



Connecting to an EDB Ark™ Cluster

Version 2.0

September 13, 2016

**Connecting to an EDB Ark Cluster
by EnterpriseDB® Corporation
Copyright © 2016 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	Connection Pre-Requisites.....	6
2.1	Selecting a Port for a Client Connection.....	6
2.2	Downloading your Private Key	7
3	Using ssh to Access a Server	8
4	Copying a file to an EDB Ark Cluster	9
5	Connecting with the Postgres Enterprise Manager Client	10
6	Connecting to EDB Ark with psql or edb-psql.....	13

1 Introduction

EDB Ark automatically provisions PostgreSQL or EDB Postgres Advanced Server (Advanced Server) databases in single instances, high-availability clusters, or application development sandboxes in an OpenStack environment. This tutorial will walk you through the process of using some of the client applications that are distributed with Advanced Server and PostgreSQL to connect to a node that resides in an OpenStack environment.

We will demonstrate:

- Connecting to the EDB Ark host operating system with `ssh`.
- Connecting to an Ark cluster with the EDB Postgres Enterprise Manager (PEM) client.
- Connecting to an Ark cluster with the `psql` client.
- Copying a file to the EDB Ark host with `scp`.

This tutorial assumes:

- You have already obtained, installed and configured an Ark cluster.
- You are using a Linux workstation.
- The security group for your cluster has been modified to allow client connections on the appropriate ports.

If you need more information about obtaining EDB Ark, or registering a user, visit the EnterpriseDB website at:

<http://www.enterprisedb.com/cloud-database>

This tutorial uses the term *Postgres* to mean an EDB Postgres Advanced Server or PostgreSQL database.

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 Connection Pre-Requisites

Before connecting to an instance that resides in the cloud, an administrator must open the required ports for connections from a client, and you must download the ssh key associated with the cluster.

2.1 Selecting a Port for a Client Connection

Each cluster is associated with a unique OpenStack security group. A security group specifies the addresses from which the cluster members will accept connections. Before connecting to an Ark cluster to perform database server maintenance activities, or connecting with SSH or SCP, an Administrative user with access to the OpenStack Dashboard must modify the security group to allow connections from the connecting client on the required port.

Cluster: sales
 Creation Date: Sat Feb 27 17:30:17 GMT 2016
 DB Username: enterprisedb
 Owner: ppcd.service
 Email: acctg@enterprisedb.com
 Size: 5gb (encrypted)
 Region: uk
 Virtual Network: General VM Network
 Hardware: m1.small
 Engine Version: Postgres Plus Advanced Server 9.5 64bit on CentOS/RHEL 6
 OS/SW update on creation: true
 Monitor Load Balancer Health
 Cluster healing mode:
 Replace failed master with a new master
 Replace failed master with existing replica

Auto-Scaling Thresholds
 % of Storage Size used: 65
 # of Server Connections: 95

DNSNAME	AZ	LBPORT	DBPORT	CXN	VM	HA	DB	UP
172.16.253.13	ox2	9999	5444	4	✓	✓	✓	✓
192.168.1.225	ox		5444	1	✓	✓	✓	✓
192.168.1.226	ox		5444	1	✓	✓	✓	✓

Backup Settings (Greenwich Mean Time)
 Backup Window: 12:00am - 2:00am
 Backup Retention: 7
 Continuous Archiving (Point-in-Time Recovery)

Figure 2.1 – The Details panel of the Clusters Tab.

By default, only port 9999 on the master server node of an Ark cluster is open for client connections. Port 9999 is the load balancing port; connect to this port when executing queries against the database or modifying data.

If you are modifying a database object, invoking administrative functions, or performing server maintenance activities you should connect to the Postgres server's listener port, identified in the DBPORT column, on the Details panel of the Clusters tab (see Figure 2.1). By default, this is port 5444 on an Advanced Server cluster, and port 5432 on a PostgreSQL cluster.

ssh and scp sessions will connect to the Ark cluster on port 22.

2.2 Downloading your Private Key

EDB Ark uses public-key authentication to authenticate a client connecting to the operating system of any hosted node. Each cluster has a cluster-specific key.



To download your private key, navigate to the `Clusters` tab of the EDB Ark console, and click the `Download SSH Key` icon located to the left of the window. The `Accessing Your Cluster Instance` popup opens (see Figure 2.2).

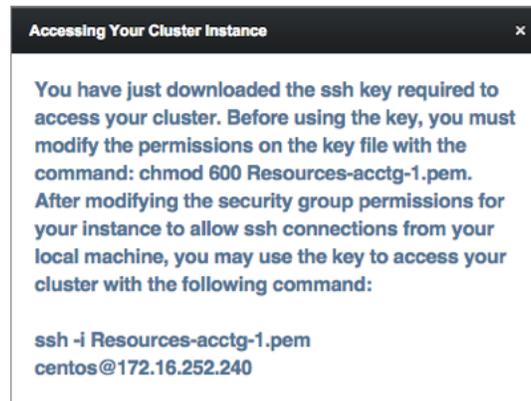


Figure 2.2 – Accessing Your Cluster Instance.

The popup displays information that will help you connect to your cluster, including the key name, the key file permissions required, and the command you can use in a terminal window when connecting.

After downloading your private `ssh` key, open a terminal window, navigate to the location of the `ssh` key file, and modify the key permissions with the command:

```
$ chmod 0600 ssh_key
```

Where:

`ssh_key` specifies the complete path and name of the EDB Ark `ssh` private key file.

After setting the permissions for the key file, you can use the key to connect to a host operating system on any node of an Ark cluster.

3 Using ssh to Access a Server

EDB Ark creates an `ssh` key when you create a new cluster; each cluster has a unique key. Before connecting to a Postgres instance that resides on the cloud via an `ssh` encrypted connection, you must download the `ssh` key (see Section [2.2](#)), and adjust the privileges on the key file.

After adjusting the privileges on the key file, open a terminal window, and specify the location of the `ssh` key, the identity of the connecting user, and the host address in the `ssh` command:

```
ssh -i ssh_key user@host_address
```

Where:

ssh_key specifies the complete path and name of the EDB Ark `ssh` private key file.

user specifies the user account that is connecting to the cluster. When you download the `ssh` key for your cluster, the user name will be provided on the EDB Ark popup.

host_address specifies the IP address of the node to which you wish to connect. You can locate the address in the `DNSNAME` column, on the `Details` panel of the `Clusters` tab in the EDB Ark console.

After connecting with `ssh`, you can:

- Stop, start, or restart the Postgres server.
- Use the Postgres Client Applications.
- Invoke Postgres Server Applications.

Please note: Postgres Server applications must be invoked by the Postgres cluster owner (identified when creating an Ark cluster as the `Master User`). If you are using a PostgreSQL server, the default user name is `postgres`; if you are using Advanced Server, the default user name is `enterprisedb`. To change your identity after connecting via `ssh`, use the `sudo su` command:

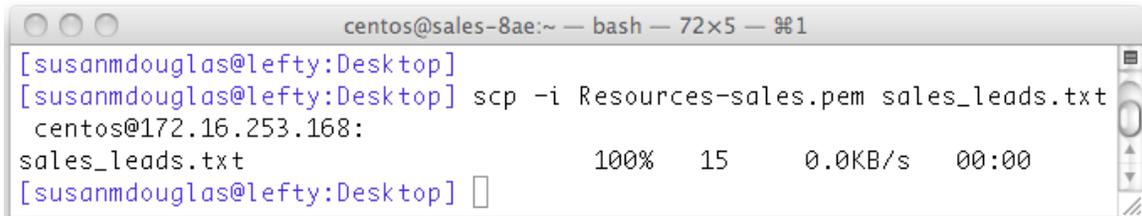
```
# sudo su user_name
```

Where *user_name* is the identity you would like to assume.

4 Copying a file to an EDB Ark Cluster

You can use the `scp` command to copy a file to an Ark cluster; include the `-i` option to specify the location of the key file associated with your cluster (see Figure 4.1):

```
scp -i /path/ssh_key file_name user@host_address:
```



```
centos@sales-8ae:~ — bash — 72x5 — %1
[susanmdouglas@lefty:Desktop]
[susanmdouglas@lefty:Desktop] scp -i Resources-sales.pem sales_leads.txt
centos@172.16.253.168:
sales_leads.txt          100%  15    0.0KB/s  00:00
[susanmdouglas@lefty:Desktop] █
```

Figure 4.1 - Moving a file to EDB Ark.

Where:

path specifies the location of your EDB Ark `ssh` certificate on the system from which you are connecting.

ssh_key specifies the name of the EDB Ark `ssh` private key file.

file_name specifies the name of the file you are copying to EDB Ark.

user specifies the name of the connecting user.

host_address specifies the IP address of the master node of the Ark cluster; the address is displayed in the `DNSNAME` column, on the `Details` panel of the `Clusters` tab in the EDB Ark console.

The colon (`:`) at the end of this command specifies that the file will be copied to the `root` directory on the cluster's primary node. You can specify a file destination by adding a destination path after the colon.

5 Connecting with the Postgres Enterprise Manager Client

The EDB Postgres Enterprise Manager (PEM) client provides a powerful graphical interface that you can use to create and manage database objects (and privileges) on a local Postgres installation or an EDB Ark cluster. The PEM client is installed with the Advanced Server graphical installer, and is also available for PostgreSQL users. For more information about installing the PEM client, visit:

<http://www.enterprisedb.com/products-services-training/products/postgres-enterprise-manager>

The PEM client should be installed on and invoked from a local workstation. Before connecting with the PEM client, an administrator must ensure that the cluster's security group will allow connections from the host of the PEM client.

To open the PEM client, navigate through the Start menu (on Linux) or the Apps menu (on Windows), selecting Postgres Enterprise Manager v6 from the Postgres Plus Advanced Server 9.5 menu (see Figure 5.1).

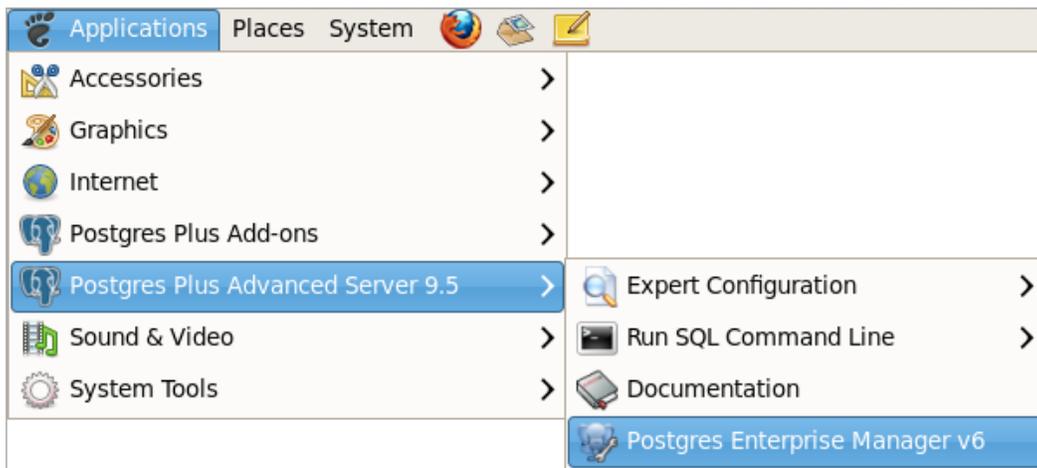


Figure 5.1 – Opening the PEM Client.

Before using the PEM client to access a cluster, you must register the server. You can use the New Server Registration dialog to register the server.

To open the New Server Registration dialog, select Add Server from the File menu.

The screenshot shows a 'New Server Registration' dialog box with the following fields and values:

- Name: (empty text box)
- Host: (empty text box)
- Port: 5444
- Service: (empty text box)
- Maintenance DB: postgres (dropdown menu)
- Username: enterprisedb
- Password: (empty text box)
- Store password:
- Store on PEM Server:
- Colour: (empty text box)
- Group: Servers (dropdown menu)
- Team: (empty text box)

Buttons at the bottom: Help, OK, Cancel.

Figure 5.2 - Connecting to the EDB Ark host.

Provide information about the server in the New Server Registration dialog (see Figure 5.2):

- Specify a descriptive name of the EDB Ark cluster in the `Name` field.
- Provide the IP address of the server in the `Host` field. You can find the IP address in the `DNSNAME` column on the `Details` panel for the cluster on the EDB Ark console.
- Specify the `Port` through which you wish to connect to the server.
If you are modifying a database or invoking administrative functions, you should connect to the master node's listener port, identified in the `DBPORT` column, on the `Details` panel of the `Clusters` tab. Before connecting to the server's listener port, an OpenStack administrator must modify the cluster's security group to allow connections from the connecting client.
- Select a maintenance database using the drop-down listbox in the `Maintenance DB` field. Select `edb` if you are connecting to an Advanced Server database, or `postgres` if you are connecting to a PostgreSQL database.
- Specify the role name that the PEM client should use when connecting in the `Username` field.
- Provide the password associated with that role, in the `Password` field.

Click OK to connect; once connected, the server will appear in the tree control in the PEM Object browser (shown in Figure 5.2).

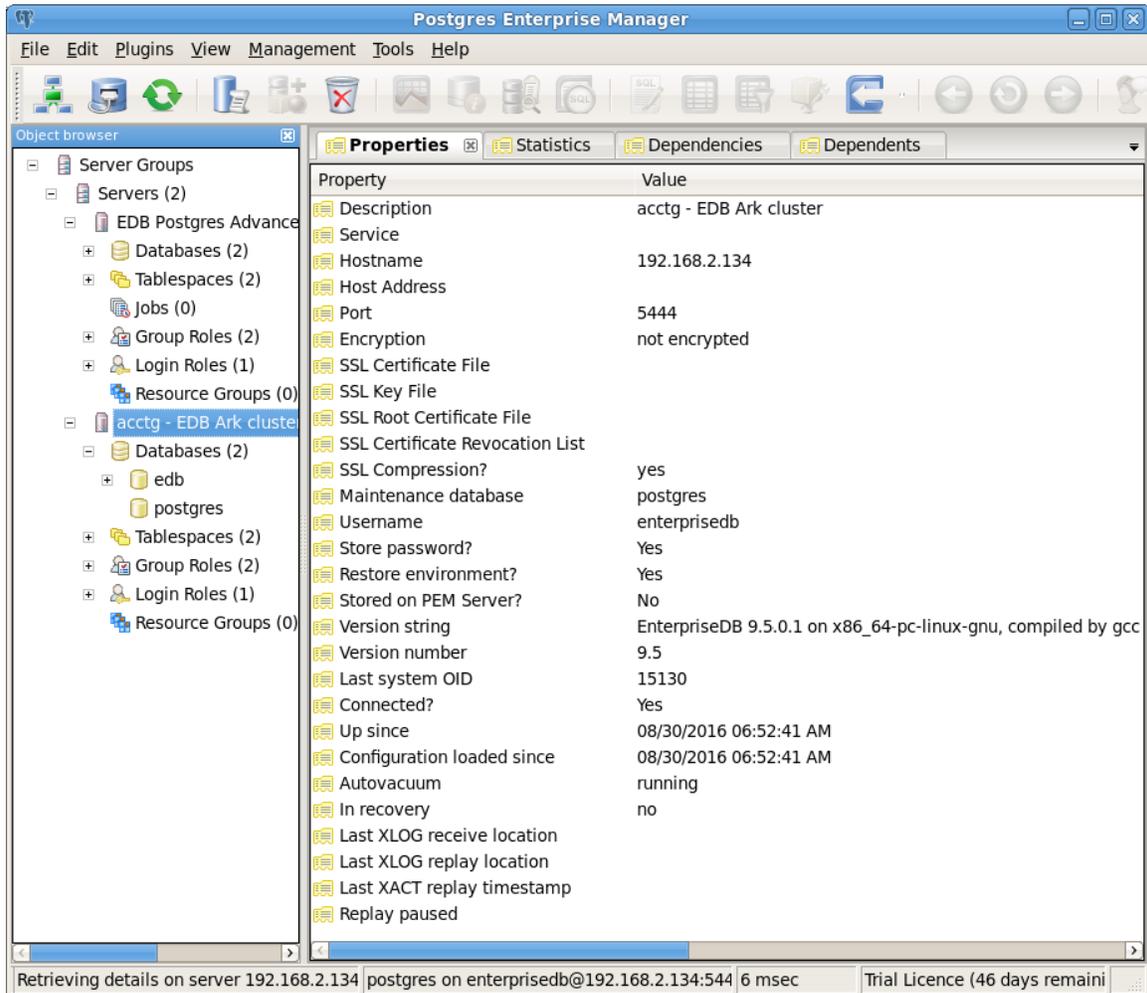


Figure 5.2 - The PEM client window, showing local and EDB Ark servers.

Now, you are ready to use the point-and-click functionality of the PEM client to create and manage database objects that reside on the node of your EDB Ark cluster to which you have connected.

The PEM client offers context-driven help; click the `Help` button provided on each dialog to access online documentation with information about the current dialog.

6 Connecting to EDB Ark with `psql` or `edb-psql`

`psql` is the Postgres command-line client; `edb-psql` is the Advanced Server command-line client. After connecting to a Postgres server hosted on EDB Ark via `psql` or `edb-psql`, you can invoke SQL commands or use meta-commands to:

- Execute queries
- Insert, update, and delete data
- Create and manage database objects (tables, indexes, views, etc.)
- Create user roles and manage privileges
- Review object and role attributes
- Invoke scripts containing complex (or simple) commands

`edb-psql` offers complete compatibility with `psql`, while adding the ability to process syntax compatible with Oracle for anonymous blocks, `CREATE FUNCTION` statements, `CREATE PROCEDURE` statements and `CREATE PACKAGE` statements on an Advanced Server database. You can also use `edb-psql` to manage Advanced Server extensions (such as SQL Protect), or use features compatible with Oracle databases not found in PostgreSQL.

By default, an Ark cluster is only open to connections via port 9999 on the master node. Port 9999 is a good choice if you are connecting for the purpose of querying the database, but if you are modifying database objects, or performing administrative functions, you should connect directly to the server's listener port.

Some administrative functions, if executed over port 9999, may be directed to the incorrect node of a multi-node cluster where they may not have the intended effect, or may return an invalid value.

The listener port number is displayed in the `DBPORT` column of the `Details` panel of the `Clusters` tab.

Before connecting to the server's listener port, an OpenStack administrator must modify the security group to allow connections from the host of your client application.

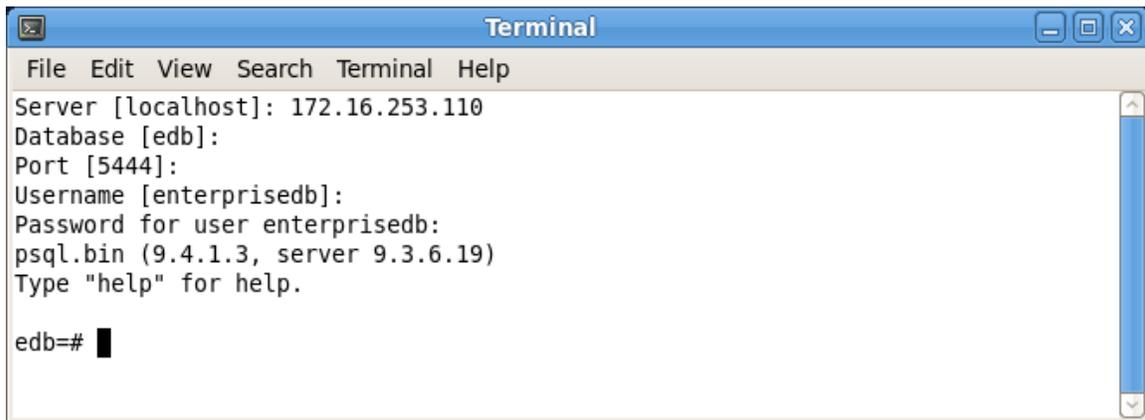
Connecting with edb-psql (or psql) From a Local Workstation

After installing Advanced Server or PostgreSQL on a local workstation, you can use `psql` to perform administrative tasks on an EDB Ark cluster.

To open an `edb-psql` console on an Advanced Server workstation that resides on a Linux host, navigate through the Applications (or Start) menu to the PostgreSQL Plus Advanced Server menu; then, open the Run SQL Command Line menu, and select EDB-PSQL.

To open the client on a Windows host, select `edb-psql` from the Apps menu.

To open a `psql` console on a PostgreSQL workstation, navigate through the Applications (or Start) menu to the PostgreSQL menu, and select SQL Shell (`psql`).



```

Terminal
File Edit View Search Terminal Help
Server [localhost]: 172.16.253.110
Database [edb]:
Port [5444]:
Username [enterprisedb]:
Password for user enterprisedb:
psql.bin (9.4.1.3, server 9.3.6.19)
Type "help" for help.
edb=#

```

Figure 6.1 - The EDB-PSQL command line utility.

Provide connection information for the EDB Ark server to which you are connecting:

- When prompted for a `Server`, enter the IP address or DNS name of the EDB Ark server. The IP address is displayed in the `DNSNAME` column on the `Details` panel of the `Clusters` tab of the EDB Ark management console.
- When prompted for a `Database`, enter the name of the database to which you wish to connect. By default, an Advanced Server cluster is created with a database named `edb`; a PostgreSQL cluster is created with a database named `postgres`.
- When prompted for a `Port`, enter the port on which the server is listening. For database queries, you can use port `9999`; if you are modifying database objects or performing administrative functions, you should use the server's listener port (`5444` for an Advanced Server cluster, `5432` for a PostgreSQL cluster).

- When prompted for a `Username`, enter the role you wish to use when connecting to the server. The name of the database superuser is specified in the `Master User` field when defining an EDB Ark server cluster. By default, the Advanced Server database superuser is `enterprisedb`. The default superuser of a PostgreSQL database is `postgres`.
- When prompted for a `Password`, enter the password associated with that role. The database superuser's password is specified in the `Master Password` field when defining an EDB Ark server cluster.

After connecting, the `edb-psql` (or `psql`) prompt will display the name of the database to which you are connected (as shown in Figure 6.1).

For information about using `psql` and the `psql` meta-commands, please see the Postgres documentation at:

<http://www.postgresql.org/docs/9.5/static/app-psql.html>

For more information about the Postgres SQL commands, please see the Postgres documentation at:

<http://www.postgresql.org/docs/9.5/static/sql-commands.html>