

Upgrade Databases for PostgreSQL

Contents:

- *Prerequisites*
- *Backup*
- *Acquire Distribution and Port Information*
- *Database Prep*
 - *Before you begin*
 - *Using non-default ports*
- *Upgrade PostgreSQL 9.6 to 12*
 - *PostgreSQL path references*
 - *Upgrade for CentOS/RHEL - PostgreSQL 9.6 to 12*
 - *Upgrade for Ubuntu - PostgreSQL 9.6 to 12*
 - *Upgrade on Amazon RDS - PostgreSQL 9.6 to 12*
- *Transfer Settings*
 - *Verify and Cleanup*

This section describes the requirements and pre- and post-upgrade steps for upgrading the PostgreSQL database server of the Trifacta® databases for a working deployment of the Trifacta platform.

When you upgrade to a new release of the platform, the underlying databases and their structures are automatically migrated to the new format. However, the database releases are not touched.

Prerequisites

NOTE: Responsibility for performing the actual upgrade of the database server is the responsibility of the customer. Please review this section before you perform the database server upgrade.

- The installing user must have write permissions to the directory from which the commands are executed.
- The installing user must have sudo privileges.
- Verify that you know the host and port number for each database.

Backup

Before you begin, please verify that you have a valid backup of each Trifacta database. See *Backup and Recovery* in the Admin Guide.

Acquire Distribution and Port Information

Before you begin, you must:

- Acquire the port information for the current database. Typically, this value is 5432.

NOTE: By default, PostgreSQL and the platform use port 5432 for communication. If that port is not available at install/upgrade time, the next available port is used, which is typically 5433. This change may occur if a previous version of PostgreSQL is on the same server. When a non-default port number is used, the platform must be configured to use it. For more information, see *Change Database Port*.

- Acquire the latest distribution for the database software.

For more information, see *Install Databases for PostgreSQL*.

For more information, see *Install Databases for MySQL*.

Database Prep

This procedure describes the process for upgrading the PostgreSQL version in use by the Trifacta databases. This procedure assumes the following:

- All current databases are co-located with the software on the Trifacta node.
- Some downtime of the databases during the upgrade process is ok.

Before you begin

- If you haven't already, please back up each database and all PostgreSQL configuration files. See *Backup and Recovery* in the Admin Guide.
- Login to the node where the databases are hosted.

Using non-default ports

The default port for these commands is port 5432.

If you are using a non-standard port for either the pre- or post-upgrade versions of PostgreSQL, you can use the following parameters with the `pg_upgrade` command:

Parameter	Description
<code>--old-port=<pre-upgrade_port></code>	The port number of the pre-upgrade instance of PostgreSQL
<code>--new-port=<post-upgrade_port></code>	The port number to use for the post-upgrade instance of PostgreSQL

For more information, see <https://www.postgresql.org/docs/12/pgupgrade.html>.

Upgrade PostgreSQL 9.6 to 12

Support for PostgreSQL 9.6 is deprecated. You must upgrade to PostgreSQL 12.

- Source version: PostgreSQL 9.6
- Upgrade version: PostgreSQL 12

NOTE: The database version upgrade must be performed after you have upgraded the software.

NOTE: The database client is installed as part of this upgrade. No further action is required to enable use of the database client.

PostgreSQL path references

In PostgreSQL 10 and later, directory references to versions of PostgreSQL paths have changed.

PostgreSQL 9.6 example path:

```
/var/lib/pgsql/9.6/data/pg_hba.conf
```

PostgreSQL 12 example path:

```
/var/lib/pgsql/12/data/pg_hba.conf
```

For more information, see https://wiki.postgresql.org/wiki/YUM_Installation.

Upgrade for CentOS/RHEL - PostgreSQL 9.6 to 12

For more information, see <https://www.postgresql.org/docs/>.

Upgrade for Ubuntu - PostgreSQL 9.6 to 12

For more information, see <https://www.postgresql.org/docs/>.

Upgrade on Amazon RDS - PostgreSQL 9.6 to 12

NOTE: After upgrading to PostgreSQL on Amazon RDS, you may encounter an I/O error when importing data in which the storage location has the wrong hash version. The solution is to reindex tables. For more information, see <https://aws.amazon.com/blogs/database/postgresql-12-a-deep-dive-into-some-new-functionality/>.

For more information, see https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.PostgreSQL.html.

Transfer Settings

For all operating systems, you must transfer the settings from your old version of PostgreSQL to the new one.

1. Update the port number and any other settings in the platform configuration. To apply this configuration change, login as an administrator to the Trifacta node. Then, edit `trifacta-conf.json`. Some of these settings may not be available through the *Admin Settings Page*. For more information, see *Platform Configuration Methods*.
 - a. Replace all instances of the PostgreSQL port number for the earlier version and replace with the desired port number for the upgrade version.

NOTE: By default, PostgreSQL and the platform use port 5432 for communication. If that port is not available at install/upgrade time, the next available port is used, which is typically 5433. This change may occur if a previous version of PostgreSQL is on the same server. When a non-default port number is used, the platform must be configured to use it. For more information, see *Change Database Port*.

- b. Make the above changes and save the file.
2. Transfer settings from the configuration files for your old database version to the new one. Please review the old and new versions of these files:

NOTE: It is risky to perform a straight copy of these configuration files. Settings may change. New ones may be introduced. Setting values specific to the installation may be overwritten in a copy. Please retain a backup of both versions of each file before migrating settings.

Path to PostgreSQL 9.6 file	Path to PostgreSQL 12 file
/var/lib/pgsql/9.6/data/pg_hba.conf	/var/lib/pgsql/12/data/pg_hba.conf
/var/lib/pgsql/9.6/data/postgresql.conf	/var/lib/pgsql/12/data/postgresql.conf

3. Start the service:

a. CentOS/RHEL:

```
sudo service postgresql-9.6 start
```

b. Ubuntu:

```
sudo service postgresql start
```

c. When the service restarts, you can check the cluster status using the following script:

```
./analyze_new_cluster.sh
```

4. After you have completed the database installation, you must review the port number of the newly installed database, which may have changed between versions of the database software. That new port number must be applied through the Trifacta software. For more information, see *Change Database Port*.

5. If all is well, restart the platform. See *Start and Stop the Platform* in the Install Guide.

Verify and Cleanup

1. Verify operations on all databases:

- Login to the application.
- Load a dataset from Flow View.
- Run a job.
- Schedule a job and execute it.
- See *Verify Operations* in the Admin Guide.

2. If all of the above tests pass, you can use the following script to delete the old PostgreSQL version and its data directory:

```
./delete_old_cluster.sh
```

3. Restart the Trifacta platform. See *Start and Stop the Platform* in the Install Guide.