does not exist) on a node that is part of the cluster. Typically, this will be the first node in the cluster. Place the downloaded server package (**esgynDB_server-2.5.0-RH6-x86_64.tar.gz** or **esgynDB_server-2.5.0-RH7-x86_64.tar.gz,** depending on the installed OS version) and the corresponding installer package (**esgynDB_pyinstaller-2.5.0-RH6.tar.gz** or **esgynDB_pyinstaller-2.5.0-RH7.tar.gz**) compressed tar files into it.

   Example

   ```
   $ mkdir $HOME/esgyndb_downloads
   $ mv <your-download-path>/esgynDB_server-2.5.0-RH6-x86_64.tar.gz
   $HOME/esgyndb_downloads
   $ mv <your-download-path>/esgynDB_pyinstaller-2.5.0-RH6.tar.gz
   $HOME/esgyndb_downloads
   $ cd $HOME/esgyndb_downloads
   ```

2. Untar the downloaded installer file. For example

   ```
   $ tar -xzf esgynDB_pyinstaller-2.5.0-RH6.tar.gz
   ```

3. Change to the installer directory:

   ```
   $ cd python-installer
   ```

4. Run the installer script. You may execute it in one of two installation modes -

Mode	Comments
Guided Setup	The installer will interactively prompt for information from the user as it works through the installation process. Recommended for the new user. `$ ./db_install.py`
Expert Setup	Required information is provided in a pre-formatted plain text configuration file. The installer is pointed to this file at invocation. The installer will not prompt for any further information. A template of the configuration file is available within the untarred installer directory: `python-installer/configs/db_config_default.ini` Make a copy of the file in your directory and populate the needed information. Recommended for the expert user and/or for unattended installs. `$ ./db_install.py –config_file <`*`installer-config-file`*`>`

NOTES:

- The installer will abort in case of any configuration errors. Correct the errors before you re-run.

- If you choose not to start Trafodion after the installation (that is, if you enter **N** for **Start Trafodion after install (Y/N)**), you will need to manually start and initialize Trafodion after **db_install.py** completes. See notes below.
- For details on manually enabling security in EsgynDB, refer to the section Enabling Security Features.

The installation process will

- Prompt for a license key

```
Add a new license file or license string (Y/N) [N]: Y
Enter full path to license file or the license string [NONE]:
```

- Install necessary RPMs on CentOS/Redhat Linux systems
- create the `trafodion` user ID
- set up passwordless ssh for the `trafodion` user ID
- copy the EsgynDB distribution files across the cluster, and
- generate startup files
- start EsgynDB, Database Connectivity Services (DCS) to access the system, and EsgynDB Manager services.

5. Once installation completes successfully, you will see this message:

```
**************************
  Installation Complete
**************************
```

6. Your EsgynDB system should now be up and running. Log on to the system as the `trafodion` user.

   If you chose not to start Trafodion after the installation, start and initialize Trafodion as follows

```
cds
sqstart
```

```
[trafodion@nap001 ~]$ sqlci
EsgynDB Enterprise Conversational Interface 2.5.0
Copyright (c) 2015-2018 Esgyn Corporation
>>initialize trafodion;
```

8.2 Manage

You use the `trafodion` user ID to perform EsgynDB management operations.

The following table provides an overview of the subsystem management scripts.

Component	Start	Stop	Status
EsgynDB (all components)	`sqstart`	`sqstop`	`sqcheck`
RMS Server	`rmsstart`	`rmsstop`	`rmscheck`
REST Server	`reststart`	`reststop`	`restcheck`
LOB Server	`lobstart`	`lobstop`	-
DCS (Database Connectivity Services)	`dcsstart`	`dcsstop`	`dcscheck`

Example: Start EsgynDB

```
$ cd $TRAF_HOME/sql/scripts
$ sqstart
$ sqcheck
```

8.3 Cloudera Manager Installer

This installer leverages Cloudera Manager for the installation, and is applicable only in cases where the Cloudera Hadoop distribution (CDH) is installed, as a Cloudera parcel.

The EsgynDB software is comprised of a Custom Service Descriptor (CSD), two Parcels, and a parcel manifest. A configuration helper script (`cm_settings.sh`) is also provided. Use of the helper script is optional, as all of the needed configurations are described in this document (see Hadoop Configuration).

EsgynDB software is provided in a single compressed tar-file download, such as **`esgyndb-CM-`*`version`*`.tgz`**.

8.3.1 Installing CSD (Cloudera Service Descriptor)

The CSD is a jar file, named **`ESGYNDB-`*`version`*`.jar`**, such as `ESGYNDB-2.5.0-1.jar`. This file must be placed on the Cloudera Manager Server host in the Local Descriptor Repository directory, which defaults to `/opt/cloudera/csd`.

If Cloudera Manager is not yet installed, the directory may be pre-created and the EsgynDB CSD placed there. When Cloudera Manager is installed, it will be automatically picked up.

If Cloudera Manager is already installed, then the server process (cloudera-scm-server) must be re-started to pick up the CSD (e.g., `sudo service cloudera-scm-server restart`).

8.3.2 Installing Parcels

The ESGYNDB_TRX parcel is a HBase plug-in for transaction processing. The ESGYNDB parcel is the main service providing the SQL engine and other database services.

The parcel files and parcel manifest file must be placed in a remote parcel repository. This may be any regular web-server location or a temporary web-server location.

To configure the parcel location, click the parcel icon in the Cloudera Manager menu bar, and then the Configuration button in the upper right. Add a new Remote Parcel Repository URL, referencing your EsgynDB parcels location.
On the main Parcels screen, click the Check for New Parcels button.

Now that Cloudera Manager knows about the EsgynDB service and where to find EsgynDB parcels, the EsgynDB service is available to be installed.

8.3.3 Adding EsgynDB to an Existing Cluster

For an existing cluster, follow these steps:

1. Go to the Parcels screen (package icon on title menu), select the cluster.
2. For both the ESGYNDB_TRX and ESGYNDB parcels,
    a. Download
    b. Distribute
    c. Activate
3. Modify the required Hadoop Configuration settings, and restart services as necessary (at least HBase).
4. Use the cluster Actions menu to Add Service, then select EsgynDB.

Note: ESGYNDB_TRX must be activated on all HBase Regionserver hosts.

8.3.4 Creating a New EsgynDB Cluster

EsgynDB can also be installed when adding a new cluster.
When selecting parcels, select CDH, ESGYNDB, and ESGYNDB_TRX.
When selecting services, select EsgynDB from the Custom Service list. The dependent services, such as HBase, will automatically be included.

In this scenario, EsgynDB service will give an error and fail to start until HBase required configuration
settings are deployed and HBase is restarted, as indicated in Hadoop Configuration.

When EsgynDB start-up gives the TRX configuration error, Cloudera Manager will give options to Retry or go Back. Instead, click on the title menu bar to go to the main screen. From there you can configure Hadoop and restart services as necessary.
That should also start EsgynDB, but you must then use the EsgynDB Actions menu to select Initialize EsgynDB MetaData.

8.3.5 Host Selection

8.3.5.1 EsgynDB Node

This is the main worker role, which runs database queries and interfaces with other Hadoop components (HDFS, HBase, Hive). EsgynDB nodes are usually placed on HBase Regionserver nodes, but may be placed on any HBase Gateway, Regionserver, or Master nodes. EsgynDB nodes also act as Hive clients, so need to be co-located with Hive Gateway or Server roles.

8.3.5.2 EsgynDB Connection Master

This role routes outside requests (JDBC/ODBC connections) to specific nodes. For availability, more than one node should be selected. These roles must be co-located with EsgynDB Nodes.

If configured, a floating IP address can be used to refer to the active DCS master node. Use of this feature will use the "sudo" command for a couple of network administration commands. Otherwise, the client driver (JDBC/ODBC) needs to be configured with a list of the server nodes instead of a single floating IP address. A third option is to maintain a floating IP address using the keepalived package. EsgynDB will provide a sample configuration file, but keepalived must be configured separately, since Cloudera Manager does not manage it.

8.3.5.3 DB Manager

This role provides an administrative interface to monitor and interact with your EsgynDB cluster. This role is automatically co-located with EsgynDB Connection Master.

8.3.6 Hadoop Configuration

EsgynDB depends on certain configuration settings for HDFS, Zookeeper, and HBase. They are all important for proper operation, but EsgynDB will not start at all without the required settings. The cluster administrator must change these values.

They can be modified manually via the Cloudera Manager web interface, or via the provided tool.

8.3.7 Required

These HBase settings are required before starting EsgynDB service:

• Region Server Advanced Configuration Snippet for hbase-site.xml

Name	Value
hbase.hregion.impl	org.apache.hadoop.hbase.regionserver.transactional.TransactionalRegion
hbase.regionserver.region.split.policy	org.apache.hadoop.hbase.regionserver.ConstantSizeRegionSplitPolicy

Note: The ESGYNDB_TRX parcel must be activated on all HBase Regionserver nodes, or these settings will cause HBase errors.

8.3.8 Important

These HBase settings are recommended:

Name	Value
hbase.regionserver.lease.period	**1 (hour)**
hbase.hregion.memstore.flush.size	512 (MiB)

hbase.hregion.memstore.block.multiplier	7
hbase.hstore.blockingStoreFiles	200
hbase.rootdir.perms	750
hbase.snapshot.master.timeoutMillis	10 (minutes)
hbase.superuser	trafodion

These HDFS settings are recommended:

dfs.namenode.acls.enabled	true

These Zookeeper settings are recommended:

maxClientCnxns	0

If Sentry service is enabled:

sentry.service.admin.groups	trafodion
sentry.service.allow.connect	trafodion

8.3.9 Automated Configuration

The required and important Hadoop settings can be set via the cm_settings.sh tool. The tool can be run on any linux system with network access to the cluster manager host. Use the -h option for a full list of options.

The tool presents the current and proposed settings of each parameter and prompts for confirmation unless a force option is given to skip the prompts.

Example Usage

```
cm_settings.sh -s https://myhost.domain.net:7180 -c "Cluster5" -u joe_admin
```

This tool requires you to authenticate with Cloudera Manager. After initial authentication, a session cookie is used for the remainder of the script and then deleted. Three options are available to provide the password:

1. Default option allows the underlying "curl" command to prompt for your password.
2. The -n option allows lookup of the password in a standard ~/.netrc file.
3. The -p option allows a password to be passed in via the command line. This option is not recommended as it is less secure.

Use the Cloudera Manager web interface to restart services as needed after using this tool.

8.3.10 EsgynDB Configuration

8.3.10.1 License Key

When installing the EsgynDB service, you will be prompted for an EsgynDB License Key.

In 2.5.x release, a valid 2.3.x or 2.4.x license may be used.

For initial installation, use the license key provided by Esgyn Corporation if this is a non-production use of EsgynDB. For initial installation for production usage, leave the license key blank and continue with installation. A short grace period is allowed to obtain a license key.

To obtain a license key, you need to send a copy of the file `/etc/trafodion/esgyndb_id` file to Esgyn. It is a very small binary file which provides a unique cluster identity.

8.3.10.2 Others

Other features may be configured after initial installation, including LDAP connection and Connectivity HA. If Hive is Sentry enabled, be sure to indicate this by setting the "Hive Sentry enabled" (traf.sentry.enabled) parameter to true via the checkbox.

8.3.11 Meta-Data Initialization

After initial installation or after installing a new version of EsgynDB, meta-data needs to be initialized. In Cloudera Manager, go to the EsgynDB service, then use the Actions menu to select Initialize EsgynDB MetaData.

It does not hurt to re-run this command. If initialization steps have already been run, then nothing is done.

If some ODBC/JDBC connections have been made (DBmanager, for example), they may continue to report that the database has not been initialized, even after a successful initialization. This may be cleared by re-starting EsgynDB service, or use the EsgynDB Actions menu to "Force Connections Restart".

8.3.12 EsgynDB Start Up

Even after Cloudera Manager has reported that EsgynDB roles have started successfully, it may take a couple of minutes for all the EsgynDB subsystems to start up and accept requests.

Detailed status may be seen using the Check EsgynDB Status command, found on the EsgynDB service Actions menu

8.3.13 Validate

Perform basic sanity checks.

1. Using the EsgynDB Enterprise Conversational Interface (sqlci). Create a table with a few records. For example:

```
[trafodion@edb001 ~]$ sqlci
EsgynDB Enterprise Conversational Interface 2.5.0
Copyright (c) 2015-2018 Esgyn Corporation
>>create table test1 (f1 int, f2 int);

--- SQL operation complete.
>>insert into test1 values(1,1);

--- 1 row(s) inserted.
>>insert into test1 values(2,2);

--- 1 row(s) inserted.
```

```
>>select * from test1;

F1            F2
-----------   -----------
          1             1
          2             2

--- 2 row(s) selected.
>>get tables;

Tables in Schema TRAFODION.SEABASE
==================================

TEST1

--- SQL operation complete.
>>exit;
```

2. Run the `mgblty_check` tool to verify all the manageability components (EsgynDB Manager, TSD, TCollector and Bosun) are running

```
[trafodion@edb001 ~]$ mgblty_check

Status of OpenTSD...(Expect 1 per node)
Total count of TSD process: 12

Status of Tcollector...(Expect 1 per node)
Total count of tcollector process:  12

Status of DBMgr...(Expect 1 per DB Manager node)
Total count of dbmgr process: 1

Status of Bosun...(Expect 1 per DB Manager node)
Total count of bosun process: 1
```

3. Open a web browser and connect to http://localhost:4205 to verify that EsgynDB Manager is started and you can successfully login. Refer to the EsgynDB Manager User Guide for more details.

4. Download and install the EsgynDB JDBC and/or ODBC drivers on your client workstation to be able to connect to EsgynDB from a client application. For instructions, refer to the Trafodion Client Installation Guide which explains how to install the JDBC and ODBC drivers, connect to an EsgynDB system, and run sample programs to test the connection. The package specific to EsgynDB is named **esgynDB_clients-2.5-RH6-x86_64.tar.gz**

8.4 Ambari Installer

8.4.1 Install Pre-Requisite Software

- If you plan to use local repository to stage the EsgynDB rpms (step 3.2), then install the following packages. Ignore this step, if you plan to stage the rpms in a central repository

```
$ yum install -y yum-utils createrepo
```

- Install the python httplib2 module. This is required for the EsgynDB upgrade steps.

```
$ pip install httplib2
```

8.4.2 Download EsgynDB RPMs

Download these 2 EsgynDB RPMs based on your OS version (RH6 or RH7).

`esgynDB_ambari-2.5.0-1.noarch.rpm (EsgynDB Ambari Management pack)`

`esgynDB-2.5.0-1.x86_64.rpm (EsgynDB server package RPM)`

8.4.3 Copy RPMs to Repository

8.4.3.1 Central Repository

If your organization hosts all the RPMs in a centralized repository, copy and stage the 2 EsgynDB RPMs there.

8.4.3.2 Local Repository

If you want to setup a temporary local repository to stage the EsgynDB rpms, perform the following steps.

8.4.3.2.1 Copy EsgynDB RPMs

Copy the 2 EsgynDB RPMs, to a temp directory.

For example, /opt/esgyn/RH6

```
esgynDB-2.5.0-1.x86_64.rpm
esgynDB_ambari-2.5.0-1.noarch.rpm
```

8.4.3.2.2 Create Local Repository for EsgynDB

Run the following commands to create a local repo metadata for EsgynDB RPMs.

```
$ cd /opt/esgyn/RH6
$ createrepo -d .
```

This should create the repodata sub-directory.

```
[root@hdp263-1 RH6] # createrepo -d .
Spawning worker 0 with 2 pkgs
Workers Finished
Gathering worker results

Saving Primary metadata
Saving file lists metadata
Saving other metadata
Generating sqlite DBs
Sqlite DBs complete
```

8.4.3.2.3 Start local webserver

Start a local webserver to host the EsgynDB repository via a web URL.

```
$ cd /opt/esgyn/RH6
python -m SimpleHTTPServer 8900
```

8.4.4 Install EsgynDB Ambari Management Pack

The instructions in this section are to be performed on the host where your Ambari Server is running and logged in as **root** user.

8.4.4.1 Create Esgyn.repo file

- Create an Esgyn.repo file under /etc/yum.repos.d with the following contents. Replace the IP address and port with the host where you are hosting the Esgyn RPMs.

```
[esgyn]
name=Esgyn
baseurl=http://<IP Address>:<port>
enabled=1
gpgcheck=0
```

8.4.4.2 Install EsgynDB Ambari management pack

- Run the following command to install the EsgynDB Ambari management pack on the Ambari server host.

```
$ yum install esgynDB_ambari
```

- Restart Ambari server

```
$ ambari-server restart
```

8.4.5 Install EsgynDB Service

8.4.5.1 Update HBase settings

Perform these 2 configuration changes before you install EsgynDB. Otherwise you will get an error during the Add Service step.

- EsgynDB recommends a value of at least 200 for hbase.hstore.blockingStoreFiles property.

- Add trafodion user to the hbase supergroup. The superuser list is comma separated.

- Click Save and restart HBase

8.4.5.2 Add EsgynDB repo URL to Ambari

- Login to Ambari and click on **Admin -> Stacks and Versions**
- Click on Versions tab
- Click on **Manage Versions**
- Click on the current version name that is installed.

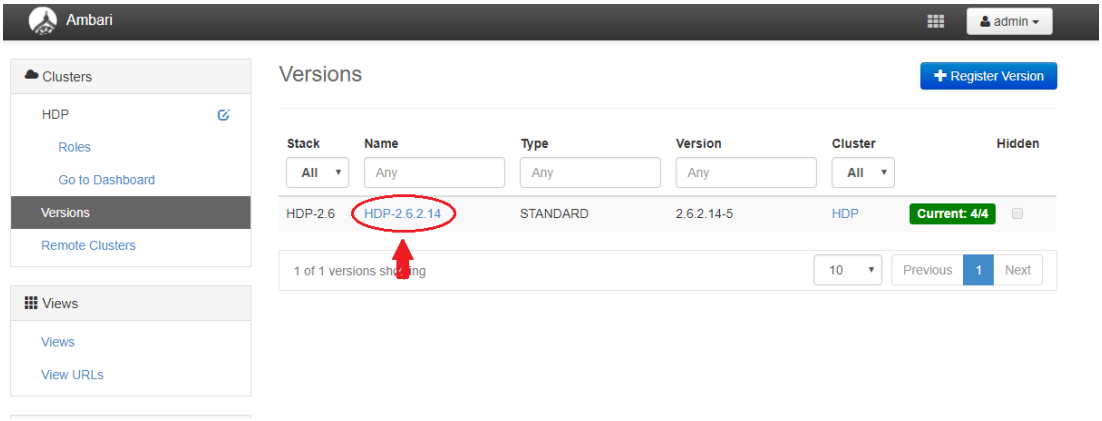

- Edit the default URL for the Esgyn repo and change it to http://<IP Address>:<port> that you set in step Start local webserver

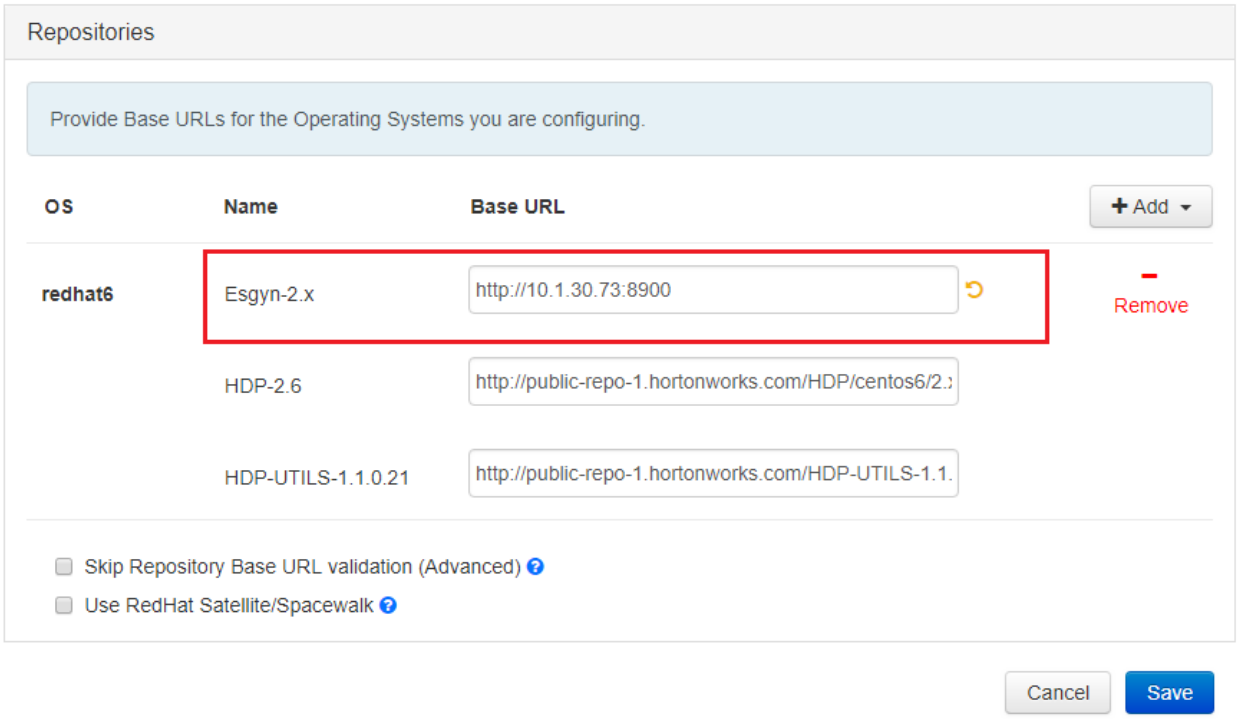

- Click Save

8.4.5.3 Add EsgynDB service

- Go to main Ambari page
- Click on **Actions -> Add Service** to launch the Add Service Wizard.

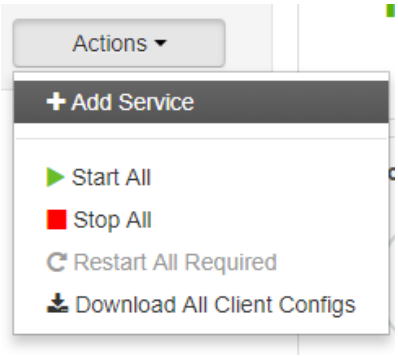

- In the Choose Service page, Select EsgynDB Service

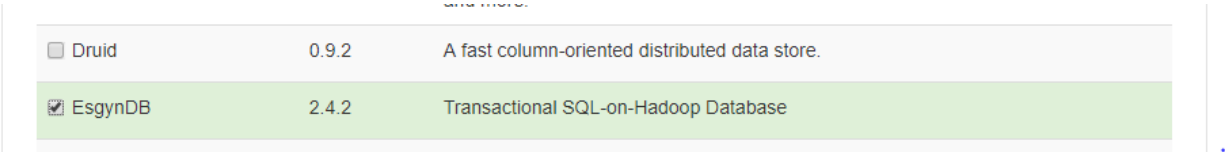

- In the Assign Masters page, Select the hosts for the EsgynDB Manager and Connection Master roles

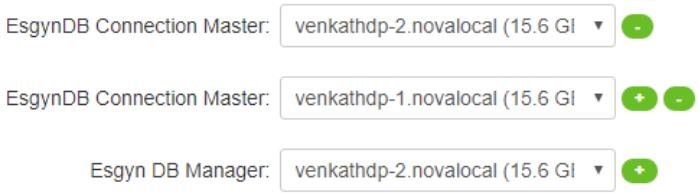

- In the Assign Slaves and Clients page, make sure EsgynDB Node is checked for all the nodes where you want to run EsgynDB

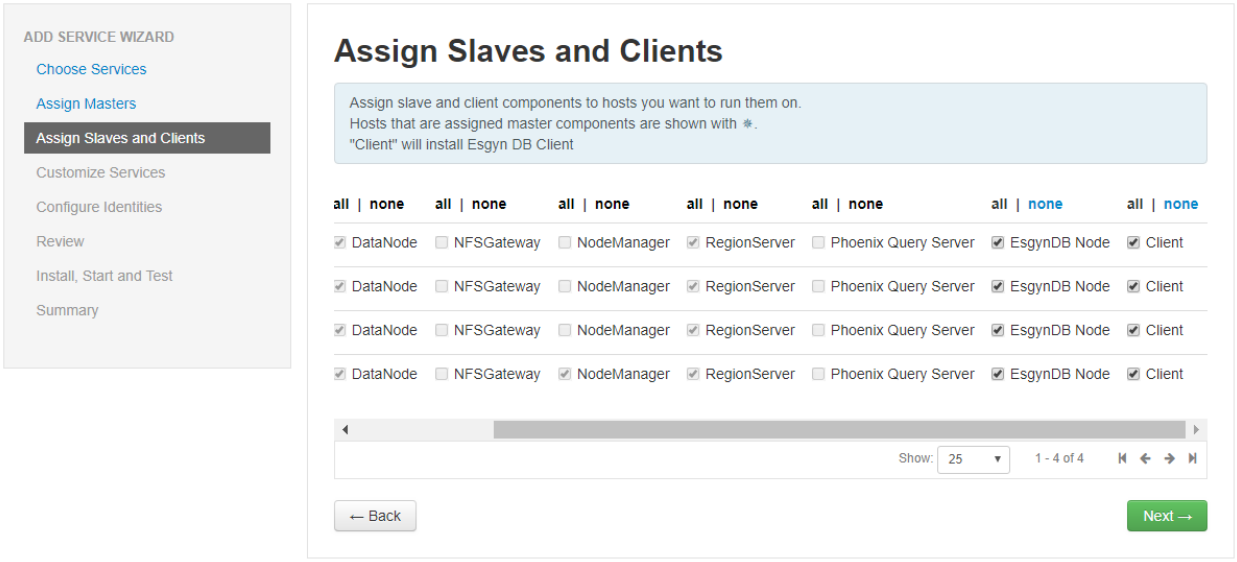

- In the Customize Services page, fill in the following configuration attributes for EsgynDB service
    - EsgynDB License Key
    - EsgynDB DB Admin User
    - EsgynDB DB Root User
    - EsgynDB Admin Password (Advanced tab)
    - LDAP Settings (optional)
        - LDAP authentication enabled
        - LDAP server list
        - LDAP port number)
        - LDAP unique identifiers
        - LDAP search user name
        - LDAP search password
        - LDAP search group base
        - LDAP search group member attribute
        - LDAP search group name attribute
        - LDAP search group object class
- In the Dependent Configurations page, accept the recommended values
- If you cluster is Kerberos enabled, review the Configure Identities page and click next
- In the Review page, check the configuration is correct and click Deploy
- Enter the KDC admin credentials if prompted.
- The install will start and deploying the EsgynDB components to the nodes

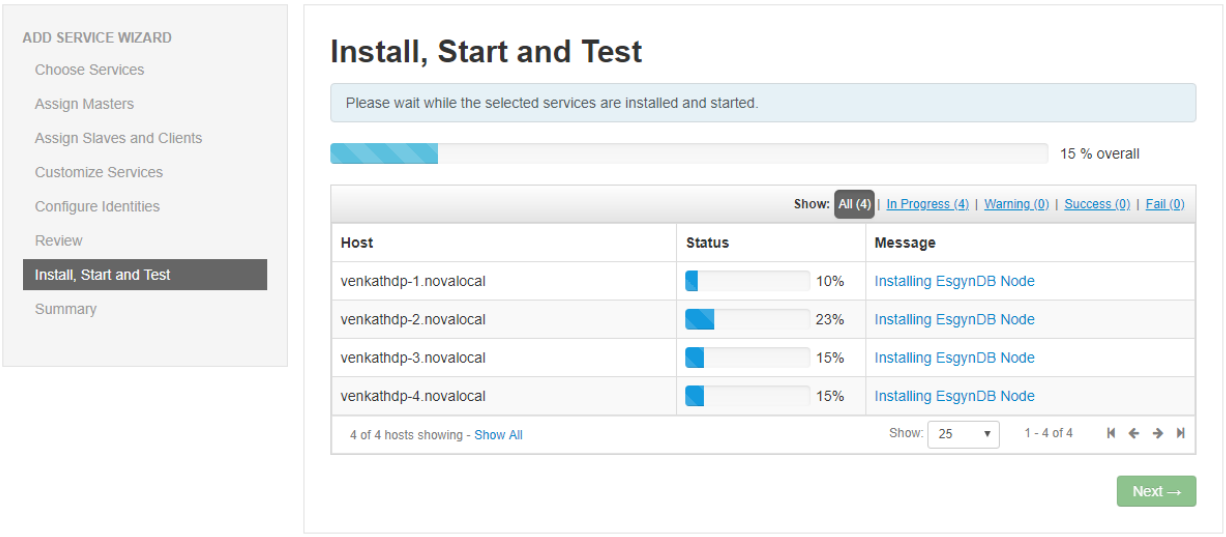

8.4.5.4 Post Install Steps

- Once the install completes, restart the affected dependent services using the **Actions -> Restart All Required** menu.

- Check the Status of EsgynDB service

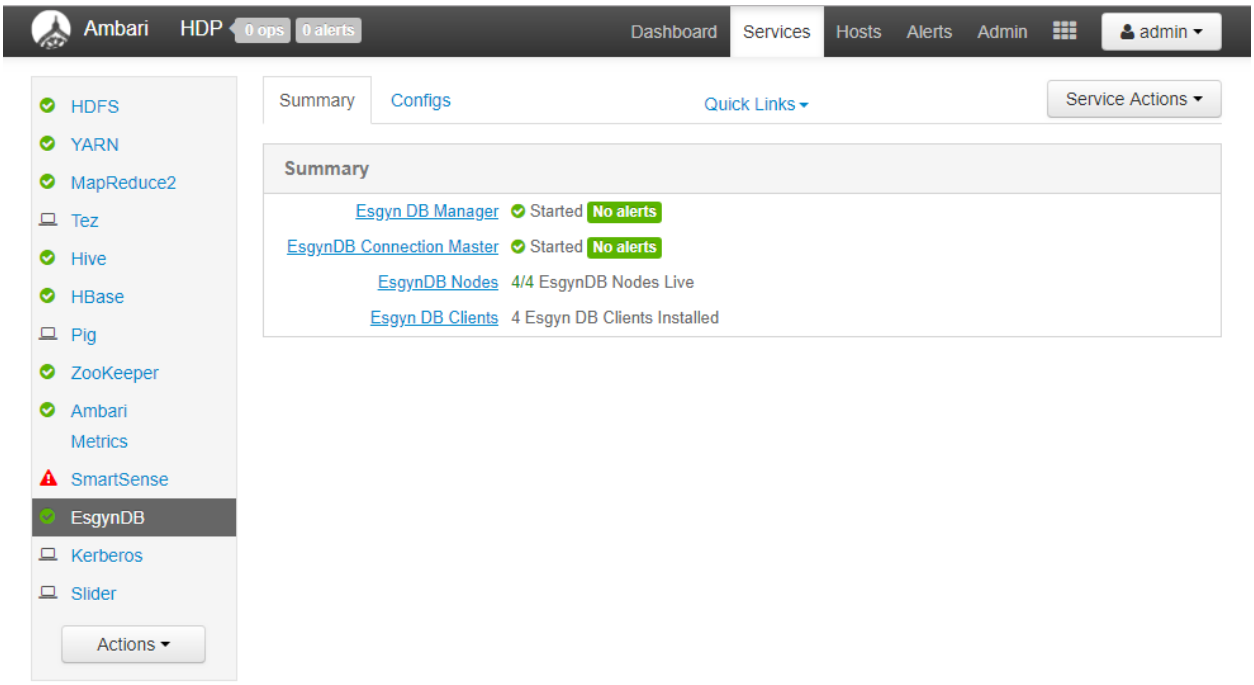

- Initialize the EsgynDB metadata using the **Service Actions -> Initialize** menu

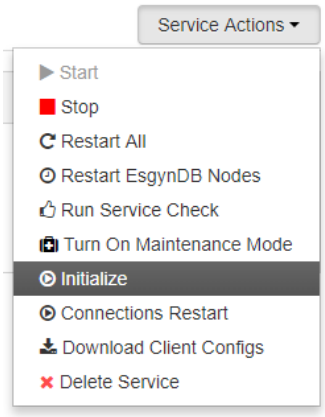

- Restart EsgynDB Connection Master using the **Service Action -> Connections Restart** menu

8.4.6 Upgrade EsgynDB Service

8.4.6.1 Stop EsgynDB Service

Stop EsgynDB service using the **Service Action -> Stop** in the Ambari console.

8.4.6.2 Download and stage EsgynDB updates

Download the 2 updated EsgynDB packages (EsgynDB Ambari management pack rpm and EsgynDB server package rpm).

```
esgynDB-2.5.0-2.x86_64.rpm
esgynDB_ambari-2.5.0-2.noarch.rpm
```

8.4.6.2.1 If using Central Repository

Propagate the updated EsgynDB RPMs to your central repository as set in the Esgyn.repo file (step 3.1)

8.4.6.2.2 If using Local Repository

Login as root user and copy the 2 updated rpms to the local directory where you are staging the EsgynDB local repo. (step 2.2.2). Recreate the repodata and start the local webserver to host the EsgynDB repo URL.

```
$ cd /opt/esgyn/RH6
$ cp /tmp/esgynDB-2.5.0-2.x86_64.rpm .
$ cp /tmp/esgynDB_ambari-2.5.0-2.noarch.rpm .
$ createrepo -d .
$ python -m SimpleHTTPServer 8900
```

8.4.6.3 Update EsgynDB Ambari management pack

- Open a new Linux shell as root user on the Ambari Server node and execute the following commands to update the EsgynDB Ambari management pack

```
$ yum erase -y esgynDB_ambari
$ yum clean all --disablerepo="*" --enablerepo="Esgyn-2.x*"
$ yum install esgynDB_ambari
```

- Restart Ambari server

```
$ ambari-server restart
```

8.4.6.4 Update EsgynDB server package

Ambari currently does not provided a facility to update custom add-on services. You'll need to run the upgrade_esgynDB.py script that is packaged with the EsgynDB Ambari management pack.

```
$ cd /opt/Esgyn
$ ./upgrade_esgynDB.py
```

- Specify the Ambari web URL.
- Enter the Ambari admin user name
- Enter the Ambari admin password

    - The password is only used to make the REST API call and is not stored anywhere
- The script will make REST calls to Ambari server and trigger yum upgrade of the EsgynDB server package on each of the hosts where EsgynDB is installed.

```
$ Enter the ambari url. (include http or https) : http://10.1.30.73:8080
Enter the ambari admin user name : admin
Enter the admin password :
Sending upgrade request to all EsgynDB hosts for component ESGYNDB_NODE
Sending upgrade request to all EsgynDB hosts for component ESGYNDB_CLIENT
Check EsgynDB upgrade results in Ambari console...
```

- Login to Ambari console and check the operations status page to monitor the progress of the EsgynDB upgrade.

8.4.6.5 Restart HBase

The symbolic links for the EsgynDB HBase TRX jar file is updated. So, restart HBase service using Ambari.

8.4.6.6 Start EsgynDB

Start EsgynDB using Ambari.

9. Uninstall

There may be times when you need to uninstall EsgynDB. Before you do so, make sure the data is saved away.

Use the Trafodion Provisioning User id to uninstall EsgynDB.

Run the commands from the first node of the cluster. Do not run them from a machine that is not part of the EsgynDB cluster.

You do not need to uninstall an existing version of EsgynDB if upgrading to a newer version.

9.1 Stop EsgynDB

Do the following

```
$ su trafodion
$ cd $TRAF_HOME/sql/scripts
$ sqstop
$ exit
```

9.2 Uninstall

Execute the `db_uninstall.py` script to completely remove EsgynDB.

10. Troubleshooting

If you are not able to start up the environment or if there are problems running trafci, then verify that all the EsgynDB processes are up and running.

`trafcheck` should indicate all processes are running.

If processes are not running as expected, then:

- `sqstop` to shut down EsgynDB. If some EsgynDB processes do not terminate cleanly, then run `ckillall`.
- `sqstart` to restart EsgynDB.

If problems persist please review logs:

$TRAF_HOME/logs: EsgynDB logs.

11. Enabling Security Features

EsgynDB supports user authentication with OpenLDAP and Active Directory (AD), integrates with Hadoop's Kerberos environment and supports authorization through database grant and revoke requests (privileges).

AD/LDAP is optional and can be configured by running the EsgynDB installer. If Kerberos is enabled in Hadoop, then the EsgynDB installer will ask questions needed to configure Kerberos for EsgynDB

- If Hadoop has enabled Kerberos, then EsgynDB must also enable Kerberos.
- If AD/LDAP is enabled, then database authorization (privilege support) is automatically enabled.
- If Kerberos is not enabled, then enabling AD/LDAP is optional.
- If multi-tenancy is enabled, then database authorization is automatically enabled.

11.1 Configuring EsgynDB for Kerberos

Kerberos is a protocol for authenticating a request for a service or operation. It uses the notion of a ticket to verify accessibility. The ticket is proof of identity encrypted with a secret key for the particular requested service. Tickets exist for a short time and then expire. Therefore, you can use the service as long as your ticket is valid (i.e. not expired). Hadoop uses Kerberos to provide security for its services, as such EsgynDB needs to function properly with Hadoop systems that have Kerberos enabled.

11.1.1 Kerberos configuration file

It is assumed that Kerberos has already been set up on all the nodes by the time EsgynDB is installed. This section briefly discusses the Kerberos configuration file for reference.

The Kerberos configuration file defaults to /etc/krb5.conf and contains, among other attributes:

```
* log location: location where Kerberos errors and other information are logged
* KDC location: host location where the KDC (Key Distribution Center) is located
* admin server location: host location where the Kerberos admin server is located
* realm: the set of nodes that share a Kerberos database
* ticket defaults: contains defaults for ticket lifetimes, encoding, and other attributes
```

You need to have access to a Kerberos administrator account to enable Kerberos for EsgynDB. The following is an example request that lists principals defined in the Kerberos database that can be used to test connectivity:

```
kadmin -p 'kdcadmin/admin' -w 'kdcadmin123' -s 'kdc.server' -q 'listprincs'
* -p (principal): please replace 'kdcadmin/admin' with your admin principal
* -w (password): please replace 'kdadmin123' with the password for the admin principal
* -s (server location): please replace 'kdc.server' with your KDC admin server location
* -q (command): defines the command to run, in this case principals are returned
```

11.1.2 Ticket Management

When Kerberos is enabled in EsgynDB, the security installation process:

- Adds a Trafodion principal in Kerberos, one per node with the name trafodion/hostname@realm.
- Creates a keytab for each principal and distributes the keytab to each node. The keytab name is the same for all nodes and defaults to a value based on the distribution, for example: etc/trafodion/keytab/trafodion.keytab.
- Performs a "kinit" on all nodes in the cluster for the trafodion user.
- Adds commands to perform "kinit" and starts the EsgynDB ticket renewal process on each node.

The Hadoop renewal service, renews the Kerberos TGT (ticket granting ticket) up until the maximum number of renewals allowed. So if your ticket lifetime is one day and the number of renewals is seven days, your ticket is good for 7 days. After the number of renewals is expired, the ESGYNDB renewal service re-initializes the ticket. Therefore, the ticket should never expire.

The EsgynDB renewal service runs on each node in the cluster and monitors the status of the Esgyn Kerberos TGT.

EsgynDB provides a script that reports the status of Kerberos TGT (trafkerberos):

```
$TRAF_HOME/sql/scripts/trafkerberos { status | stop | start }
status:  reports the status of the Kerberos ticket across all nodes:
stop:  stops the EsgynDB ticket renewal process
start:  starts or restarts the EsgynDB ticket renewal process.
```

11.1.3 Kerberos installation

The EsgynDB installation script automatically determines if Kerberos is enabled on the node. If it is enabled, then the environment variable SECURE_HADOOP is set to "Y". The installer then gathers the following information needed to integrate Kerberos with EsgynDB:

```
KDC server address
KDC admin principal
KDC admin password
Max lifetime for the trafodion user
Max renew lifetime for the trafodion user
```

KDC admin password will be saved only in configuration file **db_config.bakYYMMDD_HHMM** in the installer folder when installation is completed. You can delete this file for secure perspective. NOTE: Keytab files are auto detected by installer in CDH/HDP cluster.

11.2 Configuring LDAP

EsgynDB does not manage usernames and passwords internally but does support authentication via directory servers that use the OpenLDAP and Active Directory (AD/LDAP) protocols also known as AD/LDAP servers. You can configure the AD/LDAP servers during installation by

answering the installer's prompts. For more information, see Configuring AD/LDAP Servers for details. Installing AD/LDAP also enables database authorization (privilege support).

Once authentication and authorization are enabled, EsgynDB allows users to be registered in the database and allows privileges on objects to be granted to users and roles (which are granted to users). EsgynDB also supports component-level (or system-level) privileges, such as MANAGE_USERS, which can be granted to users and roles. See Managing Users.

If you do not enable AD/LDAP in EsgynDB, then a client interface to EsgynDB may request a user name and password, but Trafodion ignores the user name and password entered in the client interface, and the session runs as the database **root** user, **DB__ROOT**, without restrictions. If you want to restrict users, restrict access to certain users only, or restrict access to an object or operation, then you must enable security, which enforces authentication and authorization.

11.3 Configuring AD/LDAP Servers

The EsgynDB installer sets up and propagates the AD/LDAP configuration file called `.traf_authentication_config` located in `$TRAF_HOME/sql/scripts`. This file is a flat file, organized as a series of attribute/value pairs.

A sample template file is located in `$TRAF_HOME/sql/scripts/traf_authentication_config`.

11.3.1 AD/LDAP configuration file (.traf_authentication_config)

The configuration file is organized as follows:

```
SECTION: Defaults
  DefaultSectionName: local
  RefreshTime: 1800
```

TLS_CACERTFilename:

SECTION: local

Section configuration attributes

SECTION: aaa

Section configuration attributes

. . .

SECTION: zzz

Section configuration attributes

The "local" section is required. Other sections are optional and can be specified if there is a need for multiple AD/LDAP configurations. For example, users connecting as tenant1 would authenticate with AD/LDAP servers defined in section "aaa", and users connecting as tenant2 would authentication with AD/LDAP servers defined in section "bbb".

Each SECTION in the configuration file requires at a minimum:

1. AD/LDAP Host name(s)

   One or more names of hosts that support the OpenLDAP or Active Directory protocol must be specified. EsgynDB will attempt to connect to all provided host names during the authentication process. The set of usernames and passwords should be identical on all hosts to avoid unpredictable results. The attribute name is LDAPHostName.

   Example:

   ```
   LDAPHostName:  ldap.company.com
   ```

2. LDAP Port number

   Port number of the AD/LDAP server. Typically this is 389 for servers using no encryption or TLS, and 636 for servers using SSL. The attribute name is LDAPPort.

Example:

```
LDAPPort:  389
```

3. LDAP Unique Identifier

   Attribute(s) used by the directory server that uniquely identifies the username. You may provide one or more unique identifier specifiers. The attribute name is `UniqueIdentifier`.

   Example:

   ```
   UniqueIdentifier:  uid=,ou=users,dc=com
   ```

4. Encryption level

   A numeric value indicating the encryption scheme used by your AD/LDAP server. The attribute name is LDAPSSL and values are:

   0: Encryption not used

   1: SSL

   2: TLS

   Example:

   ```
   LDAPSSL: 2
   ```

   If your AD/LDAP server uses TLS you must specify a file containing the certificate used to encrypt the password. By default, the EsgynDB software looks for this file in `$TRAF_HOME/cacerts`, but you must specify a fully qualified filename, or set the environment variable CACERTS_DIR to another directory. To specify the file containing the certificate, you set the value of the attribute TLS_CACERTFilename, located in the Defaults section.

   Examples:

   ```
   TLS_CACERTFilename: mycert.pem
   TLS_CACertFilename: /usr/etc/cert.pem
   ```

5. Encryption level

   Some AD/LDAP servers require a known username and password to search the directory of usernames. If your environment has that requirement, provide these “search” values.

   Examples:

   ```
   LDAPSearchDN: lookup@company.com
   LDAPSearchPwd: Lookup123
   ```

6. If the configured AD/LDAP server requires a user group to perform name lookup, the following attributes are required:

   LDAPSearchGroup attributes

   a. LDAPSearchGroupBase
   b. LDAPSearchGroupObjectClass
   c. LDAPSearchGroupMemberAttr
   d. LDAPSearchGroupNameAttr

There are additional optional attributes that can be used to customize EsgynDB authentication. Please see the sample configuration file $TRAF_HOME/sql/scripts/traf_authentication_setup for more details.

The EsgynDB installation script only sets up the "local" section and does not setup group attributes. To configure multiple sections or group support, the .ldap_authentication_config file has to be modified and then copied to all nodes.

11.3.2 ldapconfigcheck script

You can test the authentication configuration file for syntactic errors using the utility ldapconfigcheck. If you have loaded the EsgynDB environment (sqenv.sh), then the utility will automatically check the file at $TRAF_HOME/sql/scripts/.traf_authentication_config. If not, you can specify the file to be checked. EsgynDB does not need to be running to run the utility.

```
ldapconfigcheck  [<option>]...
<option> ::= --help|-h : display usage information
             -file <config-filename>
```

If the configuration filename is not specified, the tool will look for a file using environment variables. Those environment variables and the search order are:

1. TRAFAUTH_CONFIGFILE

   A fully qualified name is expected.
2. TRAFAUTH_CONFIGDIR

   Filename `.traf_authentication_config/` is appended to the specified directory
3. TRAF_HOME

   `/sql/scripts/.traf_authentication_config` is appended to the value of TRAF_HOME.

Example:

```
ldapconfigcheck –file myconfigfile
File myconfigfile is valid.
```

If an error is found, the line number with the error is displayed along with the error.

11.3.3 ldapcheck utility

You can test the AD/LDAP connection using the utility `ldapcheck`. To use this utility the EsgynDB environment must be loaded (`sqenv.sh`), but the EsgynDB instance does not need to be running.

To test the connection only, you can specify any username or group name, and a name lookup is performed using the attributes in `.traf_authentication_config`.

```
ldapcheck [option]...
option ::= --help|-h

           --username=<LDAP-username>
           --password[=<password>]

  or       --groupname=<LDAP-groupname>
           --confignumber=<config-section-number>
           --configname=<config-section-name>

           --verbose
```

You can get additional error detail by including the `--verbose` option. The ldapcheck utility logs events to the $TRAF_HOME/logs directory in file names with the following format: dbsecurity_<host>_<pid>.log

If you supply a password, `ldapcheck` attempts to authenticate the specified username and password. The example below shows the password for illustrative purposes, but to avoid typing the password on the command line, leave the password blank (`--password=`) and the utility will prompt for the password with no echo.

Example:

```
ldapcheck --username=user1 --password=user1passwd
Authentication request: externalName user1, configName 'local' (configNumber
0), result 0 (Authentication successful)

Member of group: group1
```

11.4 Generating a Server Certificate

EsgynDB uses certificates to encrypt and decrypt passwords as part of authenticating users. EsgynDB also uses certificates as part of its HTTPS support for web applications. By default, a self-signed certificate is generated using OpenSSL and stored on each node of the cluster in the `$HOME/sqcert` directory. Optionally, a CA signed certificate can be used.

11.4.1 Self-signed certificates

Self-signed certificates are an identity certificate that is signed by the same entity whose identity it certifies. At installation and upgrade time, a self-signed certificate is generated and placed in the `$HOME/sqcert` directory on each node in the cluster. The certificate is valid for 365 days and stored in the following files:

- `server.crt` – is the certificate
- `server.key` – is the private key
- `server.keystore` – is a Java KeyStore (JKS) which is a repository of security certificates used for instance in SSL encryption.

Note: A Java Keystore is a container for authorization certificates or public key certificates and is often used by Java-based applications for encryption, authentication, and serving over HTTPS. Its entries are protected by a keystore password. A keystore entry is identified by an *alias*, and it consists of keys and certificates that form a trust chain.

There is no specific action required for self-signed certificates. When support personnel get notified that the current active certificate (if self-signed was active on the cluster) has expired then these self-signed certificates need to be manually regenerated. To regenerate self-signed certificates then run script `sqcertgen a`nd `sqcertget gen_keystore` and restart connectivity and manageability services (`dcsstop, mgblty_stop, dcsstart, mgblty_start`).

Validity of the certificate is known by running the script `certcheck`.

11.4.2 Generate a CSR to obtain a signed certificate from Certificate Authority (CA)

Using the existing server private key, the following command generates a CSR (Certificate Signing Request): `sqcertgen gen_csr`

Next, send this CSR to the Certificate Authority (CA) to be signed. Once the CSR has been signed, you will have a real signed certificate. This signed certificate can then be distributed to the cluster.

11.4.3 CA signed certificates

Certificate Authority (CA) is an entity that issues digital certificates. A digital certificate certifies the ownership of a public key by the named subject of the certificate. This allows others (relying parties) to rely upon signatures or on assertions made about the private key that corresponds to the certified public key. A CA acts as a trusted third party—trusted both by the subject (owner) of the certificate and by the party relying upon the certificate. The format of these certificates is specified by the X.509 standard.

Distributing CA signed certificate is a post install step after EsgynDB has been completed. To distribute signed certificate, then run the script `distcacert casigned<pem file name>` and restart connectivity and manageability services (`dcsstop, mgblty_stop, dcsstart, mgblty_start`).

Both the public (`server.crt`) and private (`server.key`) files should be placed in the directory `$HOME/sqcert`.

11.5 Managing Users

Kerberos is enabled for installations that require a secure Hadoop environment. AD/LDAP is enabled to enforce authentication for any user connecting to Trafodion. The Trafodion database enforces privileges on the database, database schemas, database objects (table, views, etc) and database operations. Privileges are enforced when authorization is enabled. When AD/LDAP or Kerberos is enabled, authorization is automatically enabled.

To determine the status of authentication and authorization, bring up sqlci and perform "env;".

```
>>env;
---------------------------------
Current Environment
---------------------------------
AUTHENTICATION     enabled
AUTHORIZATION      enabled
CURRENT DIRECTORY  /opt/trafodion/esgynDB-2.3.0
. . .
```

Once authorization is enabled, there are two predefined database users called DB__ROOT and DB__ADMIN. These users are associated with your specified AD/LDAP username that was set up during install. Please connect to the database as one of these users setup required schemas, users, roles, and privileges.

To learn more about how to register users, grant object and component privileges, and manage users and roles, please see the EsgynDB SQL Reference Manual.

12. Securing the installation

The following optional steps can enhance the security of your EsgynDB installation.

12.1 Secure Linux

Refer to instructions from Red Hat or CentOS to secure your Linux installation.

OS	Version	Link
RedHat Enterprise Linux	7.x	https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/index
RedHat Enterprise Linux	6.x	https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/security_guide/index
CentOS	Generic	https://wiki.centos.org/HowTos/OS_Protection#head-f80d332aeea03f57d34d7a5c09493a7d69cce177

12.2 Secure Hadoop

Refer to instructions from your Hadoop distribution vendor – Cloudera or Hortonworks.

Hadoop Distribution	Link
Hortonworks' HDP 2.x	https://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.4.0/bk_Security_Guide/content/ch_hdp-security-guide-overview.html
Cloudera CDH 5.x	https://www.cloudera.com/documentation/cdh/5-0-x/CDH5-Security-Guide/CDH5-Security-Guide.html

12.3 Secure Jetty Server

A number of components in EsgynDB serve web pages using the Jetty web server module. This server is configured to use HTTPS and enabled to use strong SSL ciphers. Refer to the Jetty security documentation.

- Configure Jetty Connectors
- Configure Security

12.4 Upgrade passwords

Change the default passwords with stronger passwords wherever applicable. Refer to this section on user-ids and passwords for specific users that are used by EsgynDB.

12.5 Secure ports

The following ports are typically required to be opened to external applications

Application	Port Range	Description
DCS Master	23400	Open a range of consecutive ports depending on the number of configured MXOSRVRs
DB Manager	4206	

12.6 Secure AWS Installation

You can secure your EsgynDB installation on AWS by implementing the following recommendations.

12.6.1 Restrict access to Ambari or Cloudera Manager

Ensure the Hadoop manageability tools such as Ambari or Cloudera Manager are only accessible from a defined set of IP addresses. This may be your client machine or machines from within your corporate network.

For example, this configuration grants access to Ambari from client 76.244.44.66, and restricts access from other IP addresses.

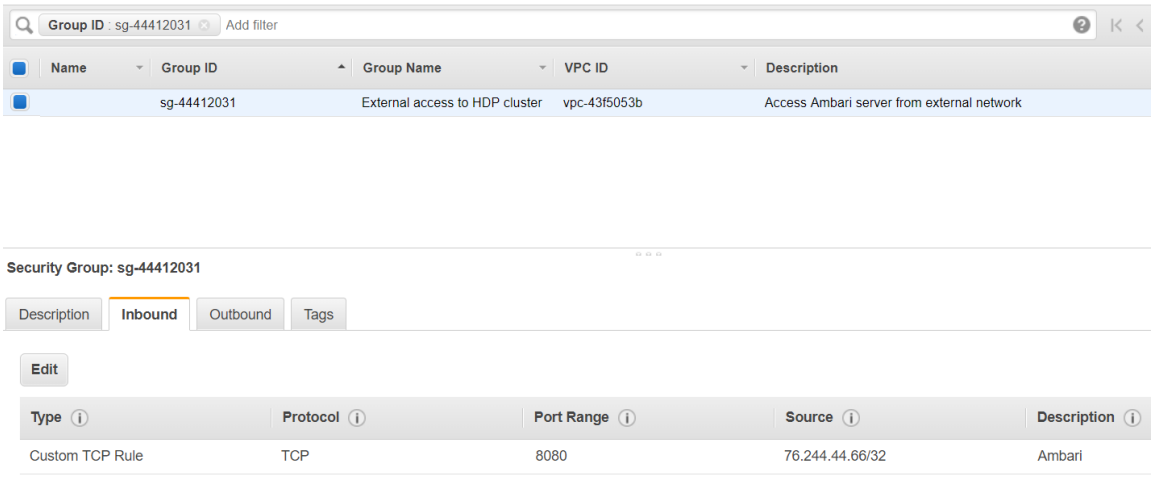

12.6.2 Restrict access to EsgynDB components

Configure AWS security rules to restrict access to the DCS subsystem and EsgynDB Manager.

For example, this configuration grants access only from client 76.244.44.66

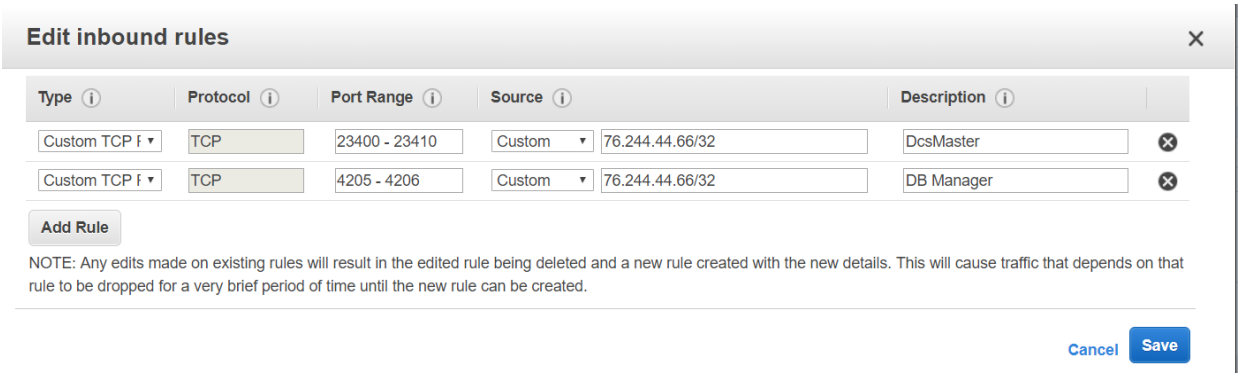

12.6.3 Access within EsgynDB instance

EsgynDB components need to communicate across nodes within the instance, while certain components such as DCSMaster and MXOSRVR processes, need to be available for access from external clients. Configure the rules accordingly, using the following example as a template.

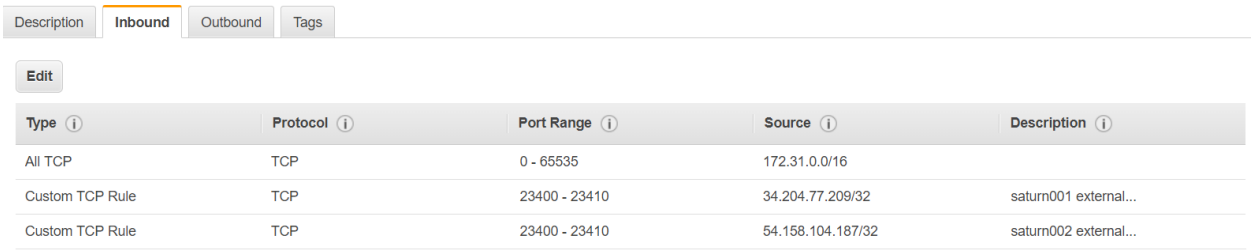

Description | Inbound | Outbound | Tags

Edit

Type	Protocol	Port Range	Source	Description
All TCP	TCP	0 - 65535	172.31.0.0/16	
Custom TCP Rule	TCP	23400 - 23410	34.204.77.209/32	saturn001 external...
Custom TCP Rule	TCP	23400 - 23410	54.158.104.187/32	saturn002 external...

12.6.4 Summary

The example screenshot below captures the summary of all inbound rules as set up on an EsgynDB cluster running Hortonworks' HDP. Access to this cluster is available only from the source address and applicable ports, but not from any other IP addresses.

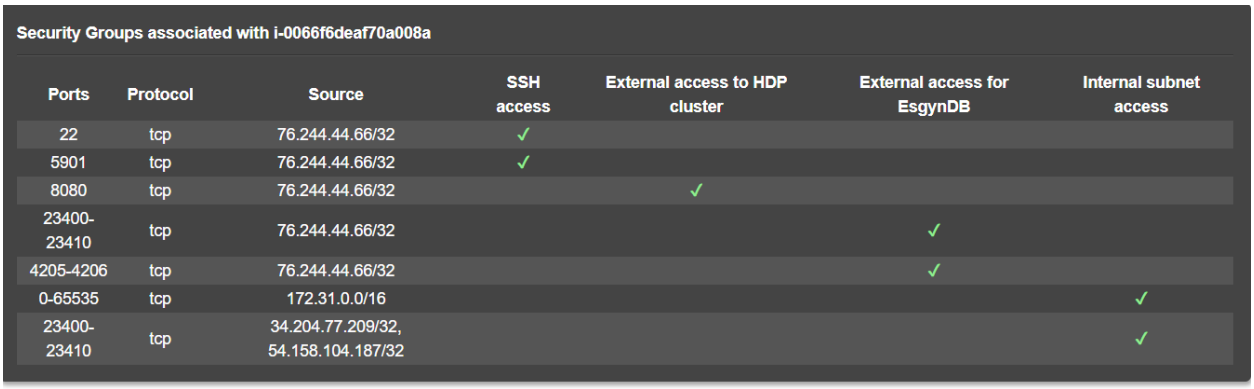

Security Groups associated with i-0066f6deaf70a008a

Ports	Protocol	Source	SSH access	External access to HDP cluster	External access for EsgynDB	Internal subnet access
22	tcp	76.244.44.66/32	✓			
5901	tcp	76.244.44.66/32	✓			
8080	tcp	76.244.44.66/32		✓		
23400-23410	tcp	76.244.44.66/32			✓	
4205-4206	tcp	76.244.44.66/32			✓	
0-65535	tcp	172.31.0.0/16				✓
23400-23410	tcp	34.204.77.209/32, 54.158.104.187/32				✓

12.6.5 Final Steps

Using the AWS console,

- create the Elastic IP

- create the network interface; check the **Allow Re-association** button
- Associate the Elastic IP with network interface
- Create the user ID, and generate the access keys
- Create the policy and attach the policy to the user

12.6.6 Best Practices

Guidance for best practices in setting up IAM, user policies and roles for accessing Amazon's EC2 cloud is available here: http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html

Appendix A: Authentication Configuration File

By default the EsgynDB authentication configuration file is located at `$TRAF_HOME/sql/scripts/.traf_authentication_config`.

Attributes Supported in .traf_authentication_config

This is a list of the attributes supported in the file `.traf_authentication_config`. For each attribute, a description and example is included.

Attribute Name	Purpose	Example Value	Notes
LDAPHostName	Host name of the local LDAP server	ldap.master.com	If more than one LDAPHostName entry is provided, Trafodion will attempt to connect with each LDAP server before returning an authentication error. Also see the description related to RetryCount and RetryDelay entries.
LDAPPort	Port number of the local LDAP server	345	Must be numeric value. Related to LDAPSSL entry. Standard port numbers for

			OpenLDAP are as follows: Non-secure: 389 SSL: 636 TLS: 389
LDAPSearchDN	If a search user is needed, the search user distinguished name is specified here.	cn=aaabbb, dc=demo, dc=net	If anonymous search is allowed on the local server, this attribute does not need to be specified or can be specified with no value (blank). To date, anonymous search is the normal approach used.
LDAPSearchPWD	Password for the LDAPSearchDN value. See that entry for details.	welcome	None
LDAPSSL	A numeric value specifying whether the local LDAP server interface is unencrypted or TLS or SSL. Legal values are 0 for unencrypted, 1 for SSL, and 2 for TLS. For SSL/TLS, see the section	0	None

	below on Encryption Support.		
UniqueIdentifier	The directory attribute that contains the user's unique identifier.	uid=,ou=Users,dc=demo, dc=net	To account for the multiple forms of DN supported by a given LDAP server, specify the UniqueIdentifier parameter multiple times with different values. During a search, each UniqueIdentifier is tried in the order it is listed in the configuration file.
LDAPNetworkTimeout	Specifies the timeout (in seconds) after which the next LDAPHostName entry will be tried, in case of no response for a connection request. This parameter is similar to NETWORK_TIMEOUT in ldap_conf(5). Default value is 30 seconds.	20	The value must be a positive number or -1. Setting this to -1 results in an infinite timeout.
LDAPTimelimit	Specifies the time to wait when performing a search on the LDAP server for the username.	15	The server may still apply a lower server-side limit on the

	The number must be a positive integer. This parameter is similar to TIMELIMIT in ldap_conf(5). Default value is 30 seconds.		duration of a search operation.
LDAPTimeout	Specifies a timeout (in seconds) after which calls to synchronous LDAP APIs will abort if no response is received. This parameter is similar to TIMEOUT in ldap_conf(5). Default value is 30 seconds.	15	The value must be a positive number or -1. Setting this to -1 results in an infinite timeout.
RetryCount	Number of attempts to establish a successful LDAP connection. Default is 5 retries before returning an error.	10	When a failed operation is retried, it will be attempted with each configured LDAP server, until the operation is successful or the number of configured retries is exceeded.
RetryDelay	Specifies the number of seconds to delay between retries. Default value is 2 seconds. See	1	None

	description of RetryCount.		
PreserveConnection	Specifies whether the connection to LDAP server will be maintained (YES) or closed (NO) once the operation finishes. Default value is NO.	YES	None
RefreshTime	Specifies the number of seconds that must have elapsed before the configuration file is reread. Default is 1800 (30 minutes).	3600	If set to zero, the configuration file is never read. The connectivity servers must be restarted for changes to take effect if this value is zero. This attribute is not specific to either configuration and must be defined in the DEFAULTS section.
TLS_CACERTFilename	Specifies the location of the certificate file for the LDAP server(s). Filename can either be fully qualified or relative to $CACERTS_DIR.	cert.pem	This attribute applies to both configurations. If a configuration does not require a certificate, this attribute is ignored. This attribute must be

			defined in the DEFAULTS section.
DefaultSectionName	Specifies the configuration type that will be assigned to a user by the REGISTER USER command if no authentication type is specified. In the initial Trafodion release, only one configuration is supported.	LOCAL	This attribute must be defined in the DEFAULTS section. If the DefaultSectionName attribute is specified, a section by that name (or equivalent) must be defined in .traf_ldapconfig. Legal values are LOCAL and ENTERPRISE. This syntax is likely to change.

Appendix B: ldapconfigcheck Utility

The utility `ldapconfigcheck` validates the syntactic correctness of a EsgynDB authentication configuration file. EsgynDB does not need to be running to run the utility.

```
ldapconfigcheck  [<option>]...
<option> ::= --help|-h : display usage information
             -file <config-filename>
```

Considerations

If the configuration filename is not specified, the tool will look for a file using environment variables. Those environment variables and the search order are:

1. TRAFAUTH_CONFIGFILE

   A fully qualified name is expected.
2. TRAFAUTH_CONFIGDIR

   Filename `.traf_authentication_config/` is appended to the specified directory
3. TRAF_HOME

   `/sql/scripts/.traf_authentication_config` is appended to the value of TRAF_HOME.

Errors

One of the following is output when the tool is run. Only the first error encountered is reported.

Code	Text
0	File *filename* is valid.
1	File *filename* not found
2	File: *filename* Invalid attribute name on line *line-number*
3	File: *filename*

	Missing required value on line *line-number*
4	File: *filename* Value out of range on line *line-number*
5	File: *filename* Open of traf_authentication_config file failed
6	File: *filename* Read of traf_authentication_config file failed
7	No file provided. Either specify a file parameter or verify environment variables.
8	TLS was requested in at least one section, but TLS_CACERTFilename was not provided
9	Missing host name in at least one section. Each LDAP connection configuration section must provide at least one hostname.
10	Missing unique identifier in at least one section. Each LDAP connection configuration section must provide at least one unique identifier.
11	At least one LDAP connection configuration section must be specified.
12	Internal error parsing . traf_authentication_config.