



EDB Postgres™ Advanced Server Installation Guide for Linux

EDB Postgres™ Advanced Server 11

March 7, 2019

EDB Postgres™ Advanced Server Installation Guide for Linux
by EnterpriseDB® Corporation
Copyright © 2009 - 2019 EnterpriseDB Corporation. All rights reserved.

EnterpriseDB Corporation, 34 Crosby Drive Suite 100, Bedford, MA 01730, USA
T +1 781 357 3390 **F** +1 978 589 5701 **E** info@enterprisedb.com **www**.enterprisedb.com

Table of Contents

1	Introduction.....	4
1.1	Typographical Conventions Used in this Guide.....	5
2	Requirements Overview.....	6
2.1	Supported Platforms.....	6
2.2	Linux Installation Prerequisites.....	7
3	Limitations.....	8
4	Using a Package Manager to Install Advanced Server.....	9
4.1	Installing an RPM Package.....	11
4.1.1	Advanced Server RPMs Packages.....	13
4.1.2	Performing a Minor Version Update of an RPM Installation.....	16
4.2	Installing Advanced Server on a SLES Host.....	17
4.3	Installing Advanced Server on a Debian or Ubuntu Host.....	23
4.3.1	Advanced Server Debian Packages.....	24
4.4	Configuring a Package Installation.....	27
4.4.1	Creating a Database Cluster and Starting the Service.....	27
4.4.2	Using a Service Configuration File on CentOS or Redhat 6.x.....	29
4.4.2.1	Specifying Cluster Options with INITDBOPTS.....	30
4.4.3	Modifying the Data Directory Location on CentOS or Redhat 7.x.....	32
4.5	Starting Multiple Postmasters with Different Clusters.....	33
4.6	Creating an Advanced Server Repository on an Isolated Network.....	35
5	Installation Troubleshooting.....	37
6	Managing an Advanced Server Installation.....	38
6.1	Starting and Stopping Advanced Server and Supporting Components.....	38
6.1.1	Controlling a Service on CentOS or RHEL 7.x.....	39
6.1.2	Controlling a Service on CentOS or RHEL 6.x.....	40
6.1.3	Controlling a Service on Debian 9x or Ubuntu 18.04.....	41
6.1.4	Using pg_ctl to Control Advanced Server.....	42
6.1.5	Configuring Component Services to AutoStart at System Reboot.....	43
6.2	Modifying the postgresql.conf File.....	44
6.3	Modifying the pg_hba.conf File.....	46
6.4	Connecting to Advanced Server with psql.....	47
7	Uninstalling Advanced Server.....	48
7.1	Uninstalling an RPM Package.....	48
7.2	Uninstalling Advanced Server Components on a Debian or Ubuntu Host.....	49

1 Introduction

The EDB Postgres Advanced Server Installation Guide is a comprehensive guide to installing EDB Postgres Advanced Server (Advanced Server). In this guide you will find detailed information about:

- Software prerequisites for Advanced Server 11 installation on a Linux host.
- Using a package manager to install and update Advanced Server and its supporting components or utilities on a Linux host.
- Managing an Advanced Server installation.
- Configuring an Advanced Server package installation.
- Uninstalling Advanced Server and its components.

1.1 *Typographical Conventions Used in this Guide*

Certain typographical conventions are used in this manual to clarify the meaning and usage of various commands, statements, programs, examples, etc. This section provides a summary of these conventions.

In the following descriptions, a *term* refers to any word or group of words that are language keywords, user-supplied values, literals, etc. A term's exact meaning depends upon the context in which it is used.

- *Italic font* introduces a new term, typically in the sentence that defines it for the first time.
- Fixed-width (mono-spaced) font is used for terms that must be given literally such as SQL commands, specific table and column names used in the examples, programming language keywords, etc. For example, `SELECT * FROM emp;`
- *Italic fixed-width font* is used for terms for which the user must substitute values in actual usage. For example, `DELETE FROM table_name;`
- A vertical pipe | denotes a choice between the terms on either side of the pipe. A vertical pipe is used to separate two or more alternative terms within square brackets (optional choices) or braces (one mandatory choice).
- Square brackets [] denote that one or none of the enclosed terms may be substituted. For example, [a | b] means choose one of “a” or “b” or neither of the two.
- Braces { } denote that exactly one of the enclosed alternatives must be specified. For example, { a | b } means exactly one of “a” or “b” must be specified.
- Ellipses ... denote that the preceding term may be repeated. For example, [a | b] ... means that you may have the sequence, “b a a b a”.

2 Requirements Overview

The following sections detail the supported platforms and installation requirements for EDB Postgres Advanced Server 11.

2.1 Supported Platforms

The Advanced Server 11 RPM packages are supported on the following platforms:

64 bit Linux:

- Red Hat Enterprise Linux (x86_64) 6.x and 7.x
- CentOS (x86_64) 6.x and 7.x
- OEL Linux 6.x and 7.x
- PPC-LE 8 running RHEL or CentOS 7.x
- SLES 12
- Debian 9x
- Ubuntu 18.04
- IBM Power 9

2.2 Linux Installation Prerequisites

You can use an RPM package to install Advanced Server and supporting components on a Linux host. Before installing Advanced Server, please review the following prerequisites.

Installing EPEL

Before installing Advanced Server, you may be required to install the EPEL (Extra Packages for Enterprise Linux) repository. You can use `yum` to install the package:

```
yum install epel-release
```

If `yum` cannot access a repository that contains `epel-release`, you will get an error message:

```
No package epel available.  
Error: Nothing to do
```

If you receive this error, you can download the EPEL rpm package, and install it manually. To manually install EPEL, download the rpm package, assume superuser privileges, navigate into the directory that contains the package, and install EPEL with the command:

```
yum install epel-release
```

For more information about installing EPEL, visit:

https://fedoraproject.org/wiki/EPEL#How_can_I_use_these_extra_packages.3F

Linux-specific Software Requirements

You must install `xterm`, `konsole`, or `gnome-terminal` before executing any console-based program installed by EnterpriseDB installers.

Migration Toolkit or EDB*Plus Installation Prerequisites

Before using an RPM to install Migration Toolkit or EDB*Plus, you must first install Java version 1.7 or later for Migration Toolkit and Java version 1.8 or later for EDB*Plus. On a Linux system, you can use the `yum` package manager to install Java. Open a terminal window, assume superuser privileges, and enter:

```
# yum install java
```

Follow the onscreen instructions to complete the installation.

3 Limitations

The following limitation applies to the EDB Postgres Advanced Server:

- The LLVM JIT package is supported on RHEL 7.x and SLES. LLVM JIT is not supported on RHEL 6.x and PPC-LE 64 (running RHEL or CentOS 7.x).

4 Using a Package Manager to Install Advanced Server

You can use the yum package manager to install Advanced Server or Advanced Server supporting components. yum will attempt to satisfy package dependencies as it installs a package, but requires access to the Advanced Server repositories. If your system does not have access to a repository via the Internet, you can use RPM to install an individual package or create a local repository, but you may be required to manually satisfy package dependencies.

You must have credentials for the EnterpriseDB repository to install Advanced Server; to submit a request for credentials, click an Access Repository button on the following page:

<https://www.enterprisedb.com/advanced-downloads>

For a list of the RPM installers and Debian packages available from EnterpriseDB for Advanced Server and supporting components, see Section [4.1.1](#).

Installing the server package creates a database superuser named `enterprisedb`. The user is assigned a user ID (UID) and a group ID (GID) of 26. The user has no default password; use the `passwd` command to assign a password for the user. The default shell for the user is `bash`, and the user's home directory is `/var/lib/edb/as11`.

By default, Advanced Server logging is configured to write files to the `log` subdirectory of the `data` directory, rotating the files each day and retaining one week of log entries. You can customize the logging behavior of the server by modifying the `postgresql.conf` file; for more information about modifying the `postgresql.conf` file, please see Section [6.2](#).

The RPM installers place Advanced Server components in the directories listed in the table below:

EDBAS Component	Path to Installation Directory
Executables	<code>/usr/edb/as11/bin</code>
Libraries	<code>/usr/edb/as11/lib</code> or <code>/usr/edb/as11/lib64</code>
Documentation	<code>/usr/edb/as11/share/doc</code>
Contrib	<code>/usr/edb/as11/share/contrib</code>
Data	<code>/var/lib/edb/as11/data</code>
Backup area	<code>/var/lib/edb/as11/backups</code>
Templates	<code>/usr/edb/as11/share</code>
Procedural Languages	<code>/usr/edb/as11/lib</code> or <code>/usr/edb/as11/lib64</code>
Development Headers	<code>/usr/edb/as11/include</code>
Shared data	<code>/usr/edb/as11/share</code>
Regression tests	<code>/usr/edb/as11/lib/pgxs/src/test/regress</code>
SGML Documentation	<code>/usr/edb/as11/share/doc</code>

The Debian package place Advanced Server and supporting components in the directories listed in the following table:

EDB Debian Package Component	Path to Installation Directory
Server	/usr/lib/edb-as/11/
Data and Configuration Directory	/var/lib/edb-as/11/main /etc/edb-as/11/main/
pgAgent	/usr/lib/edb-as/11
Pgpool	/usr/edb/pgpool3.7/
Postgis	/usr/lib/edb-as/11/
PGSNMPD	/usr/lib/edb-as/11
Slony Replication	/usr/lib/edb-as/11
pgBouncer	/usr/edb/pgbouncer1.9/
pgBouncer Configuration Files	/etc/edb/pgbouncer1.9/pgbouncer.ini /etc/edb/pgbouncer1.9/userlist.txt
SQL-Profiler	/usr/lib/edb-as/11/lib
SQL-Protect	/usr/lib/edb-as/11/lib
SSLUTILS	/usr/lib/edb-as/11/lib
PL-PERL	/usr/lib/edb-as/11/lib
PL-PYTHON	/usr/lib/edb-as/11/lib
PLTCL	/usr/lib/edb-as/11/lib
EFM	/usr/edb/efm-3.2
JDBC	/usr/edb/jdbc
MTK	/usr/edb/migrationtoolkit/

4.1 Installing an RPM Package

Before installing Advanced Server or a supporting component via an RPM package over the web, you must create the repository configuration file (`edb-repo`). The repository configuration file contains connection and authentication information for the EnterpriseDB repository. To prepare your system to perform an RPM installation:

1. Assume superuser privileges and use either `rpm` or `yum` to create the repository configuration file:

```
rpm -Uvh https://yum.enterprisedb.com/edbrepos/edb-repo-latest.noarch.rpm
```

or

```
yum install -y https://yum.enterprisedb.com/edbrepos/edb-repo-latest.noarch.rpm
```

2. Use your choice of editor to modify the repository configuration file, enabling each repository from which you will install packages, and providing your credentials. The repository configuration file is named `edb.repo`; it resides in `/etc/yum.repos.d`.

To enable a repository, change the value of the `enabled` parameter to 1 and replace the user name and password placeholders in the `baseurl` specification with your user name and the repository password.

Installing Advanced Server

To install Advanced Server 11, enable and provide connection credentials for the `edbas11` repository and the `enterprisedb-dependencies` repository:

```
[edbas11]
name=EnterpriseDB Advanced Server 11 $releasever -
$basearch
baseurl=https://<username>:<password>@yum.enterprisedb.com/
11/redhat/rhel-$releasever-$basearch
enabled=0
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/ENTERPRISEDB-GPG-KEY
```

```
[enterprisedb-dependencies]
name=EnterpriseDB Dependencies $releasever - $basearch
baseurl=https://<username>:<password>@yum.enterprisedb.com/
dependencies/redhat/rhel-$releasever-$basearch
enabled=0
```

```

gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/ENTERPRISEDB-GPG-KEY

```

Installing supporting components

The repository configuration file also contains an entry for the `enterprisedb-tools` repository. Enable the `[enterprisedb-tools]` and the `[enterprisedb-dependencies]` entries in the `edb.repo` file when installing Advanced Server supporting components:

```

[enterprisedb-tools]
name=EnterpriseDB Tools $releasever - $basearch
baseurl=https://<username>:<password>@yum.enterprisedb.com/
tools/redhat/rhel-$releasever-$basearch
enabled=0
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/ENTERPRISEDB-GPG-KEY

```

3. After modifying applicable entries in the repository configuration file, save the configuration file and exit the editor.

Then, you can use the `yum install` command to install Advanced Server or supporting components. For example, to install the server and its core components, use the command:

```
yum install edb-as11-server
```

When you install an RPM package that is signed by a source that is not recognized by your system, `yum` may ask for your permission to import the key to your local server. If prompted, and you are satisfied that the packages come from a trustworthy source, enter a `y`, and press `Return` to continue.

After installing Advanced Server, you must configure the installation; see Section [4.4](#), *Configuring a Package Installation*, for details.

For information about the available packages, see Section [4.1.1](#).

During the installation, `yum` may encounter a dependency that it cannot resolve. If it does, it will provide a list of the required dependencies that you must manually resolve.

4.1.1 Advanced Server RPMs Packages

The tables that follow list the RPM packages that are available from EnterpriseDB.

You can also use the `yum search` command to access a list of the packages that are currently available from your configured repository. Open a command line, assume superuser privileges, and enter:

```
yum search package
```

Where *package* is the search term that specifies the name (or partial name) of a package. The repository search will return a list of available packages that include the specified search term.

Please note: The available package list is subject to change.

Package Name	Package Installs
<code>edb-as11-server</code>	This package installs core components of the Advanced Server database server.
<code>edb-as11-server-client</code>	The <code>edb-as11-server-client</code> package contains client programs and utilities that you can use to access and manage Advanced Server.
<code>edb-as11-server-contrib</code>	The <code>edb-as11-contrib</code> package installs contributed tools and utilities that are distributed with Advanced Server. Files for these modules are installed in: Documentation: <code>/usr/edb/as11/share/doc</code> Loadable modules: <code>/usr/edb/as11/lib</code> Binaries: <code>/usr/edb/as11/bin</code>
<code>edb-as11-server-core</code>	The <code>edb-as11-server-core</code> package includes the programs needed to create the core functionality behind the Advanced Server database.
<code>edb-as11-server-devel</code>	The <code>edb-as11-server-devel</code> package contains the header files and libraries needed to compile C or C++ applications that directly interact with an Advanced Server server and the <code>ecpg</code> or <code>ecpgPlus</code> C preprocessor.
<code>edb-as11-server-docs</code>	The <code>edb-as11-server-docs</code> package installs the readme file.
<code>edb-as11-server-edb-modules</code>	This package installs supporting modules for Advanced Server
<code>edb-as11-server-indexadvisor</code>	This package installs Advanced Server's Index Advisor feature. The Index Advisor utility helps determine which columns you should index to improve performance in a given workload.
<code>edb-as11-server-libs</code>	The <code>edb-as11-server-libs</code> package provides the essential shared libraries for any Advanced Server client program or interface.
<code>edb-as11-server-llvmjit</code>	This package contains support for Just in Time (JIT) compiling parts of EDBAS queries.
<code>edb-as11-server-pldebugger</code>	This package implements an API for debugging PL/pgSQL functions on Advanced Server.
<code>edb-as11-server-plperl</code>	The <code>edb-as11-server-plperl</code> package installs the PL/Perl procedural language for Advanced Server. Please note that the <code>edb-as11-server-plperl</code> package is dependent on the platform-supplied version of Perl.

EDB Postgres Advanced Server Installation Guide for Linux

Package Name	Package Installs
edb-as11-server-plpython	The <code>edb-as11-server-plpython</code> package installs the PL/Python procedural language for Advanced Server. Please note that the <code>edb-as11-server-plpython</code> package is dependent on the platform-supplied version of Python.
edb-as11-server-pltcl	The <code>edb-as11-pltcl</code> package installs the PL/Tcl procedural language for Advanced Server. Please note that the <code>edb-as11-pltcl</code> package is dependent on the platform-supplied version of TCL.
edb-as11-server-sqlprofiler	This package installs Advanced Server's SQL Profiler feature. SQL Profiler helps identify and optimize SQL code.
edb-as11-server-sqlprotect	This package installs Advanced Server's SQL Protect feature. SQL Protect provides protection against SQL injection attacks.
edb-as11-server-sslutils	This package installs functionality that provides SSL support.
edb-as11-server-cloneschema	This package installs the EDB Clone Schema extension. For more information about EDB Clone Schema, see the EDB Postgres Advanced Server Guide.
edb-as11-server-parallel-clone	This package installs functionality that supports the EDB Clone Schema extension.
edb-as11-edbplus	The <code>edb-edbplus</code> package contains the files required to install the EDB*Plus command line client. EDB*Plus commands are compatible with Oracle's SQL*Plus.
edb-as11-pgagent	This package installs pgAgent; pgAgent is a job scheduler for Advanced Server. Before installing this package, you must install EPEL; for detailed information about installing EPEL, see Section 2.2.
edb-as11-pgsnmpd	SNMP (Simple Network Management Protocol) is a protocol that allows you to supervise an apparatus connected to the network.
edb-as11-pljava	This package installs PL/Java, providing access to Java stored procedures, triggers and functions via the JDBC interface.
edb-as11-pgpool37-extensions	This package creates pgPool extensions required by the server.
edb-as11-postgis-2.5	This package installs POSTGIS meta RPMs.
edb-as11-slony-replication	This package installs the meta RPM for Slony-I.
edb-as11-slony-replication-core	Slony-I builds a master-slave system that includes all features and capabilities needed to replicate large databases to a reasonably limited number of slave systems.
edb-as11-slony-replication-docs	This package contains the Slony project documentation (in pdf form).
edb-as11-slony-replication-tools	This package contains the Slony <code>altperl</code> tools and utilities that are useful when deploying Slony replication environments. Before installing this package, you must install EPEL; for detailed information about installing EPEL, see Section 2.2.
edb-as11-hdfs_fdw	The Hadoop Data Adapter allows you to query and join data from Hadoop environments with your Postgres or Advanced Server instances. It is YARN Ready certified with HortonWorks, and provides optimizations for performance with predicate pushdown support.
edb-as11-mongo_fdw	This EnterpriseDB Advanced Server extension implements a Foreign Data Wrapper for MongoDB.
edb-as11-mysql8_fdw edb-as11-mysql5_fdw	This EnterpriseDB Advanced Server extension implements a Foreign Data Wrapper for MySQL.
libevent-edb libiconv-edb libcicu-edb	These packages contain supporting library files.

The following table lists the packages for Advanced Server 11 supporting components that are stored in the `Tools` repository:

Package Name	Package Installs
edb-pgpool37	This package contains the pgPool-II installer. pgPool provides connection pooling for Advanced Server installations.
edb-jdbc	The <code>edb-jdbc</code> package includes the <code>.jar</code> files needed for Java programs to access an Advanced Server database.
edb-migrationtoolkit	The <code>edb-migrationtoolkit</code> package installs Migration Toolkit, facilitating migration to an Advanced Server database from Oracle, PostgreSQL, MySQL, Sybase and SQL Server.
edb-oci	The <code>edb-oci</code> package installs the EnterpriseDB Open Client library, allowing applications that use the Oracle Call Interface API to connect to an Advanced Server database.
edb-oci-devel	This package installs the OCI include files; install this package if you are developing C/C++ applications that require these files.
edb-odbc	This package installs the driver needed for applications to access an Advanced Server system via ODBC.
edb-odbc-devel	This package installs the ODBC include files; install this package if you are developing C/C++ applications that require these files.
edb-pgbouncer19	This package contains PgBouncer (a lightweight connection pooler). This package requires the libevent package.
ppas-xdb	This package contains the xDB installer; xDB provides asynchronous cross-database replication. For more information, visit http://www.enterprisedb.com/faq-xdb-multi-master
ppas-xdb-console	This package provides support for xDB.
ppas-xdb-libs	This package provides support for xDB.
ppas-xdb-publisher	This package provides support for xDB.
ppas-xdb-subscriber	This package provides support for xDB.
edb-pem	The <code>edb-pem</code> package installs Management Tool that efficiently manages, monitor, and tune large Postgres deployments from a single remote GUI console.
edb-pem-agent	This package is an agent component of Postgres Enterprise Manager.
edb-pem-docs	This package contains documentation for various languages, which are in HTML format.
edb-pem-server	This package contains server components of Postgres Enterprise Manager.
edb-efm32 edb-efm31 edb-efm30	This package installs EDB Failover Manager that adds fault tolerance to database clusters to minimize downtime when a master database fails by keeping data online in high availability configurations.
edb-bart edb-bart20	This package installs the Backup and Recovery Tool (BART) to support online backup and recovery across local and remote PostgreSQL and EDB Advanced Servers.
libevent-edb libiconv-edb libevent-edb-devel	These packages contain supporting library files.

4.1.2 Performing a Minor Version Update of an RPM Installation

If you used an RPM package to install Advanced Server or its supporting components, you can use yum to perform a minor version upgrade to a more recent version. To review a list of the package updates that are available for your system, open a command line, assume root privileges, and enter the command:

```
yum check-update package_name
```

Where *package_name* is the search term for which you wish to search for updates. Please note that you can include wild-card values in the search term. To use yum update to install an updated package, use the command:

```
yum update package_name
```

Where *package_name* is the name of the package you wish to update. Include wild-card values in the update command to update multiple related packages with a single command. For example, use the following command to update all packages whose names include the expression `edb`:

```
yum update edb*
```

Please note that the yum update command will only perform an update between minor releases; to update between major releases, you must use `pg_upgrade`.

For more information about using yum commands and options, enter `yum --help` on your command line, or visit:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/ch-yum.html

4.2 Installing Advanced Server on a SLES Host

You must have credentials that allow access to the EnterpriseDB repository to install Advanced Server on a SLES host. To request credentials for the repository, visit the Advanced Downloads page at:

<https://www.enterprisedb.com/advanced-downloads>

You can use the zypper package manager to install Advanced Server or Advanced Server supporting components on an SLES host. zypper will attempt to satisfy package dependencies as it installs a package, but some components require access to specific repositories that are not hosted at EnterpriseDB.

To add Advanced Server or supporting components, use the following commands to add EnterpriseDB repository configuration files to your SLES host:

```
zypper addrepo https://zypp.enterprisedb.com/suse/epas11-
sles.repo
zypper addrepo https://zypp.enterprisedb.com/suse/epas-sles-
tools.repo
zypper addrepo https://zypp.enterprisedb.com/suse/epas-sles-
dependencies.repo
```

Each command creates a repository configuration file in the `/etc/zypp/repos.d` directory. The files are named:

- `edbas11suse.repo`
- `edbasdependencies.repo`
- `edbastools.repo`

After creating the repository configuration files, use the `zypper refresh` command to refresh the metadata on your SLES host to include the EnterpriseDB repositories.

```
/etc/zypp/repos.d # zypper refresh
Repository 'Devel project of Java packages for openSUSE:Factory
(SLE_12_SP3)' is up to date.
Repository 'SLE-Module-Legacy12-Pool' is up to date.
Repository 'SLE-Module-Legacy12-Updates' is up to date.
Repository 'SLES12-12-0' is up to date.
Repository 'SLES12-Pool' is up to date.
Repository 'SLES12-Updates' is up to date.
Repository 'SLE-SDK12-Pool' is up to date.
Repository 'SLE-SDK12-Updates' is up to date.
Repository 'SUSE-PackageHub-12' is up to date.
Repository 'SUSE-PackageHub-12-Pool' is up to date.
Repository 'wxWidgets Packages (SLE_12_SP3)' is up to date.
```

```
Repository 'Various compilers (SLE_12_SP3)' is up to date.  
Retrieving repository 'EDB Postgres Advanced Server 11 12 -  
x86_64' metadata -----[\  
]
```

```
Authentication required for  
'https://zypp.enterprisedb.com/11/suse/suse-12-x86_64'  
User Name:  
Password:
```

```
Retrieving repository 'EDB Postgres Advanced Server 11 12 -  
x86_64' metadata.....[done]  
Building repository 'EDB Postgres Advanced Server 11 12 - x86_64'  
cache.....[done]  
Retrieving repository 'EDB Postgres Advanced Server Dependencies  
12 - x86_64' metadata.....[done]  
Building repository 'EDB Postgres Advanced Server Dependencies 12  
- x86_64' cache.....[done]  
Retrieving repository 'EDB Postgres Advanced Server Tools 12 -  
x86_64' metadata.....[done]  
Building repository 'EDB Postgres Advanced Server Tools 12 -  
x86_64' cache.....[done]  
All repositories have been refreshed.
```

When prompted for a User Name and Password, provide your connection credentials for the EnterpriseDB repository. If you need credentials, contact EnterpriseDB at:

<https://www.enterprisedb.com/general-inquiry-form>

Before installing EDB Postgres Advanced Server or supporting components, you must also add SUSEConnect and the SUSE Package Hub extension to the SLES host and register the host with SUSE, allowing access to SUSE repositories. Use the commands:

```
zypper install SUSEConnect  
SUSEConnect -p PackageHub/12/x86_64  
SUSEConnect -p sle-sdk/12/x86_64
```

For detailed information about registering a SUSE host, visit:

<https://www.suse.com/support/kb/doc/?id=7016626>

The following sections provide information about the specific repositories that are not hosted by EnterpriseDB, but that provide software that satisfies dependencies for each listed component. Please note that after adding repository access, you should use the `zypper refresh` command to process the repository content:

```
zypper refresh
```

Java Components:

Please note that only OpenJDK (version 1.8) is supported on SLES hosts of Java components. Before using an RPM package to add Advanced Server or a supporting component that requires Java, use zypper to add supporting repository locations to your system.

Use the commands:

```
zypper addrepo
http://download.opensuse.org/repositories/Java:/Factory/SLE_12_SP
2/Java:Factory.repo
```

```
zypper refresh
```

Then, to install the server and its core components, invoke the command:

```
zypper install edb-as11
```

Slony:

If you are using a package to add Slony to a SLES host, please note that you must add access to the Perl repository:

For SLES 12 SP2:

```
zypper addrepo
http://download.opensuse.org/repositories/devel:languages:perl/SLE
E_12_SP2/devel:languages:perl.repo
```

For SLES 12 SP3:

```
zypper addrepo
http://download.opensuse.org/repositories/devel:languages:perl/SLE
E_12_SP3/devel:languages:perl.repo
```

Then:

```
zypper refresh
```

PostGIS:

When using an RPM package to add PostGIS to a SLES host, you must add the following repository:

```
zypper addrepo
"http://download.opensuse.org/repositories/server:/Kolab:/3.3/SLE
_12/server:Kolab:3.3.repo"
```

Then:

```
zypper refresh
```

pgAgent:

When adding the pgAgent packages to a SLES host, you must add access to the wxWidgets repository:

For SLES 12 SP2:

```
SUSEConnect -p sle-module-legacy/12/x86_64
SUSEConnect -p sle-sdk/12/x86_64
zypper addrepo
http://download.opensuse.org/repositories/X11:/wxWidgets/SLE_12_S
P2/X11:wxWidgets.repo
```

For SLES 12 SP3:

```
SUSEConnect -p sle-module-legacy/12/x86_64
SUSEConnect -p sle-sdk/12/x86_64
zypper addrepo
http://download.opensuse.org/repositories/X11:/wxWidgets/SLE_12_S
P3/X11:wxWidgets.repo
```

Then:

```
zypper refresh
```

PEM Server and PEM Agent:

When adding the PEM Server and PEM Agent to a SLES host, you must add access to the wxWidgets repository:

For SLES 12 SP2:

```
SUSEConnect -p sle-module-legacy/12/x86_64
SUSEConnect -p sle-sdk/12/x86_64
zypper addrepo
https://download.opensuse.org/repositories/Apache:Modules/SLE_12_
SP2/Apache:Modules.repo
zypper addrepo
http://download.opensuse.org/repositories/Cloud:/OpenStack:/Newto
n:/cisco-apic:/2.3.1/SLE_12_SP2/ pem_opensuse_boost
```

For SLES 12 SP3:

```
SUSEConnect -p sle-module-legacy/12/x86_64
```

```
SUSEConnect -p sle-sdk/12/x86_64
zypper addrepo
https://download.opensuse.org/repositories/Apache:Modules/SLE_12_
SP3/Apache:Modules.repo
zypper addrepo
http://download.opensuse.org/repositories/Cloud:/OpenStack:/Newto
n:/cisco-apic:/2.3.1/SLE_12_SP3/ pem_opensuse_boost
```

Then:

```
zypper refresh
zypper install edb-pem-server
```

The PEM server installer includes a script (`configure-pem-server.sh`) to help automate the configuration process for RPM installations. The script is installed in the `/usr/edb/pem/bin` directory. To invoke the script, use the command:

```
/usr/edb/pem/bin/configure-pem-server.sh
```

BART:

When installing BART on a SLES host, you must add the SUSE SDK as well as the Archiving repository:

```
SUSEConnect -p sle-sdk/12/x86_64
zypper addrepo
https://download.opensuse.org/repositories/Archiving/SLE_12_SP3/A
rchiving.repo
```

Then:

```
zypper refresh
```

LLVM JIT:

When installing LLVM JIT on a SLES host, you must add the following repository to your system:

```
zypper addrepo
https://download.opensuse.org/repositories/devel:/tools:/compiler
/SLE_12_SP3/devel:tools:compiler.repo
```

Then:

```
zypper refresh
```

Updating Components on a SLES Host:

To update components installed with zypper, use the `zypper update` command.

4.3 Installing Advanced Server on a Debian or Ubuntu Host

To install Advanced Server on a Debian or Ubuntu host, you must have credentials that allow access to the EnterpriseDB repository. To request credentials for the repository, visit the EnterpriseDB [Advanced Downloads](https://www.enterprisedb.com/advanced-downloads) page at:

<https://www.enterprisedb.com/advanced-downloads>

The following steps will walk you through using the EnterpriseDB apt repository to install a debian package. When using the commands, replace the `username` and `password` with the credentials provided by EnterpriseDB.

1. Assume superuser privileges:

```
sudo su -
```

2. Configure the EnterpriseDB repository:

```
sh -c 'echo "deb
https://username:password@apt.enterprisedb.com/$(lsb_release -
cs)-edb/ $(lsb_release -cs) main" >
/etc/apt/sources.list.d/edb-$(lsb_release -cs).list'
```

3. Add support to your system for secure APT repositories:

```
apt-get install apt-transport-https
```

4. Add the EBD signing key:

```
wget -q -O - https://username:password
@apt.enterprisedb.com/edb-deb.gpg.key | apt-key add -
```

5. Update the repository metadata:

```
apt-get update
```

6. Install Debian package:

```
apt-get install edb-as11
```

Note: The OpenJDK version 10 is supported on a Debian or Ubuntu host of Java components. Before using a debian package to add Advanced Server or a supporting component such as MTK or EDB*Plus that requires Java, make sure to install Java version 8 on your system to support EDB Java-based components.

4.3.1 Advanced Server Debian Packages

The table that follows lists some of the Debian packages that are available from EnterpriseDB.

You can also use the `apt list` command to access a list of the packages that are currently available from your configured repository. Open a command line, assume superuser privileges, and enter:

```
apt list edb*
```

Please note: the available package list is subject to change.

Package Name	Package Installs
<code>edb-as11-server</code>	This package installs core components of the Advanced Server database server.
<code>edb-as11-server-client</code>	The <code>edb-as11-server-client</code> package contains client programs and utilities that you can use to access and manage Advanced Server.
<code>edb-as11-server-core</code>	The <code>edb-as11-server-core</code> package includes the programs needed to create the core functionality behind the Advanced Server database.
<code>edb-as11-server-dev</code>	The <code>edb-as11-server-dev</code> package contains the header files and libraries needed to compile C or C++ applications that directly interact with an Advanced Server server and the <code>ecpg</code> or <code>ecpgPlus</code> C preprocessor.
<code>edb-as11-server-doc</code>	The <code>edb-as11-server-docs</code> package installs the readme file.
<code>edb-as11-server-edb-modules</code>	This package installs supporting modules for Advanced Server
<code>edb-as11-server-indexadvisor</code>	This package installs Advanced Server's Index Advisor feature. The Index Advisor utility helps determine which columns you should index to improve performance in a given workload.
<code>edb-as11-server-pldebugger</code>	This package implements an API for debugging PL/pgSQL functions on Advanced Server.
<code>edb-as11-server-plpython</code>	The <code>edb-as11-server-plpython</code> package installs the PL/Python procedural language for Advanced Server. Please note that the <code>edb-as11-server-plpython</code> package is dependent on the platform-supplied version of Python.
<code>edb-as11-server-pltcl</code>	The <code>edb-as11-pltcl</code> package installs the PL/Tcl procedural language for Advanced Server. Please note that the <code>edb-as11-pltcl</code> package is dependent on the platform-supplied version of TCL.
<code>edb-as11-server-sqlprofiler</code>	This package installs Advanced Server's SQL Profiler feature. SQL Profiler helps identify and optimize SQL code.
<code>edb-as11-server-sqlprotect</code>	This package installs Advanced Server's SQL Protect feature. SQL Protect provides protection against SQL injection attacks.
<code>edb-as11-server-sslutils</code>	This package installs functionality that provides SSL support.
<code>edb-as11-server-cloneschema</code>	This package installs the EDB Clone Schema extension. For more information about EDB Clone Schema, see the EDB Postgres Advanced Server Guide.
<code>edb-as11-server-parallel-clone</code>	This package installs functionality that supports the EDB Clone Schema extension.
<code>edb-as11-edbplus</code>	The <code>edb-edbplus</code> package contains the files required to install the

EDB Postgres Advanced Server Installation Guide for Linux

Package Name	Package Installs
	EDB*Plus command line client. EDB*Plus commands are compatible with Oracle's SQL*Plus.
edb-as11-pgsnmpd	SNMP (Simple Network Management Protocol) is a protocol that allows you to supervise an apparatus connected to the network.
edb-as11-pljava	This package installs PL/Java, providing access to Java stored procedures, triggers and functions via the JDBC interface.
edb-as11-pgadmin4	pgAdmin 4 provides a graphical management interface for Advanced Server and PostgreSQL databases.
edb-as11-pgadmin-apache	Apache support module for pgAdmin 4.
edb-as11-pgadmin4-common	pgAdmin 4 supporting files.
edb-as11-pgadmin4-doc	pgAdmin 4 documentation module.
edb-as11-pgpool37-extensions	This package creates pgPool extensions required by the server.
edb-as11-postgis-2.5	This package installs POSTGIS support for geospatial data.
edb-as11-postgis-2.5-scripts	This package installs POSTGIS support for geospatial data.
edb-as11-postgis-doc-2.5	This package provides support for POSTGIS.
edb-as11-postgis-gui-2.5	This package provides support for POSTGIS.
edb-as11-postgis-jdbc	This package provides support for POSTGIS.
edb-as11-postgis-scripts	This package provides support for POSTGIS.
edb-as11-pgagent	This package installs pgAgent; pgAgent is a job scheduler for Advanced Server. Before installing this package, you must install EPEL; for detailed information about installing EPEL, see Section 2.2 .
edb-as11-slony-replication	This package installs the meta RPM for Slony-I.
edb-as11-slony-replication-core	This package contains core portions of Slony-I to build a master-slave system that includes all features and capabilities needed to replicate large databases to a reasonably limited number of slave systems.
edb-as11-slony-replication-docs	This package contains the Slony project documentation (in pdf form).
edb-as11-slony-replication-tools	This package contains the Slony altperl tools and utilities that are useful when deploying Slony replication environments. Before installing this package, you must install EPEL; for detailed information about installing EPEL, see Section 2.2 .
edb-as11-hdfs-fdw	The Hadoop Data Adapter allows you to query and join data from Hadoop environments with your Postgres or Advanced Server instances. It is YARN Ready certified with HortonWorks, and provides optimizations for performance with predicate pushdown support.
edb-as11-hdfs-fdw-doc	Documentation for the Hadoop Data Adapter.
edb-as11-mongo-fdw	This EnterpriseDB Advanced Server extension implements a Foreign Data Wrapper for MongoDB.
edb-as11-mongo-fdw-doc	Documentation for the Foreign Data Wrapper for MongoDB.
edb-as11-mysql-fdw	This EnterpriseDB Advanced Server extension implements a Foreign Data Wrapper for MySQL.
edb-pgpool37	This package contains the pgPool-II installer. pgPool provides connection pooling for Advanced Server installations.
edb-jdbc	The edb-jdbc package includes the .jar files needed for Java programs to access an Advanced Server database.
edb-migrationtoolkit	The edb-migrationtoolkit package installs Migration Toolkit, facilitating migration to an Advanced Server database from Oracle, PostgreSQL, MySQL, Sybase and SQL Server.
edb-pgbouncer19	This package contains PgBouncer (a lightweight connection

EDB Postgres Advanced Server Installation Guide for Linux

Package Name	Package Installs
	pooler). This package requires the libevent package.
edb-efm34 edb-efm33 edb-efm32 edb-efm31	This package installs EDB Failover Manager that adds fault tolerance to database clusters to minimize downtime when a master database fails by keeping data online in high availability configurations.

4.4 Configuring a Package Installation

The packages that install the database server component create a service configuration file (on version 6.x hosts) or unit file (on version 7.x hosts), and service startup scripts.

4.4.1 Creating a Database Cluster and Starting the Service

The PostgreSQL `initdb` command creates a database cluster. After installing Advanced Server, you must manually configure the service and invoke `initdb` to create your cluster. When invoking `initdb`, you can:

- Specify environment options on the command line.
- Include the `service` command on RHEL or CentOS 6.x, and use a service configuration file to configure the environment.
- Include the `systemd` service manager on RHEL or CentOS 7.x use a service configuration file to configure the environment.

To review the `initdb` documentation, visit:

<https://www.postgresql.org/docs/11/static/app-initdb.html>

After specifying any options in the service configuration file, you can create the database cluster and start the service; these steps are platform specific.

On RHEL or CentOS 6.x

To create a database cluster in the `PGDATA` directory that listens on the port specified by the `PGPORT` environment variable specified in the service configuration file (described in Section 4.4.2), assume `root` privileges, and invoke the `service` script:

```
service edb-as-11 initdb
```

You can also assign a locale to the cluster when invoking `initdb`. By default, `initdb` will use the value specified by the `$LANG` operating system variable, but if you append a preferred locale when invoking the script, the cluster will use the alternate value. For example, to create a database cluster that uses simplified Chinese, invoke the command:

```
service edb-as-11 initdb zh_CH.UTF-8
```

After creating a database cluster, start the database server with the command:

```
service edb-as-11 start
```

On RHEL or CentOS 7.x

To invoke `initdb` on a RHEL or CentOS 7.x system, with the options specified in the service configuration file, assume the identity of the operating system superuser:

```
su - root
```

To initialize a cluster with the non-default values, you can use the `PGSETUP_INITDB_OPTIONS` environment variable by invoking the `edb-as-11-setup` cluster initialization script that resides under `EPAS_Home/bin`.

To invoke `initdb` export the `PGSETUP_INITDB_OPTIONS` environment variable with the following command:

```
PGSETUP_INITDB_OPTIONS="-E UTF-8" /usr/edb/as11/bin/edb-as-11-setup initdb
```

After creating the cluster, use `systemctl` to start, stop, or restart the service:

```
systemctl { start | stop | restart } edb-as-11
```

On Debian 9x or Ubuntu 18.04

You can initialize multiple clusters using the bundled scripts. To create a new cluster, assume `root` privileges, and invoke the bundled script:

```
/usr/bin/epas_createcluster 11 main2
```

To start a new cluster, use the following command:

```
/usr/bin/epas_ctlcluster 11 main2 start
```

To list all the available clusters, use the following command:

```
/usr/bin/epas_lsclusters
```

Note: The data directory is created under `/var/lib/edb-as/11/main2` and configuration directory is created under `/etc/edb-as/11/main/`.

4.4.2 Using a Service Configuration File on CentOS or Redhat 6.x

On a CentOS or RedHat version 6.x host, the RPM installer creates a service configuration file named `edb-as-11.sysconfig` in `/etc/sysconfig/edb/as11` (see Figure 4.1). Please note that options specified in the service configuration file are only enforced if `initdb` is invoked via the `service` command; if you manually invoke `initdb` (at the command line), you must specify the other options (such as the location of the data directory and installation mode) on the command line.

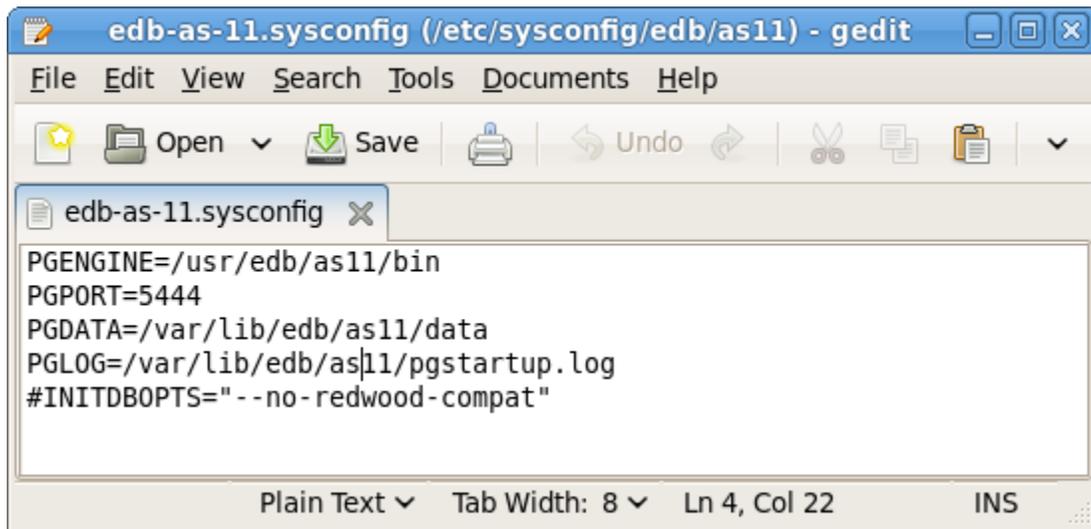


Figure 4.1 -The Advanced Server service configuration file.

The file contains the following environment variables:

- `PGENGINE` specifies the location of the engine and utility executable files.
- `PGPORT` specifies the listener port for the database server.
- `PGDATA` specifies the path to the data directory.
- `PGLOG` specifies the location of the log file to which the server writes startup information.
- Use `INITDBOPTS` to specify any `initdb` option or options that you wish to apply to the new cluster. For more information, see Section [4.4.2.1](#).

You can modify the `edb-as-11.sysconfig` file before using the `service` command to invoke the `startup` script to change the listener port, data directory location, startup log location or installation mode. If you plan to create more than one instance on the same system, you may wish to copy the `edb-as-11.sysconfig` file (and the associated `edb-as-11` startup script) and modify the file contents for each additional instance that resides on the same host.

4.4.2.1 Specifying Cluster Options with INITDBOPTS

You can use the `INITDBOPTS` variable to specify your cluster configuration preferences. By default, the `INITDBOPTS` variable is commented out in the service configuration file; unless modified, when you run the service startup script, the new cluster will be created in a mode compatible with Oracle databases. Clusters created in this mode will contain a database named `edb`, and have a database superuser named `enterprisedb`.

To create a new cluster in PostgreSQL mode, remove the pound sign (`#`) in front of the `INITDBOPTS` variable, enabling the `--no-redwood-compat` option. Clusters created in PostgreSQL mode will contain a database named `postgres`, and have a database superuser named `postgres`.

You may also specify multiple `initdb` options. For example, the following statement:

```
INITDBOPTS="--no-redwood-compat -U alice --locale=en_US.UTF-8"
```

Creates a database cluster (without compatibility features for Oracle) that contains a database named `postgres` that is owned by a user named `alice`; the cluster uses UTF-8 encoding.

In addition to the cluster configuration options documented in the PostgreSQL core documentation, Advanced Server supports the following `initdb` options:

```
--no-redwood-compat
```

Include the `--no-redwood-compat` keywords to instruct the server to create the cluster in PostgreSQL mode. When the cluster is created in PostgreSQL mode, the name of the database superuser will be `postgres`, the name of the default database will be `postgres`, and Advanced Server's features compatible with Oracle databases will not be available to the cluster.

```
--redwood-like
```

Include the `--redwood-like` keywords to instruct the server to use an escape character (an empty string (`' '`)) following the `LIKE` (or PostgreSQL-compatible `ILIKE`) operator in a SQL statement that is compatible with Oracle syntax.

```
--icu-short-form
```

Include the `--icu-short-form` keywords to create a cluster that uses a default ICU (International Components for Unicode) collation for all databases in the cluster. For more information about Unicode collations, please refer to the *EDB Postgres Advanced Server Guide* available at:

<http://www.enterprisedb.com/products-services-training/products/documentation>

For more information about using `initdb`, and the available cluster configuration options, see the PostgreSQL Core Documentation available at:

<https://www.postgresql.org/docs/11/static/app-initdb.html>

You can also view online help for `initdb` by assuming superuser privileges and entering:

```
/path_to_initdb_installation_directory/initdb --help
```

Where *path_to_initdb_installation_directory* specifies the location of the `initdb` binary file.

4.4.3 Modifying the Data Directory Location on CentOS or Redhat 7.x

On a CentOS or RedHat version 7.x host, the unit file is named `edb-as-11.service` and resides in `/usr/lib/systemd/system`. The unit file contains references to the location of the Advanced Server data directory. You should avoid making any modifications directly to the unit file because it may be overwritten during package upgrades.

By default, data files reside under `/var/lib/edb/as11/data` directory. To use a data directory that resides in a non-default location, create a copy of the unit file under the `/etc` directory:

```
cp /usr/lib/systemd/system/edb-as-11.service
   /etc/systemd/system/
```

After copying the unit file to the new location, modify the service file (`/etc/systemd/system/edb-as-11.service`) with your editor of choice, correcting any required paths.

Then, use the following command to reload `systemd`, updating the modified service scripts:

```
systemctl daemon-reload
```

Then, start the Advanced Server service with the following command:

```
systemctl start edb-as-11
```

4.5 Starting Multiple Postmasters with Different Clusters

You can configure Advanced Server to use multiple postmasters, each with its own database cluster. The steps required are version specific to the Linux host.

On RHEL or CentOS 6.x

The `edb-as11-server-core` RPM contains a script that starts the Advanced Server instance. The script can be copied, allowing you to run multiple services, with unique data directories and that monitor different ports. You must have `root` access to invoke or modify the script.

The example that follows creates a second instance on an Advanced Server host; the secondary instance is named `secondary`:

1. Create a hard link in `/etc/rc.d/init.d` (or equivalent location) to the `edb-as-11` service (named `secondary-edb-as-11`):

```
ln edb-as-11 secondary-edb-as-11
```

Be sure to pick a name that is not already used in `/etc/rc.d/init.d`.

2. Create a file in `/etc/sysconfig/edb/as11/` named `secondary-edb-as-11`. This file is where you would typically define `PGDATA` and `PGOPTS`. Since `$PGDATA/postgresql.conf` will override many of these settings (except `PGDATA`) you might notice unexpected results on startup.
3. Create the target `PGDATA` directory.
4. Assume the identity of the Advanced Server database superuser (`enterprisedb`) and invoke `initdb` on the target `PGDATA`. For information about using `initdb`, please see the PostgreSQL Core Documentation available at:

<https://www.postgresql.org/docs/11/static/app-initdb.html>
5. Edit the `postgresql.conf` file to specify the port, address, TCP/IP settings, etc. for the `secondary` instance.
6. Start the postmaster with the following command:

```
service secondary-edb-as-11 start
```

On RHEL or CentOS 7.x

The `edb-as11-server-core` RPM for version 7.x contains a unit file that starts the Advanced Server instance. The file allows you to start multiple services, with unique data directories and that monitor different ports. You must have `root` access to invoke or modify the script.

The example that follows creates an Advanced Server installation with two instances; the secondary instance is named `secondary`:

1. Make a copy of the default file with the new name. As noted at the top of the file, all modifications must reside under `/etc`. You must pick a name that is not already used in `/etc/systemd/system`.

```
cp /usr/lib/systemd/system/edb-as-11.service
   /etc/systemd/system/secondary-edb-as-11.service
```

2. Edit the file, changing `PGDATA` to point to the new data directory that you will create the cluster against.
3. Create the target `PGDATA` with user `enterprisedb`.
4. Run `initdb`, specifying the setup script:

```
/usr/edb/as11/bin/edb-as-11-setup initdb secondary-edb-as-11
```

5. Edit the `postgresql.conf` file for the new instance, specifying the port, the IP address, TCP/IP settings, etc.
6. Make sure that new cluster runs after a reboot:

```
systemctl enable secondary-edb-as-11
```

7. Start the second cluster with the following command:

```
systemctl start secondary-edb-as-11
```

4.6 *Creating an Advanced Server Repository on an Isolated Network*

You can create a local repository to act as a host for the Advanced Server RPM packages if the server on which you wish to install Advanced Server (or supporting components) cannot directly access the EnterpriseDB repository. Please note that this is a high-level listing of the steps requires; you will need to modify the process for your individual network.

To create and use a local repository, you must:

1. Use yum to install the `epel-release`, `yum-utils`, and `createrepo` packages:

```
yum install epel-release
yum install yum-utils
yum install createrepo
```

2. Create a directory in which to store the repository:

```
mkdir /srv/repos
```

3. Copy the RPM installation packages to your local repository. You can download the individual packages or use a tarball to populate the repository.
4. Sync the RPM packages and create the repository.

```
reposync -r edbas11 -p /srv/repos
createrepo /srv/repos
```

5. Install your preferred webserver on the host that will act as your local repository, and ensure that the repository directory is accessible to the other servers on your network.
6. On each isolated database server, configure yum to pull updates from the mirrored repository on your local network. For example, you might create a repository configuration file called `/etc/yum.repos.d/edb-repo` with connection information that specifies:

```
[edbas11]
name=EnterpriseDB Advanced Server 11
baseurl=https://yum.your_domain.com/edbas11
enabled=1
gpgcheck=0
```

After specifying the location and connection information for your local repository, you can use yum commands to install Advanced Server and its supporting components on the isolated servers. For example:

```
yum install edb-as11-server
```

For more information about creating a local yum repository, visit:

<https://wiki.centos.org/HowTos/CreateLocalRepos>

5 Installation Troubleshooting

Difficulty Displaying Java-based Applications

If you encounter difficulty displaying Java-based server features (controls or text not being displayed correctly, or blank windows), upgrading to the latest `libxcb-xlib` libraries should correct the problem on most distributions. Please visit the following link for other possible work-arounds:

http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=6532373

The Installation Fails to Complete Due to Existing data Directory Contents

If an installation fails to complete due to existing content in the data directory, the server will write an error message to the server logs:

```
A data directory is neither empty, or a recognisable data
directory.
```

If you encounter a similar message, you should confirm that the data directory is empty; the presence of files (including the system-generated `lost+found` folder) will prevent the installation from completing. Either remove the files from the data directory, or specify a different location for the data directory before re-invoking the installer to complete the installation.

6 Managing an Advanced Server Installation

Unless otherwise noted, the commands and paths noted in the following section assume that you have performed an installation using the native packages.

6.1 Starting and Stopping Advanced Server and Supporting Components

A service is a program that runs in the background and requires no user interaction (in fact, a service provides no user interface); a service can be configured to start at boot time, or manually on demand. Services are best controlled using the platform-specific operating system service control utility. Many of the Advanced Server supporting components are services.

The following table lists the names of the services that control Advanced Server and services that control Advanced Server supporting components:

Advanced Server Component Name	Linux Service Name	Debian Service Name
Advanced Server	edb-as-11	edb-as11-main
pgAgent	edb-pgagent-11	edb-as11-pgagent
PgBouncer	edb-pgbouncer-19	edb-pgbouncer19
pgPool-II	edb-pgpool-37	edb-pgpool37
Slony	edb-slony-replication-11	edb-as11-slony-replication
EFM	efm-3.2	efm-3.2

You can use the Linux command line to control Advanced Server's database server and the services of Advanced Server's supporting components. The commands that control the Advanced Server service on a Linux platform are host specific.

6.1.1 Controlling a Service on CentOS or RHEL 7.x

If your installation of Advanced Server resides on version 7.x of RHEL and CentOS, you must use the `systemctl` command to control the Advanced Server service and supporting components.

The `systemctl` command must be in your search path and must be invoked with superuser privileges. To use the command, open a command line, and enter:

```
systemctl action service_name
```

Where:

action

action specifies the action taken by the service command. Specify:

- `start` to start the service.
- `stop` to stop the service.
- `restart` to stop and then start the service.
- `status` to discover the current status of the service.

service_name

service_name specifies the name of the service.

6.1.2 Controlling a Service on CentOS or RHEL 6.x

On version 6.x of RHEL or CentOS Linux, you can control a service at the command line with the `service` command. The `service` command can be used to manage an Advanced Server cluster, as well as the services of component software installed with Advanced Server.

Using the `service` command to change the status of a service allows the Linux service controller to keep track of the server status (the `pg_ctl` command does not alert the service controller to changes in the status of a server). The command must be in your search path and must be invoked with superuser privileges. Open a command line, and issue the command:

```
service service_name action
```

The Linux `service` command invokes a script (with the same name as the service) that resides in `/etc/init.d`. If your Linux distribution does not support the `service` command, you can call the script directly by entering:

```
/etc/init.d/service_name action
```

Where:

`service_name`

`service_name` specifies the name of the service.

`action`

`action` specifies the action taken by the service command. Specify:

- `start` to start the service.
- `stop` to stop the service.
- `condstop` to stop the service without displaying a notice if the server is already stopped.
- `restart` to stop and then start the service.
- `condrestart` to restart the service without displaying a notice if the server is already stopped.
- `try-restart` to restart the service without displaying a notice if the server is already stopped.
- `status` to discover the current status of the service.

6.1.3 Controlling a Service on Debain 9x or Ubuntu 18.04

By default, a server is running with the peer or md5 permission on a Debian or Ubuntu host. You can change the `auth_type` from `pg_hba.conf` file available under:

```
vi /etc/edb-as/11/main/pg_hba.conf
```

To know more about modifying the `pg_hba.conf` file, see Section [6.3](#).

Please note, all the server configuration files are available under `/etc/edb-as/11/main`.

If your installation of Advanced Server resides on version 9x of Debian or 18.04 of Ubuntu, assume superuser privileges and invoke the following command (using bundled scripts) to:

1. Discover the current status of a service:

```
/usr/bin/epas_ctlcluster 11 main status
```

2. Stop a service:

```
/usr/bin/epas_ctlcluster 11 main stop
```

3. Restart a service:

```
/usr/bin/epas_ctlcluster 11 main restart
```

4. Reload a service:

```
/usr/bin/epas_ctlcluster 11 main reload
```

5. Control the component services:

```
systemctl restart edb-as@11-main
```

6.1.4 Using `pg_ctl` to Control Advanced Server

You can use the `pg_ctl` utility to control an Advanced Server service from the command line on any platform. `pg_ctl` allows you to start, stop, or restart the Advanced Server database server, reload the configuration parameters, or display the status of a running server. To invoke the utility, assume the identity of the cluster owner, navigate into the home directory of Advanced Server, and issue the command:

```
./bin/pg_ctl -D data_directory action
```

data_directory

data_directory is the location of the data controlled by the Advanced Server cluster.

action

action specifies the action taken by the `pg_ctl` utility. Specify:

- `start` to start the service.
- `stop` to stop the service.
- `restart` to stop and then start the service.
- `reload` sends the server a `SIGHUP` signal, reloading configuration parameters
- `status` to discover the current status of the service.

For more information about using the `pg_ctl` utility, or the command line options available, please see the official PostgreSQL Core Documentation available at:

<https://www.postgresql.org/docs/11/static/app-pg-ctl.html>

Choosing Between `pg_ctl` and the `service` Command

You can use the `pg_ctl` utility to manage the status of an Advanced Server cluster, but it is important to note that `pg_ctl` does not alert the operating system service controller to changes in the status of a server, so it is beneficial to use the `service` command whenever possible.

6.1.5 Configuring Component Services to AutoStart at System Reboot

After installing, configuring, and starting the services of Advanced Server supporting components on a Linux system, you must manually configure your system to autostart the service when your system reboots. To configure a service to autostart on a Linux system, open a command line, assume superuser privileges, and enter the following command.

On a Redhat-compatible Linux system, enter:

```
/sbin/chkconfig service_name on
```

Where *service_name* specifies the name of the service.

6.2 Modifying the `postgresql.conf` File

Configuration parameters in the `postgresql.conf` file specify server behavior with regards to auditing, authentication, encryption, and other behaviors. The `postgresql.conf` file resides in the data directory under your Advanced Server installation.

```
# PostgreSQL Client Authentication Configuration File
# =====
#
# Refer to the "Client Authentication" section in the PostgreSQL
# documentation for a complete description of this file. A short
# synopsis follows.
#
# This file controls: which hosts are allowed to connect, how clients
# are authenticated, which PostgreSQL user names they can use, which
# databases they can access. Records take one of these forms:
#
# local      DATABASE USER METHOD [OPTIONS]
# host       DATABASE USER ADDRESS METHOD [OPTIONS]
# hostssl    DATABASE USER ADDRESS METHOD [OPTIONS]
# hostnossl  DATABASE USER ADDRESS METHOD [OPTIONS]
#
# (The uppercase items must be replaced by actual values.)
#
# The first field is the connection type: "local" is a Unix-domain
# socket, "host" is either a plain or SSL-encrypted TCP/IP socket,
# "hostssl" is an SSL-encrypted TCP/IP socket, and "hostnossl" is a
# plain TCP/IP socket.
#
# DATABASE can be "all", "sameuser", "samerole", "replication", a
# database name, or a comma-separated list thereof. The "all"
# keyword does not match "replication". Access to replication
# must be enabled in a separate record (see example below).
#
# USER can be "all", a user name, a group name prefixed with "+", or a
# comma-separated list thereof. In both the DATABASE and USER fields
# you can also write a file name prefixed with "@" to include names
# from a separate file.
#
# ADDRESS specifies the set of hosts the record matches. It can be a
# host name, or it is made up of an IP address and a CIDR mask that is
# an integer (between 0 and 32 (IPv4) or 128 (IPv6) inclusive) that
# specifies the number of significant bits in the mask. A host name
# that starts with a dot (.) matches a suffix of the actual host name.
# Alternatively, you can write an IP address and netmask in separate
# columns to specify the set of hosts. Instead of a CIDR-address, you
# can write "samehost" to match any of the server's own IP addresses,
# or "samenet" to match any address in any subnet that the server is
# directly connected to.
#
# METHOD can be "trust", "reject", "md5", "password", "scram-sha-256",
# "gss", "sspi", "ident", "peer", "pam", "ldap", "radius" or "cert".
# Note that "password" sends passwords in clear text; "md5" or
# "scram-sha-256" are preferred since they send encrypted passwords.
#
# OPTIONS are a set of options for the authentication in the format
```

Figure 6.1 - The `postgresql.conf` file.

Parameters that are preceded by a pound sign (#) are set to their default value (as shown in the parameter setting). To change a parameter value, remove the pound sign and enter a new value. After setting or changing a parameter, you must either *reload* or *restart* the server for the new parameter value to take effect.

Within the `postgresql.conf` file, some parameters contain comments that indicate change requires restart (see Figure 6.1). To view a list of the parameters that require a server restart, execute the following query at the `psql` command line:

```
SELECT name FROM pg_settings WHERE context = 'postmaster';
```

6.3 Modifying the `pg_hba.conf` File

Appropriate authentication methods provide protection and security. Entries in the `pg_hba.conf` file specify the authentication method or methods that the server will use when authenticating connecting clients. Before connecting to the server, you may be required to modify the authentication properties specified in the `pg_hba.conf` file.

When you invoke the `initdb` utility to create a cluster, `initdb` creates a `pg_hba.conf` file for that cluster that specifies the type of authentication required from connecting clients.

The default authentication configuration specified in the `pg_hba.conf` file is:

```
# TYPE  DATABASE        USER            ADDRESS             METHOD
# "local" is for Unix domain socket connections only
local   all             all             peer
# IPv4 local connections:
host    all             all             127.0.0.1/32       ident
# IPv6 local connections:
host    all             all             ::1/128             ident
# Allow replication connections from localhost, by a user with the
# replication privilege.
local   replication    all             peer
host    replication    all             127.0.0.1/32       ident
host    replication    all             ::1/128             ident
```

To modify the `pg_hba.conf` file, open the file with your choice of editor. After modifying the authentication settings in the `pg_hba.conf` file, use the Linux command line to restart the server and apply the changes.

For more information about authentication, and modifying the `pg_hba.conf` file, see the PostgreSQL Core Documentation at:

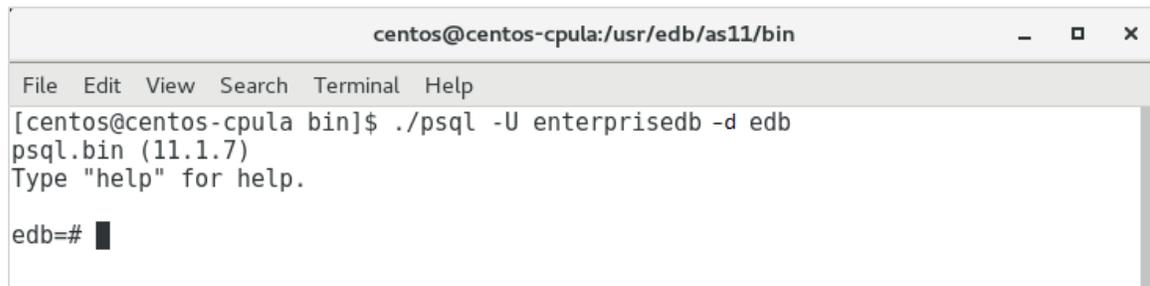
<https://www.postgresql.org/docs/11/static/auth-pg-hba-conf.html>

6.4 Connecting to Advanced Server with psql

`psql` is a command line client application that allows you to execute SQL commands and view the results. To open the `psql` client, the client must be in your search path. The executable resides in the `bin` directory, under your Advanced Server installation.

Use the following command and options to start the `psql` client (see Figure 6.2):

```
psql -d edb -U enterprisedb
```

A screenshot of a terminal window titled "centos@centos-cpula:/usr/edb/as11/bin". The terminal shows the command `./psql -U enterprisedb -d edb` being executed. The output is `psql.bin (11.1.7)` and `Type "help" for help.` The prompt `edb=#` is visible at the bottom of the terminal.

```
centos@centos-cpula:/usr/edb/as11/bin
File Edit View Search Terminal Help
[centos@centos-cpula bin]$ ./psql -U enterprisedb -d edb
psql.bin (11.1.7)
Type "help" for help.
edb=#
```

Figure 6.2 - Connecting to the server.

Where:

- d specifies the database to which `psql` will connect;
- U specifies the identity of the database user that will be used for the session.

For more information about using the command line client, please refer to the PostgreSQL Core Documentation at:

<https://www.postgresql.org/docs/11/static/app-psql.html>

7 Uninstalling Advanced Server

Note that after uninstalling Advanced Server, the cluster data files remain intact and the service user persists. You may manually remove the cluster data and service user from the system.

7.1 Uninstalling an RPM Package

You can use variations of the `rpm` or `yum` command to remove installed packages. Note that removing a package does not damage the Advanced Server `data` directory.

Include the `-e` option when invoking the `rpm` command to remove an installed package; the command syntax is:

```
rpm -e package_name
```

Where `package_name` is the name of the package that you would like to remove.

You can use the `yum remove` command to remove a package installed by `yum`. To remove a package, open a terminal window, assume superuser privileges, and enter the command:

```
yum remove package_name
```

Where `package_name` is the name of the package that you would like to remove.

Note: `yum` and `RPM` will not remove a package that is required by another package. If you attempt to remove a package that satisfies a package dependency, `yum` or `RPM` will provide a warning.

7.2 Uninstalling Advanced Server Components on a Debian or Ubuntu Host

1. To uninstall Advanced Server, invoke the following command:

```
apt-get remove edb-as11-server*
```

Please note: the configuration files and data directory remains intact.

2. To uninstall Advanced Server, configuration files, and data directory, invoke the following command:

```
apt-get purge edb-as11-server*
```