



EDB Postgres Enterprise Manager[®] EDB Ark Management Features Guide

Version 7.8

June 05, 2019

EDB Postgres Enterprise Manager EDB Ark Management Features Guide
by EnterpriseDB® Corporation
Copyright © 2013 - 2019 EnterpriseDB Corporation. All rights reserved.

EnterpriseDB Corporation, 34 Crosby Drive Suite 201, Bedford, MA 01730, USA
T +1 781 357 3390 **F** +1 978 467 1307 **E** info@enterprisedb.com **www**.enterprisedb.com

EnterpriseDB, EDB Postgres, Postgres Plus, Postgres Enterprise Manager, and DynaTune are trademarks of EnterpriseDB Corporation. Other names may be trademarks of their respective owners. © 2019.

Table of Contents

1	Introduction.....	4
1.1	What's New.....	5
1.2	Typographical Conventions Used in this Guide.....	6
1.3	Prerequisites.....	7
2	Using PEM to Manage EDB Postgres Ark.....	8
2.1	Registering an Ark Console.....	8
2.2	Managing an Ark Console.....	11
2.2.1	Downloading Console Logs.....	14
2.2.2	Reviewing the Console Deployment Options.....	15
2.3	Performing Administrative Tasks.....	19
2.3.1	Defining a Server Image.....	20
2.3.2	Defining a Database Engine.....	22
2.3.2.1	Managing Database Engines.....	25
2.3.3	Creating a RHEL Subscription.....	28
2.3.3.1	Managing RHEL Subscriptions.....	32
2.3.4	Creating a Cluster Template.....	34
2.3.4.1	Managing Templates.....	41
2.3.5	User Management.....	43
2.4	Creating and Managing a Cluster.....	45
2.4.1	Using the Create Cluster Wizard.....	46
2.4.1.1	Reviewing a List of Existing Clusters.....	54
2.4.2	Managing a Cluster.....	55
2.4.2.1	Connecting to a Cluster Instance.....	59
2.4.2.2	Taking a Cluster Backup.....	60
2.4.2.3	Scaling a Cluster.....	62
2.4.2.4	Scaling a Machine Type.....	64
2.4.2.5	Upgrading a Cluster.....	65

1 Introduction

Postgres Enterprise Manager (PEM) is designed to assist database administrators, system architects, and performance analysts when administering, monitoring, and tuning PostgreSQL and Advanced Server database servers. PEM has been designed to manage and monitor a single server or multiple servers from a single console, allowing complete control over monitored databases.

EDB Ark provisions EDB Postgres Advanced Server or PostgreSQL databases in single instances, high-availability clusters, or application development sandboxes on Amazon AWS, Azure, and OpenStack hosts. EDB Ark offers an elastic and highly scalable environment that frees DBAs and application developers from the rigors of setting up and administering modern and robust database environments.

For detailed information about EnterpriseDB cloud offerings, visit:

<https://www.enterprisedb.com/cloud>

Throughout this guide, the term *Postgres* refers to either a PostgreSQL or an Advanced Server installation, where either is appropriate.

1.1 What's New

The following features have been added to create Postgres Enterprise Manager 7.8:

- PEM 7.8 allows you to set the `enable_heartbeat_connection` to `true` during PEM agent registration; this creates a dedicated heartbeat connection between the PEM Agent and Server to constantly update the PEM server.
- The PEM 7.8 interface has improved accessibility features such as keyboard navigation enabled across dialogs and wizards, and new shortcut keys.
- PEM 7.8 adds functionality that makes the Query Tool history persistent across sessions.
- Support has been added for the `IDENTITY` columns supported in PostgreSQL and EDB Postgres Advanced Server 10+.

1.2 *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 term(s) 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”.

1.3 Prerequisites

You must modify the display preferences for your PEM web interface to display the Ark Servers node in the tree control. Navigate through the File menu to open the Preferences dialog; then, click on the Nodes panel in the Ark Servers node to control the object types that are displayed.

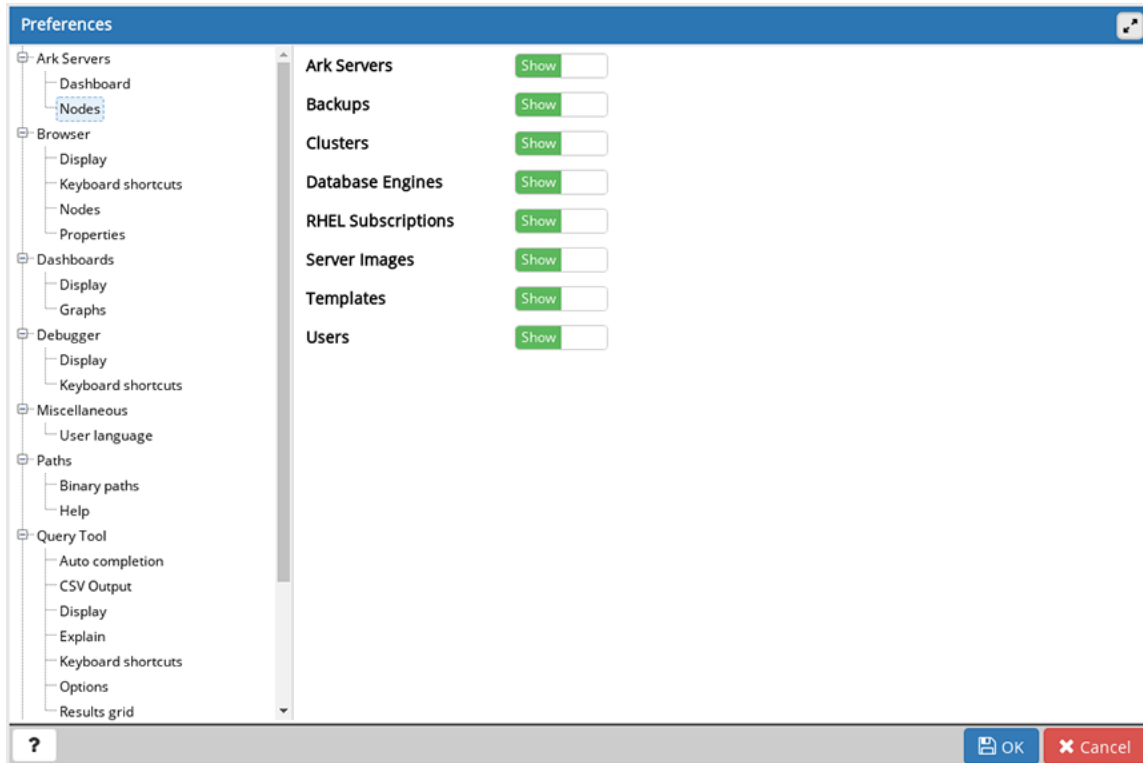


Figure 1.1 - The Nodes panel of the Ark Servers node.

Use the controls to Show or Hide each type of object that is associated with an Ark server. When you've specified your preferences, select OK to preserve your changes, or Cancel to exit without saving.

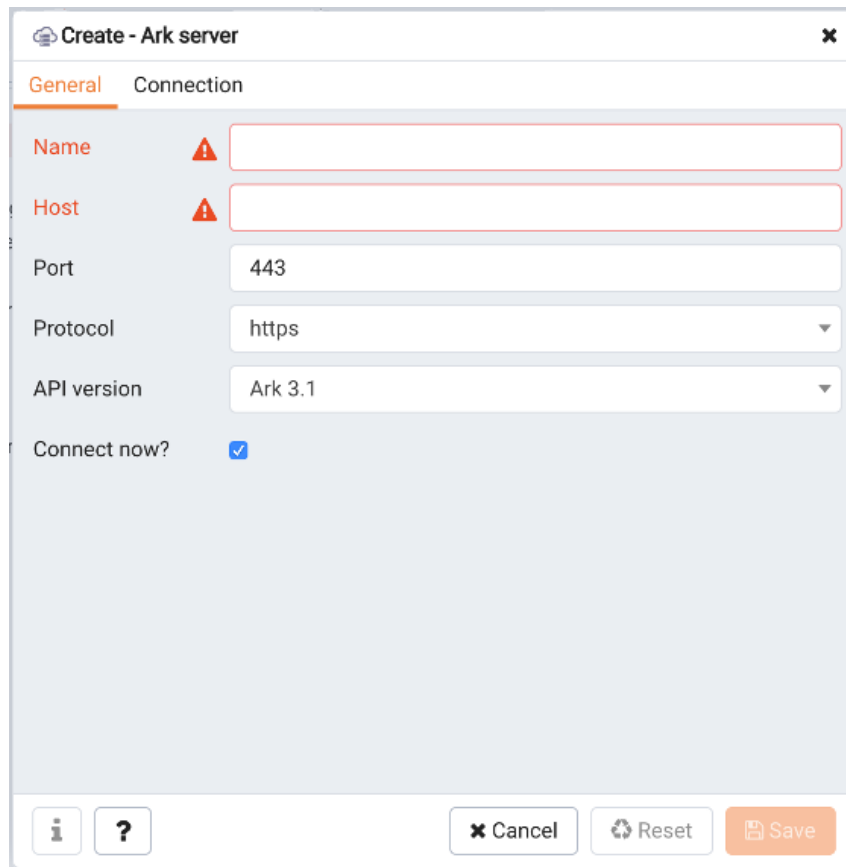
2 Using PEM to Manage EDB Postgres Ark

PEM provides a user-friendly interface that allows you to manage your Ark consoles and create and scale Ark clusters. The PEM web interface helps you manage Ark by providing a point-and-click interface accessed in your favorite browser. Object definitions on existing Ark consoles (clusters, engines, server images, etc.) will be displayed in the PEM object browser when you register an Ark console with PEM.

Before you can use PEM's management features, you must register your Ark console with the PEM server.

2.1 Registering an Ark Console

You can use the `Create - Server` dialog to register an existing Ark console with the PEM server. To access the dialog, right-click on an Ark Server node to access the context menu; navigate through the `Create` option to select `Ark Server . . .`



The screenshot shows a dialog box titled "Create - Ark server". It has two tabs: "General" and "Connection". The "General" tab is active. The form contains the following fields and controls:

- Name:** A text input field with a red warning triangle icon to its left.
- Host:** A text input field with a red warning triangle icon to its left.
- Port:** A text input field containing the value "443".
- Protocol:** A dropdown menu with "https" selected.
- API version:** A dropdown menu with "Ark 3.1" selected.
- Connect now?:** A checkbox that is checked.

At the bottom of the dialog, there are three buttons: "Cancel", "Reset", and "Save". To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Figure 2.1 - Providing general information about a console.

Use the fields on the `General` tab to describe the general properties of the Ark console:

- Use the `Name` field to specify a user-friendly name for the server. The name specified will identify the server in the PEM Browser tree control.
- Use the `Host` field to provide the name of the Ark console host.
- Use the `Port` field to specify the port that the PEM server will use when connecting to the Ark console.
- Use the `Protocol` drop-down listbox to select the connection protocol that the PEM server will use when connecting to the Ark console.
- Use the `API version` drop-down listbox to select the version of the Ark API that the PEM server will use when managing the Ark console.

Create - Ark server

General **Connection**

Username

Password

Tenant/Role

Domain name

Domain name is applicable to openstack ark console and it required when ark api version is 3.1 or greater

Save password?

Cancel **Reset** **Save**

Figure 2.2 - Providing connection information.

Use fields on the `Connection` tab to specify connection details for the Ark console:

- Specify the name of the connecting user in the `Username` field. The specified user must be an existing user in the backing database of the Ark console.
- Provide the password associated with the specified user in the `Password` field.
- Specify the name of the OpenStack tenant or `Ark Port` field.
- Use the `Domain name` field to specify the name of the OpenStack domain to which the server will connect if the Ark API version specified on the `General` tab is version 3.1 or greater.
- Check the box next to `Save password?` to save the password associated with the console registration; if do not save the password, you will be prompted for the password when you attempt to connect to the server.

When you have provided registration information for the Ark console, click the `Save` button to save the Ark console definition.

2.2 Managing an Ark Console

The PEM web interface provides dialogs that allow you to review and modify the Ark console attributes.

The `Dashboard` tab displays icons that provide quick access to Ark information and features. If the tab is accessed from the top-level `Ark Servers` node, it includes informational links that are applicable to all consoles. If the tab is accessed from a console name node, the tab includes information specific to that console:

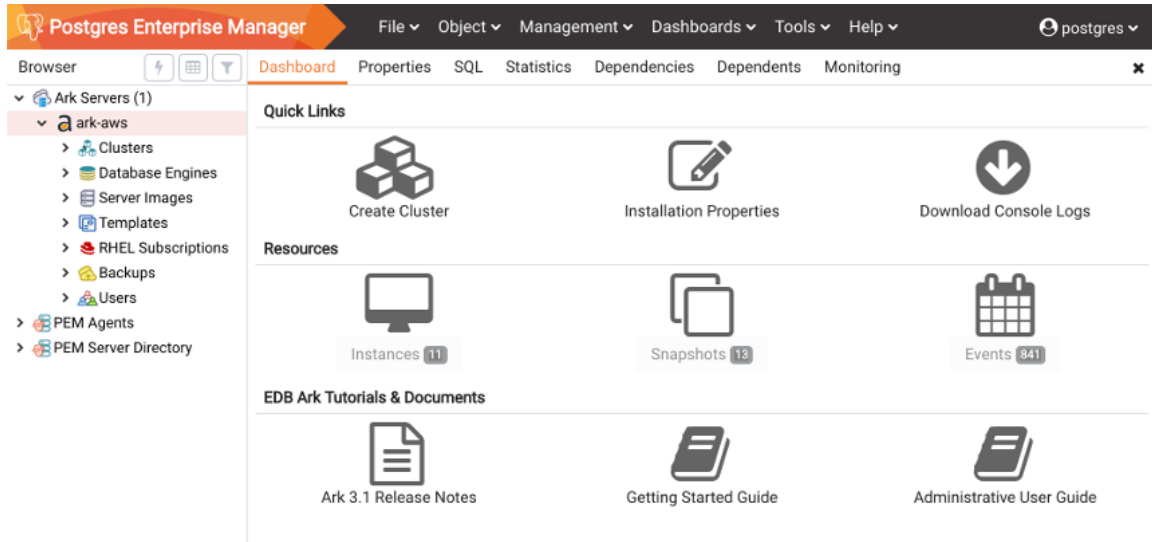


Figure 2.3 - The Dashboard tab.

The `Quick Links` icons allow you to:

- Click the `Create Cluster` icon to define a new cluster on the selected console.
- Click the `Installation Properties` icon to review the deployment properties for the selected console.
- Click the `Download Console Logs` icon to download the log files for the selected console.

The `Resources` icons provide access to information about the objects that reside on the console:

- Click the `Instances` icon to review information about the instances that reside on the console.
- Click the `Snapshots` icon to review information about console backups.

- Click the `Events` icon to review information about events that have been logged for the selected console.

The `EDB Ark Tutorials & Documents` icons provide access to information that is specific to the Ark version that is deployed on the console:

- Click the `Ark 3.0 Release Notes` icon to review the release notes.
- Click the `Getting Started Guide` icon to review a usage guide for the Ark console.
- Click the `Administrative User Guide` icon to review the Ark Administrative User Guide.

To view console properties, select the name of a console in the PEM Browser tree control and select the `Properties` tab.

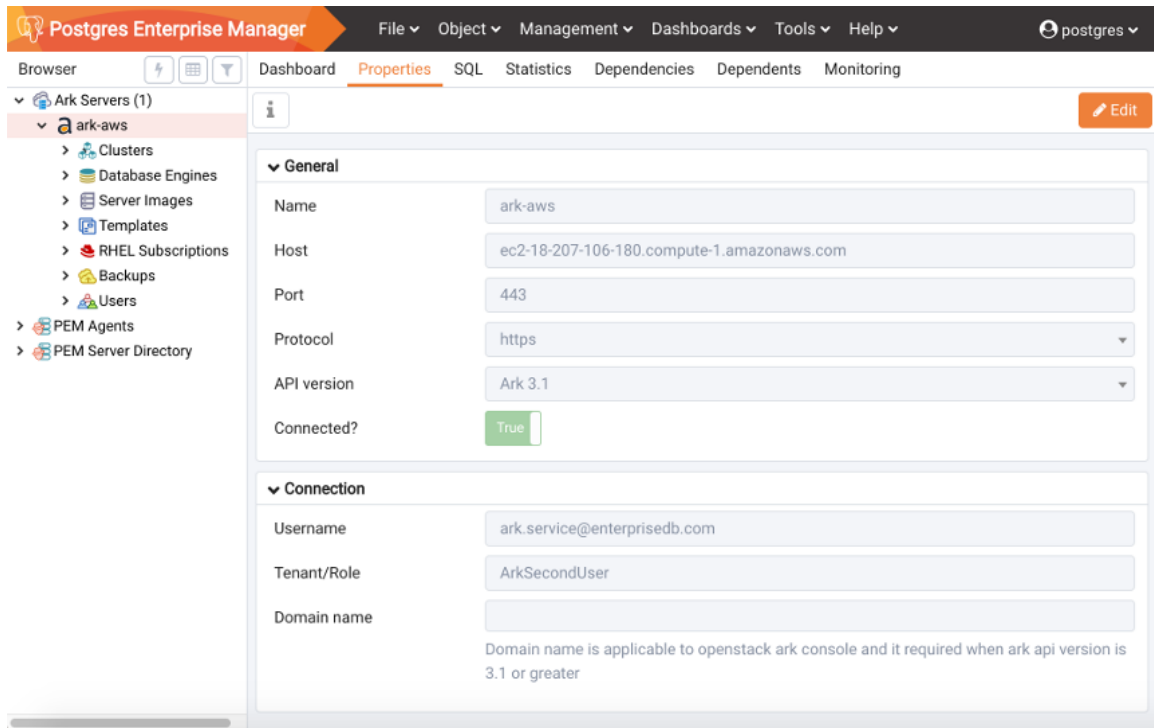


Figure 2.4 - The Properties tab.

Panels on the `Properties` tab display configuration details for the console.

The `General` panel displays:

- The `Name` field displays the console name.

- The `Host` name displays the IP address of the console.
- The `Port` field displays the port at which a browser can connect to the console.
- The `Protocol` field displays the connection protocol used to connect to the console.
- The `API version` field displays the version of the API used to manage console objects.
- The `Connected?` field displays the current state of the console.

The `Connection` panel displays:

- The `Username` field displays the name of the user that deployed the console.
- The `Tenant/Role` field displays the name of the tenant or role in which the console is hosted.
- If deployed on an OpenStack host and using API version 3.1 or later, the `Domain name` field displays the OpenStack domain in which the console is hosted.

Use options on the console's context menu to manage a console registration; right-click on the console name in the PEM browser tree control and select from the following menu options:

Option	Action
Refresh	Select <code>Refresh</code> to update the definition of the server or server objects that reside on the console.
Connect Server	Select <code>Connect Server</code> to connect to the selected Ark console. After authenticating with the Ark console host, the tree control will display a host-specific icon to the left of the console name. Expand the tree control to review and manage objects that reside on the console.
Delete/Drop	Select <code>Delete/Drop</code> to delete the registration of the selected Ark console. Please note that deleting a console registration does not remove the console from the backing database. PEM will confirm that you wish to delete the registration before removing it.
Download console logs	Select <code>Download console logs</code> to download the log files from the selected console.
Installation options	Select <code>Installation options</code> to open the Installation options dialog and review or modify console deployment properties.
Disconnect Server	Select <code>Disconnect Server</code> to disconnect from the selected Ark console.
Properties	Select <code>Properties</code> to review the properties specified when registering an Ark console; if the server is disconnected, you may modify the property values. After modifying property values, select <code>Save</code> to preserve your changes.

2.2.1 Downloading Console Logs

Use options on the console `Dashboard` panel or context menu to download a zip file that contains the server logs for the underlying application server. You can use the log file to confirm changes to server status or verify server activity. To download a console log you can either:

- Right click on the console name in the PEM Browser tree control and select `Download console logs` from the context menu.
- Highlight a console name in the tree control and navigate to the `Dashboard` tab; then, select the `Download console logs` option from the `Quick Links` panel.

A popup will ask you to confirm your selection; click `OK` to download the log file.

You can also review the console logs via an ssh session. Log files are stored in `/var/log/edb-ark`; the current log file is `/var/log/edb-ark/ark.log`. Use the Linux `tail` utility to display the most recent entries in any of the server logs. For example, to review the last 10 lines in the server log file, use SSH to connect to the console host's backing database and enter:

```
tail file_name
```

Where `file_name` specifies the complete path to the log file. Include the `-F` option to instruct the `tail` utility to display only the last 10 lines of the log file, and new log file entries as they are added to the file:

```
tail -F file_name
```

2.2.2 Reviewing the Console Deployment Options

The `Installation options` dialog displays the deployment properties of the selected console. To open the dialog, right-click on a console name and select `Installation options` from the context menu.

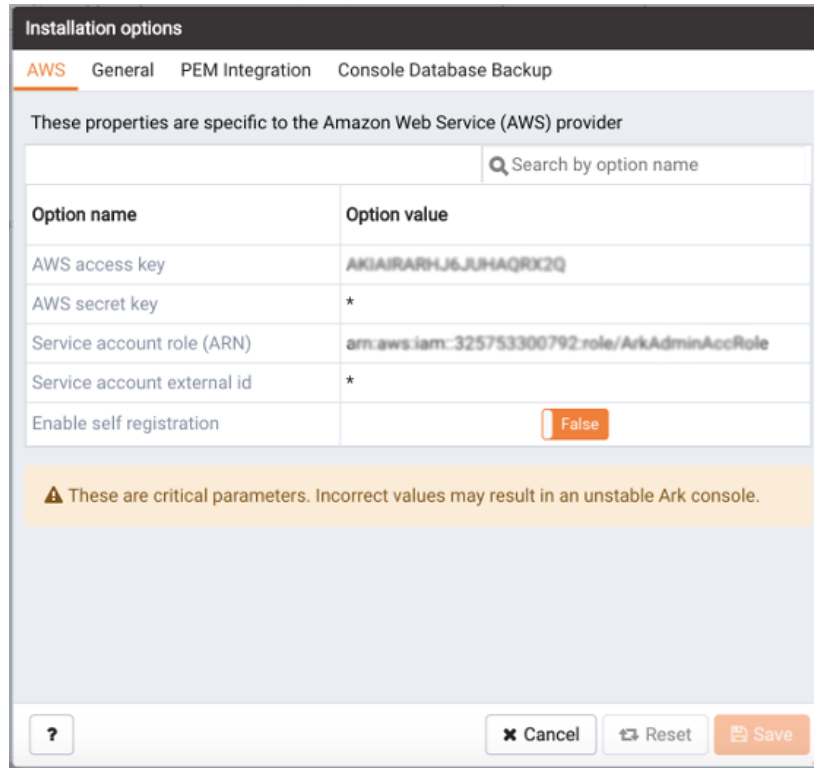


Figure 2.5 - The Ark console Installation Options dialog.

The first tab of the `Installation options` dialog displays provider-specific details about the console deployment. The properties displayed may include:

Option	Action
<code>AWS cross account accesskey</code>	The Amazon access key ID associated with the AWS role that will be used for account.
<code>AWS cross account secretkey</code>	The Amazon secret key associated with the AWS role that will be used for account administration.
<code>Service account externalid</code>	The Amazon external ID that should be used by the Ark service user.
<code>Service account rolearn</code>	The Amazon Role ARN (resource name) that should be used by the Ark service user when performing management functions on behalf of Ark.
<code>Enable self registration</code>	If the Ark console should allow self-registration for Ark users.
<code>Azure application registration ID</code>	The application ID associated with the Azure account that hosts the Ark console.

EDB Postgres Enterprise Manager EDB Ark Management Features Guide

Azure active directory ID	The directory ID associated with the Azure account that hosts the Ark console.
Azure storage account	The Azure block storage account you wish to use with this Ark server.
Azure subscription ID	The Azure account that hosts the Ark console.
OpenStack admin role	The name of the OpenStack administrative user.
OpenStack identity service endpoint	The location of the OpenStack identity service endpoint.
Service account ID	The name of the service account on the console's backing database.
Service account password	The password for the postgres user on the console's backing database.

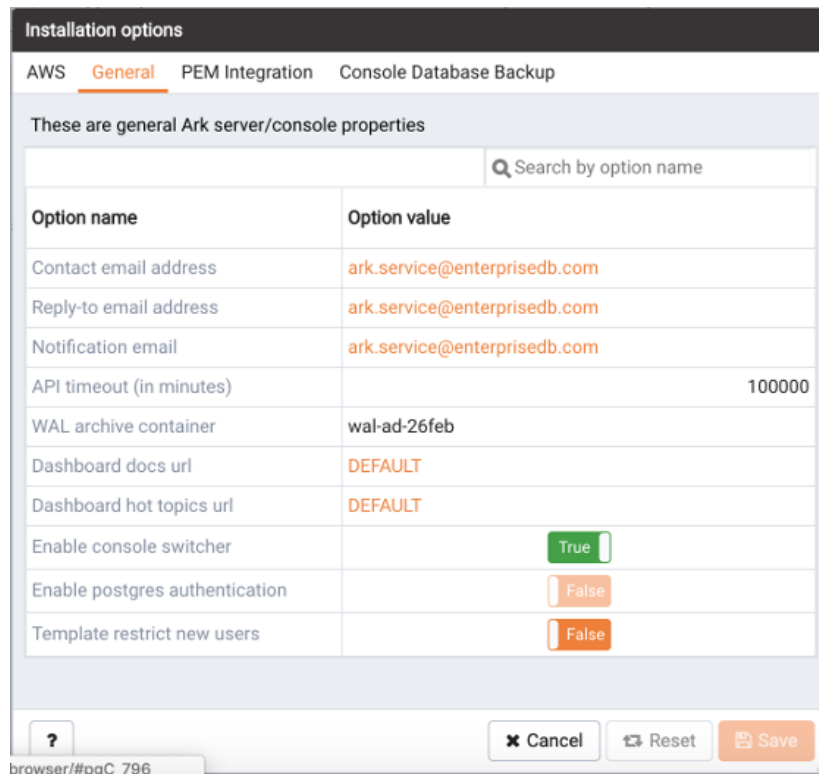


Figure 2.6 - The General tab.

The `General` tab displays details about behavior related aspects of the console deployment. Displayed information may include:

Option	Action
API timeout	The number of minutes that an authorization token will be valid for use with the API.
Contact email address	The email address that will be included in the body of cluster status notification emails.
Dashboard docs url	The location of the content that will be displayed on the Dashboard tab of the Ark console.
Dashboard hot topics url	The location of the content that will be displayed on the Dashboard tab of the Ark console.

EDB Postgres Enterprise Manager EDB Ark Management Features Guide

Database backup tenant	The name of the tenant in which the database backup will reside.
Database name	The name of the console database.
Directory to store backups	The directory in which backups will be stored.
Enable console switcher	If the console should display console switcher functionality.
Enable PostgreSQL authentication	If Ark will enforce the authentication method configured on the backing Postgres server.
Folder to store backups	The name of the backup folder within the storage bucket.
Notification email	The email address to which email notifications about the status of the Ark console will be sent.
Reply-to email address	The return email address used on cluster status notification emails.
Template restrict new users	Indicates if the Ark console will make any new user a Template Only user by default.
WAL archive container	The name of the container in which WAL archives are stored.

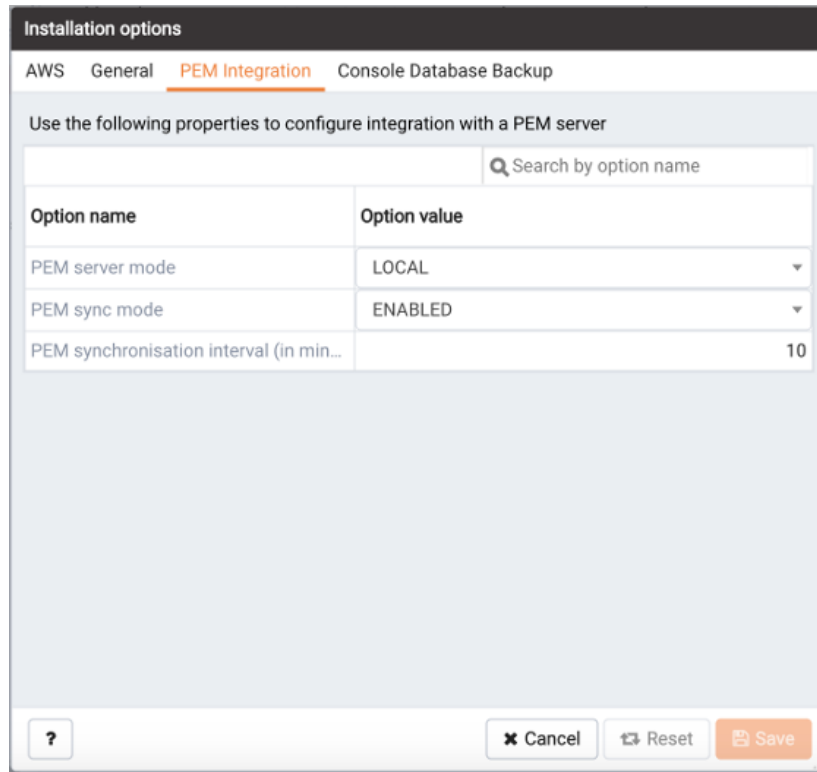


Figure 2.7 - The PEM integration tab.

The PEM integration tab displays details about behavior related to PEM. The properties displayed may include:

Option	Action
PEM server mode	The location of the PEM server; a LOCAL server resides on the monitored PEM instance, while a REMOTE server is on a different host.
PEM sync mode	If enabled, all roles/groups/tenants that are accessible by the role that creates the console are added to the PEM server as a

	member of a non-login group role.
PEM synchronisation interval	The number of minutes between PEM/Ark synchronisations.

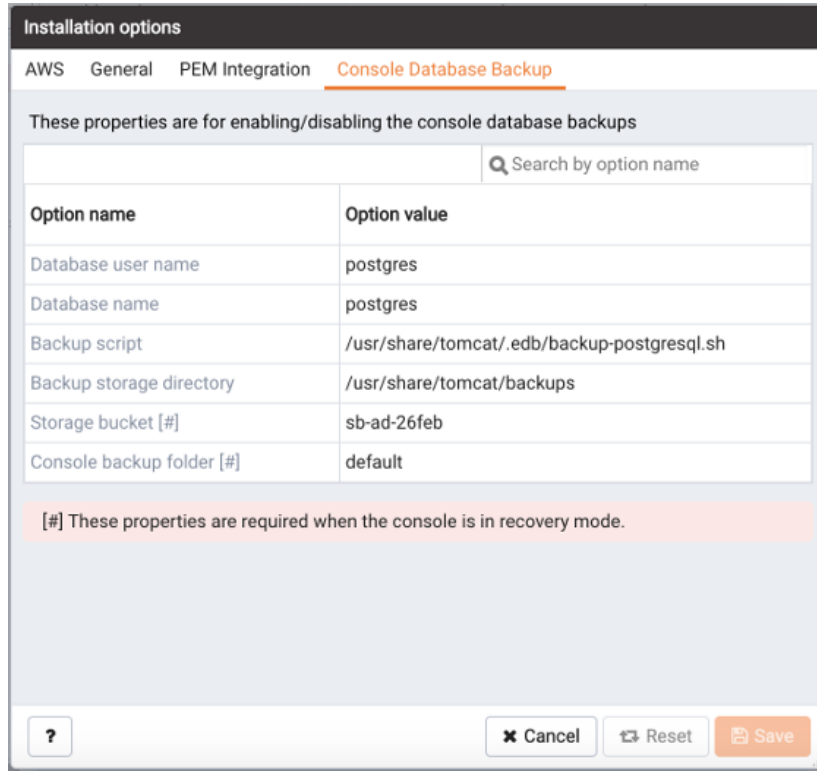


Figure 2.8 - The Console database backup tab.

The Console database backup tab displays details about backup behavior. The properties displayed may include:

Option	Action
Backup script	The complete path to the backup script on the console host.
Backup storage directory	The complete path to the backup script on the console host.
Console backup folder	The name of the console database.
Database name	The name of the console database.
Database user name	The name of the database superuser.
Backup storage directory	The directory in which backups will be stored.
Folder to store backups	The name of the backup folder within the storage bucket.
Storage bucket	The name of the container in which backups will be stored.
Storage tenant	The name of the tenant in which the database backup will reside.

2.3 Performing Administrative Tasks

You can use the PEM web interface to create and manage clusters that reside on the console host. Before creating a cluster, you must:

- define a server image.
- define a database engine.

After defining a server image and a database engine, you can use PEM to create, monitor, and scale Ark clusters.

2.3.1 Defining a Server Image

A server image definition describes the virtual machine that will reside on the console host, and provide the environment for an instance of Advanced Server or PostgreSQL. To create a new server image, right-click on the `Server Images` node, and select `Create`, then `Server Image...`

Figure 2.9 - The Create - Server Image dialog.

Use fields on the `Create - Server Image` dialog to define the server image:

- Use the `Server ID` field to provide an identifier for the server image. The `Server ID` must be unique, and may not be modified after saving the server image.
- Use the `Server description` field to provide a description of the server image.
- Use the `Image ID` field to provide the `Image ID` of the server image.
- Use the `Initial user` field to provide the name of the default operating system user. This user must have `sudo root` privileges to perform the initial provisioning of software on the node.
- Use the `System type` field drop-down listbox to select the operating system type of the server; select `CentOS` or `RHEL`.

- Set the `Statically provisioned?` slider to `Yes` to indicate that the server is statically provisioned. A statically provisioned server is a pre-installed image that contains the software required to create a database cluster.

Deleting a Server Image

To delete a server image, right-click on the image name in the PEM browser tree control, and select `Delete/Drop` from the context menu. The PEM server will prompt you to confirm that you wish to delete the server image. Please note that you cannot remove a server image that is currently in use by an engine; before deleting the server, you must delete any dependent engines.

2.3.2 Defining a Database Engine

An engine definition provides detailed information about the backing database used by a cluster; only an EDB Ark administrative user can define an engine. Once defined, all of the engines that reside within a specific tenant will be made available to all users with access to that tenant. The Ark console ships with a number of default engine definitions.

You can also create custom database engines. To create a new database engine, right-click on the Database Engines node, and select Create, then Database Engine....

Figure 2.10 - The Create - Database Engine dialog.

Use fields on the `General` tab to provide information about the engine:

- Use the `Name` field to provide a user-friendly name for the engine.
- Use the `Engine ID` field to specify the identifier of the engine. The identifier must be unique, and cannot be modified after the engine definition is saved.
- Use the `Database Type` drop-down listbox to select the database type (PostgreSQL or Advanced Server) that will be deployed on the engine. The database type cannot be modified after the engine definition is saved.

- Use the `Version` drop-down listbox to select the database server version of the engine. The version cannot be modified after the engine definition is saved.
- Use the `Server type` drop-down listbox to select the backing operating system that will be used by the engine.

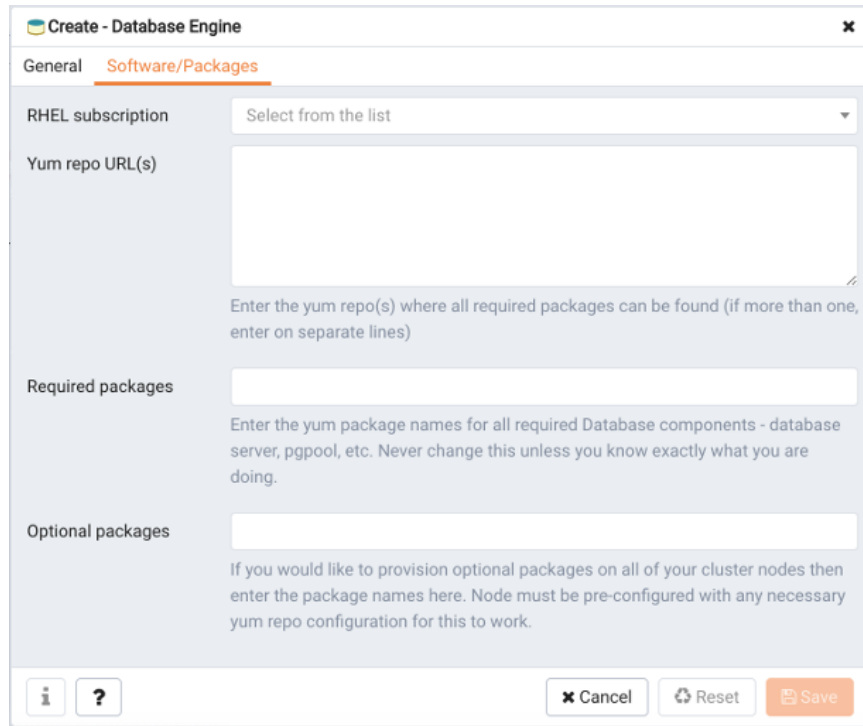


Figure 2.11 - The Software/Packages tab.

Use fields on the `Software/Packages` tab to provide information about the software that will be deployed on the engine:

- Use the `RHEL subscription` drop-down listbox to select the RHEL subscription service that will be used when provisioning the engine. You can use the `Create - RHEL Subscription` dialog to propagate the RHEL subscription list.
- Use the `Yum repo URL(s)` field to provide the names and credentials associated with one or more yum repositories from which software will be installed.
- Use the `Required packages` field to provide a list of the software packages that are required by the specified server.
- Use the `Optional packages` field to provide a list of the additional software packages that will be installed on clusters that are provisioned using the engine.

After defining the database engine click Save to add the engine to the Database Engines list in the PEM Browser tree control.

Before using the engine when deploying an Ark cluster, the engine must be enabled. You can use options on the node's context menu to manage the engine.

Modifying a Database Engine

To modify a database engine, right-click on the name of the database engine, and select `Properties...` from the context menu, or highlight the engine's name and select `Properties...` from the `Object` dialog. Use fields on the engine's properties dialog to update the properties associated with the engine.

2.3.2.1 Managing Database Engines

To view a list of the currently defined engines, select the `Database Engines` node in the PEM Browser tree control and navigate to the `Properties` tab.

Name	Engine ID	Enabled	Database type	Version	Server type	RHEL subscription
PostgreSQL 9.4 64bit on CentOS 6	PG_94_C6_ARK30	False	PostgreSQL Server	9.4		
PostgreSQL 9.4 64bit on CentOS / RHE...	PG_94_CR7_ARK30	False	PostgreSQL Server	9.4		
EDB Postgres Advanced Server 9.4 64...	PPAS_94_ARK30	False	EDB Postgres Advanced Server	9.4		
PostgreSQL 9.5 64bit on CentOS 6	PG_95_C6_ARK30	False	PostgreSQL Server	9.5		
PostgreSQL 9.5 64bit on CentOS / RHE...	PG_95_CR7_ARK30	False	PostgreSQL Server	9.5		
EDB Postgres Advanced Server 9.5 64...	PPAS_95_ARK30	False	EDB Postgres Advanced Server	9.5		
PostgreSQL 9.6 64bit on CentOS 6	PG_96_C6_ARK30	False	PostgreSQL Server	9.6		
PostgreSQL 9.6 64bit on CentOS / RHE...	PG_96_CR7_ARK30	False	PostgreSQL Server	9.6		
EDB Postgres Advanced Server 9.6 64...	PPAS_96_ARK30	False	EDB Postgres Advanced Server	9.6		
PostgreSQL 10 64bit on CentOS 6	PG_10_C6_ARK30	False	PostgreSQL Server	10		
PostgreSQL 10 64bit on CentOS / RHE...	PG_10_CR7_ARK30	False	PostgreSQL Server	10		
EDB Postgres Advanced Server 10 64b...	PPAS_10_ARK30	False	EDB Postgres Advanced Server	10		
PostgreSQL 11 64bit on CentOS 6	PG_11_C6_ARK32	False	PostgreSQL Server	11		
PostgreSQL 11 64bit on CentOS / RHE...	PG_11_CR7_ARK32	True	PostgreSQL Server	11	c7	
EDB Postgres Advanced Server 11 64b...	PPAS_11_ARK32	True	EDB Postgres Advanced Server	11	c7	
Pradip Testing API CR7 PG 10 - 150821	A_PG10CR7_150821	True	PostgreSQL Server	10	A_C7.2_150820	
INITIAL SETUP API-test-eol-false	API-test-eol-false	True	PostgreSQL Server	9.5	c7	

Figure 2.12 - The Engine node's Properties tab.

The `Properties` tab displays detailed information about the currently defined engines:

- The `Name` field displays the user-friendly name of the database engine.
- The `Engine ID` field displays the engine identifier.
- The `Enabled` field displays the current state of the database engine.
- The `Database type` field displays the type of database used by the engine.
- The `Version` field displays the server version used by the engine.
- The `Server type` field displays type of the backing operating system.
- If applicable, the `RHEL subscription` field displays the identifier of the RHEL subscription used by the engine.

You can use options on the engine's context menu to manage the engine. To access the context menu, right-click on the engine name in the PEM Browser tree control.

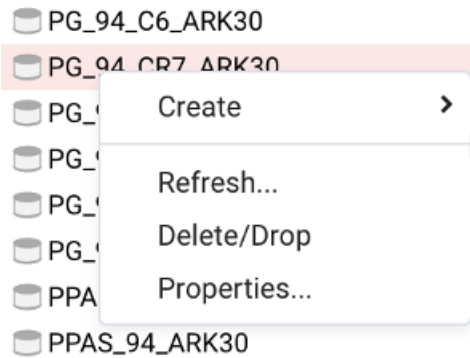


Figure 2.13 - The engine's context menu.

To manage the engine, select from the menu options:

Option	Action
Refresh	Select Refresh to update the definition of the database engine.
Delete/Drop	Select Delete/Drop to delete the engine. PEM will confirm that you wish to delete the engine before removing it.
Enable database engine	Select Enable database engine to make the engine available for use when deploying a cluster.
Disable database engine	Select Disable database engine to make the engine unavailable for use.
Properties...	Select Properties... to review or modify engine properties.

Enabling or Disabling a Database Engine

Enabled engines are available for use when provisioning clusters; disabled engines are not made available. There are two ways to enable or disable a database engine - you can:

- Right-click on the engine name in the PEM browser tree control and select Enable database engine or Disable database engine from the context menu.
- Use the Enabled slider on the database engine Properties dialog to control the state of the engine.

To open the Properties dialog, right-click on the engine name and select Properties... from the context menu or highlight the engine's name and select Properties... from the Object menu.

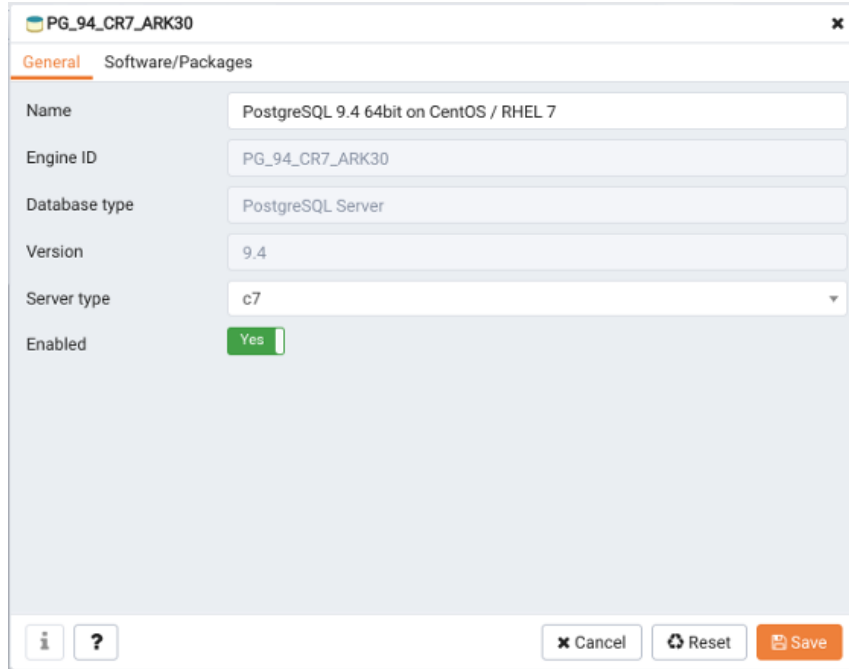


Figure 2.14 - The engine's properties dialog.

Move the `Enabled` slider to the `Yes` position to make the engine available for cluster provisioning. To disable an engine, move the slider to the `No` position.

Please note that disabling an engine has no impact on any running clusters; it simply prevents users from creating new clusters with the engine. You can use this feature to phase out the use of older engines.

Deleting a Database Engine

To delete an engine definition, right-click on the engine name in the PEM browser tree control, and select `Delete/Drop` from the context menu. The PEM server will prompt you to confirm that you wish to delete the engine before the engine is deleted. Please note that you cannot remove an engine that is referenced by one or more clusters and/or backups.

2.3.3 Creating a RHEL Subscription

You can use the PEM web interface to attach Red Hat Subscription Manager information to engines hosted on Red Hat consoles. The Red Hat Subscription Manager tracks installed products and subscriptions to implement content management with tools like yum. For information about Red Hat Subscription Manager, visit the Red Hat website at:

<https://access.redhat.com/documentation/en/red-hat-subscription-management>

When you create a new cluster that uses an engine that is associated with a Red Hat subscription, Ark registers the cluster nodes with Red Hat; when you terminate the node, the system's subscription is unregistered.

To create a RHEL subscription definition, right-click on the RHEL Subscriptions node, and select `Create`, then `RHEL Subscription`....

The screenshot shows a web-based form titled "Create - RHEL Subscription". The form is divided into three tabs: "General", "Activation", and "Software/Packages". The "General" tab is currently selected and contains the following fields:

- Subscription ID
- Username
- Password
- Server URL
- Base URL
- Organization
- Environment
- System name

At the bottom of the form, there are three buttons: "Cancel", "Reset", and "Save". There are also two small icons on the left: an information icon and a question mark icon.

Figure 2.15 - The General tab.

Use fields on the `General` tab to provide information about the subscription:

- Use the `Subscription ID` field to provide a user-friendly name for the subscription. The name will identify the subscription in the RHEL Subscription drop-down on the engine details dialog.
- Use the `Username` field to provide the name of the user account registered with the Red Hat content server.
- Use the `Password` field to provide the password associated with the user account.
- Use the `Server Url` field to provide the host name of the subscription server used by the service; if left blank, the default value of `subscription.rhn.redhat.com` will be used.
- Use the `Base Url` field to provide the host name of the content delivery server used by the service; if left blank, the default value of `https://cdn.redhat.com` will be used.
- Use the `Organization` field to provide the organization that will be registered with the Red Hat subscription system.
- Use the `Environment` field to provide the name of the environment (within the organization that will be registered).
- Use the `System name` field to provide the name of the system that will be registered.

Figure 2.16 - The Activation tab.

Use fields on the `Activation` tab to provide information about subscription behaviors:

- Use the `Activation key` field to provide the activation key of the Red Hat subscription.
- When the `Auto-attach` slider is set to `Yes`, any node associated with the subscription will automatically attach to the service.
- Use the `Pool` field to provide the pool identifier for the Red Hat subscription service.
- When the `Auto` slider is set to `Yes`, nodes provisioned with engines associated with the pool will automatically attach to the subscription service.
- Use the `Quantity` selector to specify the number of subscriptions in the subscription pool.
- Use the `Service level` drop-down listbox to provide the service level of the subscription.
- Use the `Release` field to provide the operating system minor release that will be used when identifying updates to any nodes provisioned with the subscription.

- When the `Force` slider is set to `Yes`, the node will be registered, even if it is already registered.
- Use the `Subscription type` drop-down listbox to specify the type of consumer that is being registered.

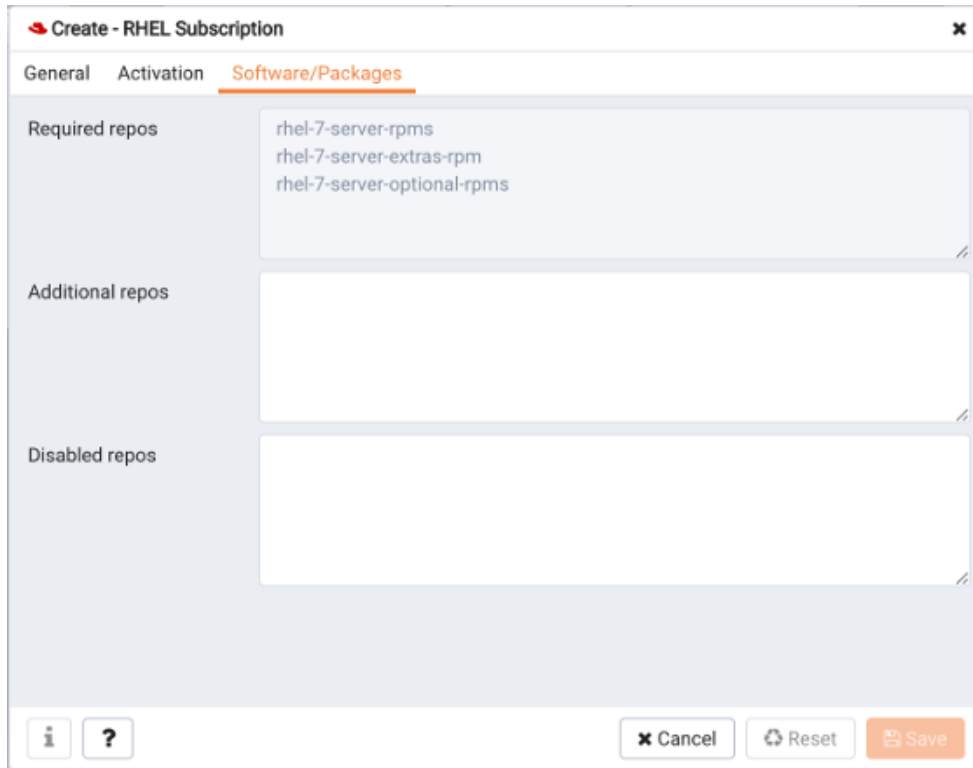


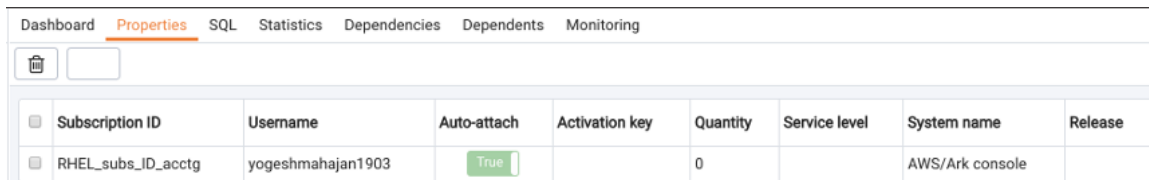
Figure 2.17 - The Software/Packages tab.

Use fields on the `Software/Packages` tab to provide information about the software that will be deployed on engines that use the repository:

- The `Required repos` list is populated by the Ark console, and displays a list of the repositories required by the subscription definition.
- Use the `Additional repos` field to provide the names of any additional repositories that should be enabled on the cluster node(s).
- Use the `Disabled repos` field to provide the names of any repositories that should be disabled on the cluster node(s).

2.3.3.1 Managing RHEL Subscriptions

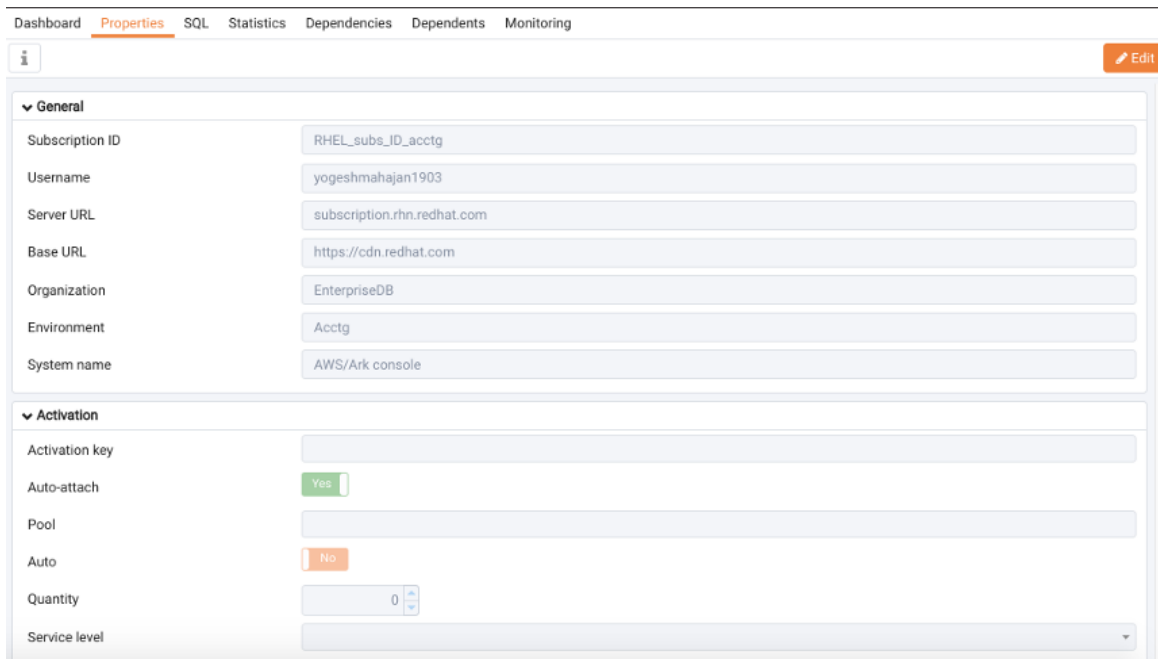
To view a list of the currently defined RHEL subscriptions, highlight the `RHEL Subscriptions` node in the PEM browser tree control and navigate to the `Properties` tab.



Subscription ID	Username	Auto-attach	Activation key	Quantity	Service level	System name	Release
RHEL_subs_ID_acctg	yogeshmahajan1903	<input checked="" type="checkbox"/>		0		AWS/Ark console	

Figure 2.18 - The RHEL Subscriptions list.

To review detailed information about a specific subscription, select the subscription name and navigate to the `Properties` tab.



Dashboard **Properties** SQL Statistics Dependencies Dependents Monitoring

General

Subscription ID: RHEL_subs_ID_acctg

Username: yogeshmahajan1903

Server URL: subscription.rhn.redhat.com

Base URL: https://cdn.redhat.com

Organization: EnterpriseDB

Environment: Acctg

System name: AWS/Ark console

Activation

Activation key:

Auto-attach: Yes

Pool:

Auto: No

Quantity:

Service level:

Figure 2.19 - RHEL Subscription properties.

Panels on the `Properties` tab display details about the subscription. Select the edit button in the header to open a dialog that allows you to edit the subscription properties. Please note that to modify subscription information, you must provide the password associated with the subscription.

- The `General` panel displays subscription information such as the registered user name, your organization information, and the URLs used to connect to the subscription.
- The `Activation` panel displays information about the RHEL subscription service and its behavior.
- The `Software/Packages` panel displays the names of repositories used by the subscription and the subscription manager command.

You can use options on the subscription's context menu to manage the subscription. To access the context menu, right-click on the subscription name in the PEM Browser tree control.

Then, to manage the subscription, select from the menu options:

Option	Action
Refresh	Select Refresh to update the definition of the RHEL subscription.
Delete/Drop	Select Delete/Drop to delete the subscription. PEM will confirm that you wish to delete the subscription definition before removing it.
Properties...	Select Properties... to modify the subscription properties.

Modifying Subscription Properties

To modify the properties associated with a RHEL subscription, right-click on a subscription name and select `Properties...` from the context menu.

Before modifying properties, you must provide the `Password` associated with the subscription. After entering a password and making any desired changes to the subscription definition, click `Save` to preserve your changes and exit the dialog.

Deleting a Subscription Definition

To delete a subscription definition, right-click on the name of a subscription in the PEM Browser tree control, and select `Delete/Drop` from the context menu. The PEM server will prompt you to confirm that you wish to delete the subscription before the subscription is deleted.

2.3.4 Creating a Cluster Template

You can use the PEM web interface to define one or more templates on each console. A template contains a predefined set of server options that determine the configuration of a database cluster. A template can simplify creation of clusters that use a common configuration, or limit user access to costly resources such as large server classes.

You must be an administrative user to create a template.

An administrative user can specify that a user is a *template only* user. A template only user must adhere to cluster definitions provided in a template when deploying an Ark cluster. A template only user will have access to only those templates that specify a role or tenant in which they have membership in the Select Roles (on Amazon or Azure) or Select Tenants (on OpenStack) section of the Add Template dialog.

If a user is specified as a Template Only user:

- They must use a template when deploying a cluster.
- They will be restricted to the scaling policies defined in the template.
- They cannot modify a manually-defined cluster created by another user.
- They can only create clusters in a server class that exists in an available template.
- They must use a template when cloning or restoring from backup.
- They may only delete backups of template created clusters.
- They may not delete last backup of a template created cluster if the cluster had been deleted (removing the last artifact of any cluster).

Figure 2.20 - The General tab.

To define a new template, right click on the `Templates` node, and select `Create then Templates...` from the context menu.

Use fields on the `General` tab to specify general behavioral preferences for the template:

- Provide a user-friendly name for the template in the `Template name` field.
- Use the `Description` field to provide a description of the template.
- Use the drop-down listbox in the `Engine version` field to select the version of the Postgres engine that you wish to use on clusters configured by the template.
- Use the drop-down listbox in the `Server class` field to specify the size of each cluster node. The server class determines the size and type (compute power and RAM) of any cluster configured by the template.
- If your cluster resides on an OpenStack host, use the drop-down listbox in the `Virtual network` field to specify the identity of the network in which clusters configured by the template should reside.

- If your cluster resides on an Amazon AMI, use the drop-down listbox in the `VPC` field to specify the identity of the network in which clusters configured by the template should reside.
- If your cluster resides on an OpenStack host, use the drop-down listbox in the `Floating IP pool` field to select the address pool in which clusters configured by the template should reside.
- Use the drop-down listbox in the `Number of nodes` field to specify the number of nodes that should be created in each cluster.
- Use the `Storage GB` field to specify the initial size of the data space (in Gigabytes).
- Set the `Disabled` slider to `Yes` to indicate that the template is disabled.

Figure 2.21 - The Backup Retention tab.

Use fields on the Backup Retention tab to specify the backup policies that will be applied for the template:

- Use the Number of automatic backups to retain field to specify the number of backups that will be retained for the cluster. When the specified number of server backups is reached, EDB Ark will delete the oldest backup to make room for a new backup.
- Use the Backup Window drop-down listbox to specify a time that it is convenient to perform a cluster backup.
- Set the Continuous Archiving (Point-in-Time Recovery) slider to Yes to enable point-in-time recovery for the cluster. When enabled, a base backup is automatically performed that can to be used to restore to a specific point in time. All subsequent automatic scheduled backups will also support point-in-time recovery.

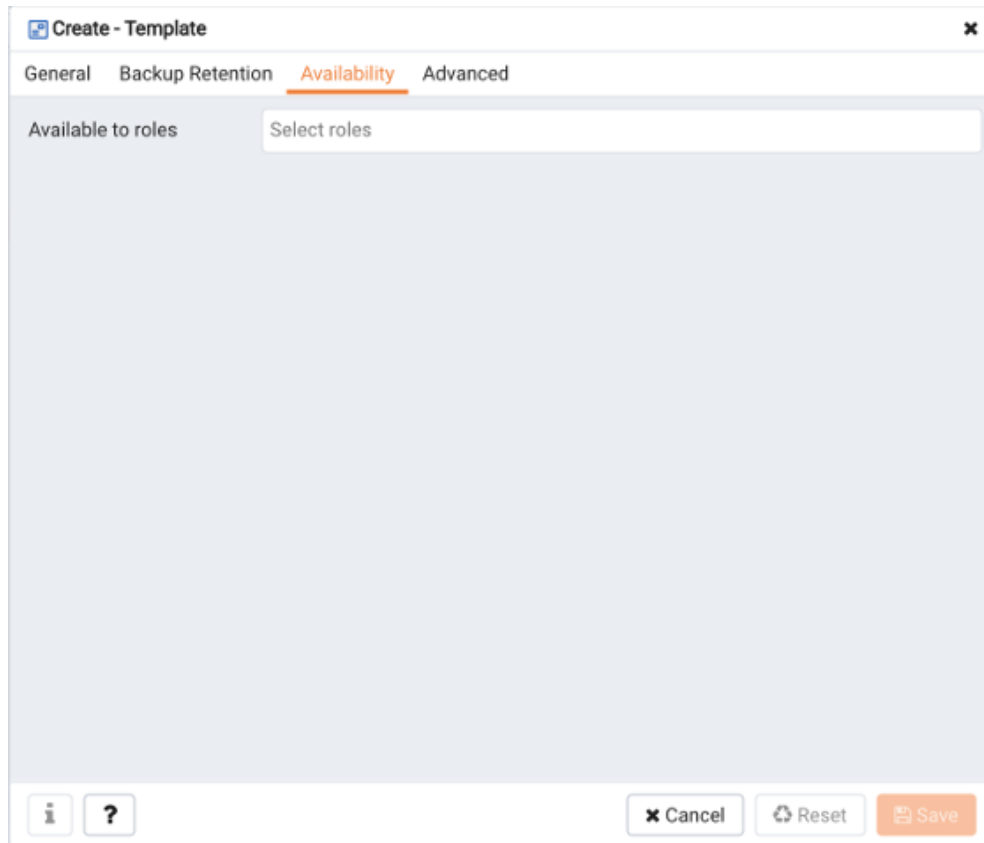


Figure 2.22 - The Availability Tab.

Use the `Availability` tab to specify the users that will have access to the template:

- Specify one or more tenants, roles, or groups that will be allowed to use the template in the `Available to tenant/role/group` field.

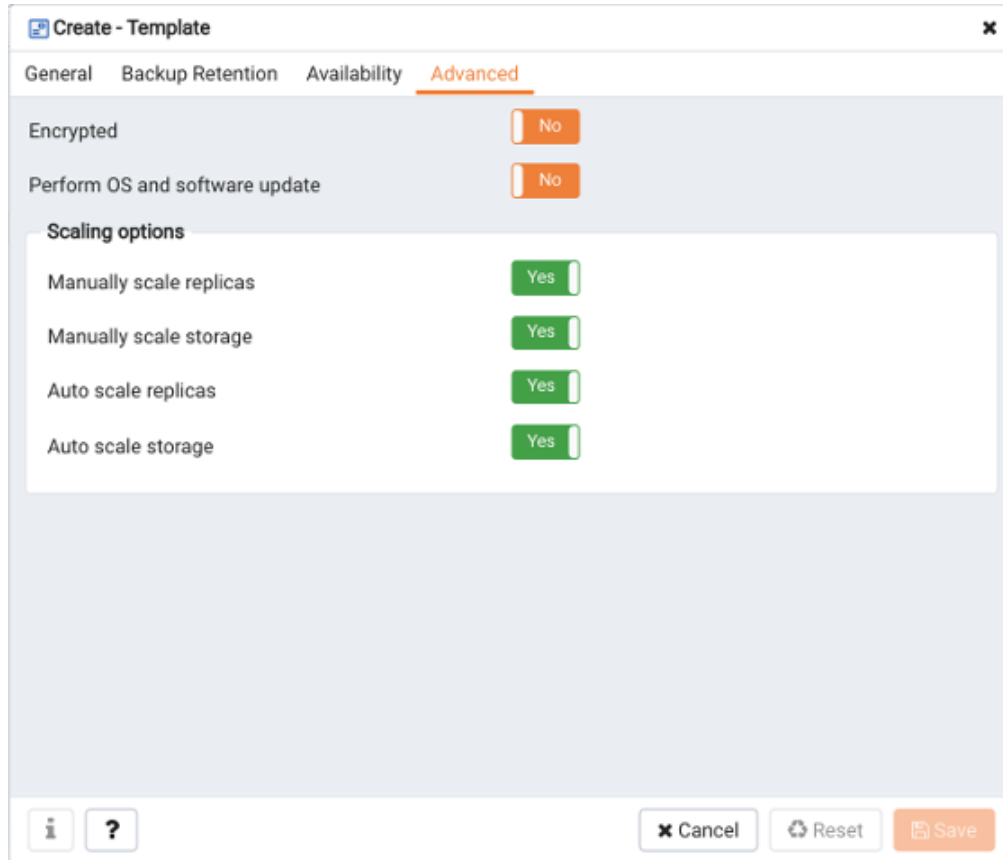


Figure 2.23 - The Advanced tab.

Use fields on the `Advanced` tab to specify additional template policies:

- Set the `Encrypted` slider to `Yes` to indicate that the cluster should be encrypted. EDB Ark uses the `aes-xts-plain` (512-bit) cipher suite to provide an encryption environment that is both secure and transparent to connecting clients. When encryption is enabled, everything residing on the cluster is encrypted except for the root filesystem.
- Set the `Perform OS and Software update` slider to `Yes` to specify that a software update should be performed whenever the cluster is provisioned. Please note: this option is disabled if the cluster uses a statically provisioned server.

Use the `Scaling Options` sliders to indicate which options will be available to template users:

- Set the `Manually scale replicas` slider to `Yes` to specify that users of this template will be allowed to manually scale replica nodes configured by this template.

- Set the `Manually scale storage` slider to `Yes` to specify that users of this template will be allowed to manually scale storage on clusters configured by this template.
- Set the `Auto scale replicas` slider to `Yes` to specify that users of this template will be able to configure automatic node scaling for clusters configured by this template.
- Set the `Auto scale storage` slider to `Yes` to specify that users of this template will be able to configure automatic storage scaling for clusters configured by this template.

2.3.4.1 Managing Templates

To view a list of the currently defined templates, select the `Templates` node in the PEM Browser tree control and navigate to the `Properties` tab.

Template name	Enabled	Engine version	Server class	Available to
Ark_user_temp_1	True	EDB Postgres Advanced Server 11 64bit on CentOS 6/7, RHEL 7	t2.medium	ArkSecondUser

Figure 2.24 - The Template Properties tab.

The `Properties` tab displays detailed information about the currently defined templates:

- The `Template name` column displays the user-friendly name of the template.
- The `Enabled` column displays `True` if the template is available for use.
- The `Engine version` column displays the engine identifier.
- The `Server class` column displays the server class available to cluster deployed with the template.
- The `Available to` column displays the names of the groups/roles/tenants that have access to the template.

You can use options on the template's context menu to manage the template. To access the context menu, right-click on the template name in the PEM Browser tree control. Then, select from the menu options:

Option	Action
Refresh	Select Refresh to update the definition of the template.
Delete/Drop	Select Delete/Drop to delete the template. PEM will confirm that you wish to delete the template before removing it.
Enable template	Select Enable template to make the template available for use when deploying a cluster.
Disable template	Select Disable template to make the template unavailable for use.
Properties...	Select Properties... to review or modify template properties.

Enabling or Disabling a Template

Enabled templates are available for use when provisioning clusters; disabled templates are not available. There are two ways to enable or disable a template - you can:

- Right-click on the engine name in the PEM browser tree control and select `Enable template` or `Disable template` from the context menu.
- Use the `Enabled` slider on the `template Properties` dialog to control the state of the template. To open the `Properties` dialog, right-click on the template name and select `Properties...` from the context menu:

Please note that disabling a template has no impact on any running clusters; it simply prevents users from creating new clusters with the template.

Modifying a Template

To modify a template definition, right-click on the template name in the PEM browser tree control, and select `Properties...` from the context menu. The template definition dialog will open, allowing you to modify the template properties. Please note that you cannot modify the `Template ID` field; the `Template ID` is a system-assigned identifier.

After modifying the template definition, click `Save` to preserve your changes and close the dialog.

Deleting a Template

To delete a template, right-click on the template name in the PEM browser tree control, and select `Delete/Drop` from the context menu. The PEM server will prompt you to confirm that you wish to delete the template before the template is deleted.

2.3.5 User Management

If allowed by the authentication model supported by your console, you can use the PEM web interface to define an Ark console user. For more information about authentication models supported by Ark, please consult the EDB Postgres Ark Administrative User's Guide, available via the `Dashboard` tab.

To define a new user, right-click on the `Users` node and select `Create`, then `User...` from the context menu.

Figure 2.25 - The Create - Users dialog.

Use fields on the `General` tab to provide information about the user:

- Use the `Login ID` field to provide the identifier that will be provided when the user logs in to the Ark server.
- Use the `Email` field to provide a notification email for the user.
- Use the `First name` field to provide the user's first name.
- Use the `Last name` field to provide the user's last name.

- Set the `Enabled` slider to `Yes` to indicate that the user account is active, and to allow the user to login.
- Set the `Admin` slider to `Yes` to specify that the user will have administrative access to the Ark console.
- Set the `Templates only` slider to `Yes` to require the user to use a template when defining a new cluster.
- Use the `Password` field to provide a password for the user.
- Use the `Confirm password` field to confirm the spelling of the user's password.

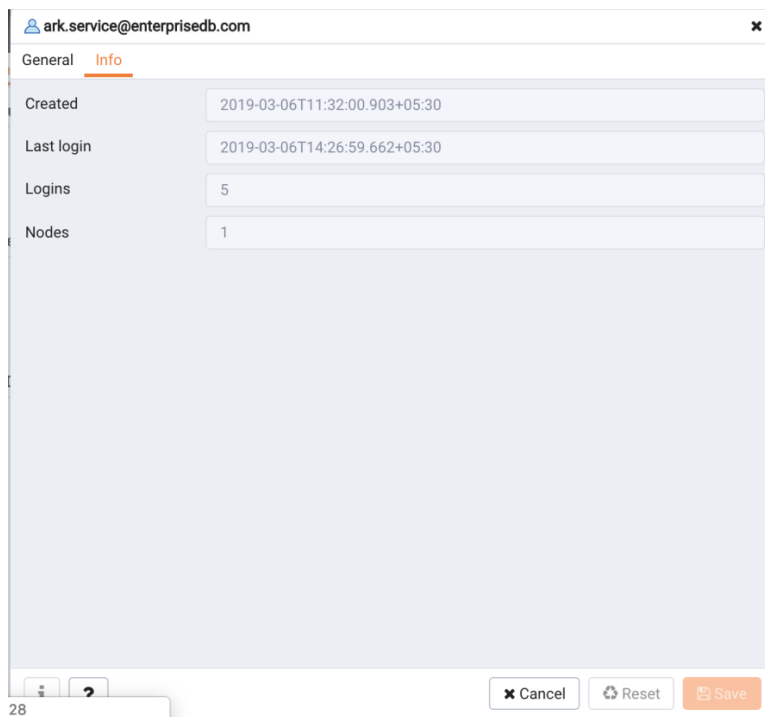


Figure 2.26 - The User Info tab.

The `Info` tab provides information about console usage by the selected user:

- The `Created` field displays the date and time that the user account was created.
- The `Last login` field displays the date and time of the most recent login by the user.
- The `Logins` field displays a count of the number of logins by the user.
- The `Nodes` field displays the number of nodes created by the user.

Deleting a User Account

To delete a user account, right-click on the name of a user, and select `Delete/Drop` from the context menu. PEM will prompt you to confirm that you wish to delete the user before the account is deleted.

2.4 Creating and Managing a Cluster

There are multiple ways to create an Ark cluster; you can:

- define a new cluster manually.
- clone an existing cluster.
- restore a backup into a new cluster.

Applying a template during cluster creation simplifies the process by providing many of the cluster creation properties.

2.4.1 Using the Create Cluster Wizard

The PEM Create cluster wizard simplifies the process of creating an Ark cluster by presenting a series of dialogs that define a cluster. The dialog content presented by the wizard will vary by platform and the steps involved in cluster creation will be different for template users than for those users defining a console manually.

To open the wizard, connect to the hosting server and expand the PEM browser tree control. Then, right-click on the `Clusters` node and navigate through the `Create` menu option and select `Cluster...`

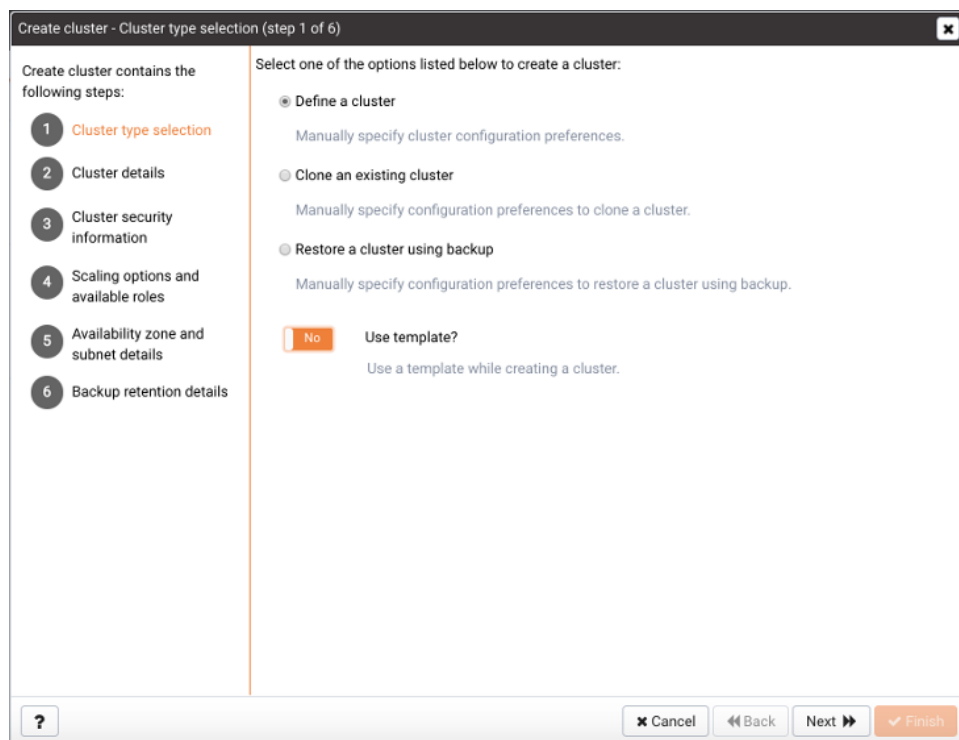


Figure 2.27 - The Cluster type selection dialog.

Use fields on the `Cluster type selection` dialog to specify how you will define the cluster:

- Select `Define a cluster` to provide detailed cluster configuration preferences.
- Select `Clone a cluster` to make a modified copy of the master node of an existing cluster.
- Select `Restore a cluster using backup` to clone a backup into the master node of a cluster. When you clone from backup, the server configuration will

match the original configuration, but the server addresses will change. When the cluster creation completes, you may need to modify parameters in the `postgresql.conf` file on the restored cluster to reflect the available memory of the new instance if the server class has changed from the original setting. After modifying the server configuration, restart the server for the changes to take effect.

A template contains a predefined set of server options that determine the configuration of a database cluster. A template can simplify creation of clusters that use a common configuration, or limit user access to costly resources such as large server classes. An administrative user can specify that a user is a *template only* user; a template only user must adhere to cluster definitions provided in a template when deploying an Ark cluster.

- Move the `Use template?` slider to `Yes` to use a template during cluster deployment.

Click `Next` to continue.

Figure 2.28 - The Cluster details dialog.

Use fields on the `Cluster details` dialog to provide deployment details for the cluster:

- Specify a name for the new cluster in the `Name` field. Please Note: EDB Ark uses the name specified in the `Name` field to identify the cluster when performing management functions. The cluster name is also part of the instance name on the

OpenStack console; you must not modify the name in the OpenStack management console. Changing the cluster name in the OpenStack console can break key EDB Ark features (i.e. failover).

- If applicable, use the `Clone from cluster` drop-down listbox to select an existing cluster to clone.
- If applicable, use the `Backup from cluster` drop-down listbox to select a backup to restore.
- The `Backup start time` field displays the time at which the backup was taken; this field is not modifiable.
- If applicable, use the `Recovery point` date selector to specify the date and time to which you wish to recover. The specified time may not precede the timestamp of the specified backup..
- If applicable, use the `Template name` drop-down listbox to select a template that will provide deployment details for the cluster.
- Use the drop-down listbox in the `Engine version` field to select the version of the Postgres engine that you wish to use.
- Use the drop-down listbox in the `Server class` field to specify the size of each cluster node. The server class determines the size and type (compute power and RAM) of each node within the cluster. You can adjust the amount of storage used by the cluster, or number of replicas in the cluster as your resource demands change. For example, you can start with a `m1.small` instance, and later, easily upgrade to a more capable `c1.medium` instance as your performance requirements dictate.
- To create a cluster on an Amazon host that resides in a private cloud, move the `Use Private IP addresses` slider to `Yes`; then use the `Virtual network` field to select the virtual private cloud into which the cluster will be deployed.
- If your cluster resides on an Azure host, use the drop-down listbox in the `Virtual network` field to specify the network in which the cluster will reside.
- If your cluster resides on an OpenStack host, use the `Virtual network` drop-down listbox to select the network in which the cluster will be deployed; then, use the `Floating IP pool` drop-down listbox to select an address pool in which the cluster will reside.

- Use the `Number of nodes` field to specify the number of server nodes that you wish to create. The name specified in the `Name` field will apply to the master node; each additional node will act as a replication server for the master node.
- Use the `Storage GB` field to specify the initial amount of the data space (in Gigabytes) available for the cluster.
- Set the `Encrypted` slider to `Yes` to indicate that the cluster should be encrypted. EDB Ark uses the `aes-xts-plain (512-bit)` cipher suite to provide an encryption environment that is both secure and transparent to connecting clients. When encryption is enabled, everything residing on the cluster is encrypted except for the root filesystem.
- If your cluster resides on an AWS host, set the `EBS Optimized` slider to `Yes` to specify that your cluster should use an Amazon EBS-optimized instance and provisioned IOPS to guarantee a level of I/O performance. The `IOPS` field is enabled for those clusters that will reside on an EBS-optimized instance. Use the `IOPS` field to specify the level of I/O performance that will be maintained for the cluster by automatic scaling. The maximum value is 30 times the size of your cluster; for example, if you have a 4 Gigabyte cluster, you can specify a maximum value of 120.
- Set the `Perform OS and software update` slider to `Yes` to instruct the server to perform a software update when the cluster is provisioned. Please note: this option is disabled if the cluster uses a statically provisioned server.

Click `Next` to continue.

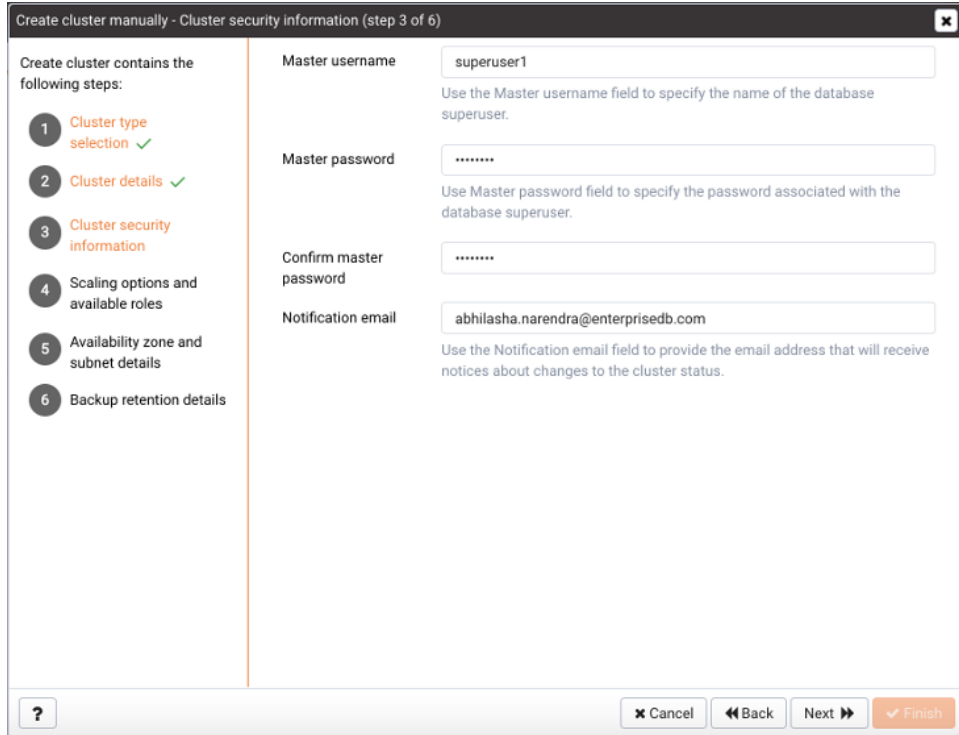


Figure 2.29 - The Cluster security information dialog.

Use fields on the Cluster security information dialog to provide connection and notification details for the cluster:

- Enter the name of the database superuser in the Master username field.
- Enter the password associated with the database superuser in the Master password field.
- Confirm the password provided in the Master password field in the Confirm master password field.
- Use the Notification email field to provide the email address that will receive notices about changes to the cluster status.

Click Next to continue.

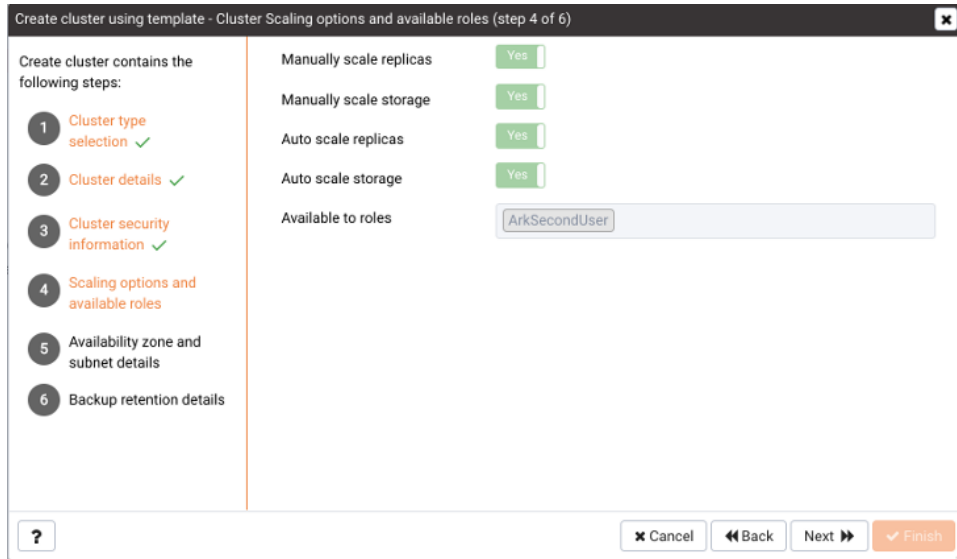


Figure 2.30 - The Cluster Scaling options and available roles dialog.

If you are using a template to define the cluster, the `Scaling options and available roles` dialog displays the scaling configuration of the new cluster as designated on the template. The information displayed on the dialog is not modifiable:

- If the `Manually scale replicas` slider is set to `Yes`, manual replica scaling will be enabled for the new cluster; cluster users will be able to manually add replicas to the cluster.
- If the `Manually scale storage` slider is set to `Yes`, manual scaling of storage will be enabled for the new cluster; cluster users will be able to manually add storage to the cluster when needed.
- If the `Auto scale replicas` slider is set to `Yes`, automatic replica scaling will be enabled for the cluster; the cluster will automatically add replicas when cluster usage reaches the threshold specified in the template.
- If the `Auto scale storage` slider is set to `Yes`, automatic scaling of storage available to the cluster will be enabled; the cluster will automatically add storage when the cluster reaches the threshold specified in the template.
- The `Available to tenant/role/group` field displays the names of the platform-specific user groups that can access the new cluster.

Click `Next` to continue.

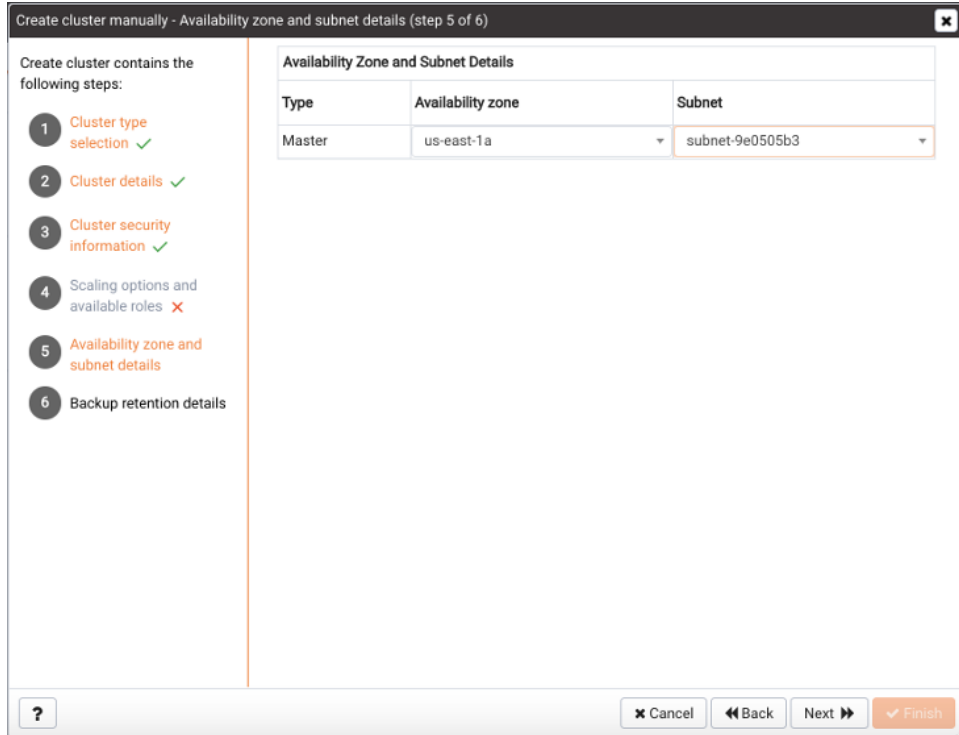


Figure 2.31 - The Availability zone and subnet details dialog.

Use fields on the Availability zone and subnet details dialog to specify an availability zone/subnet pair for each node of the cluster:

- Use the drop-down listbox in the Availability zone column to specify the data center in which the node listed in the Type column will be deployed, or accept the default to allow the host to select an availability zone.
- Use the drop-down listbox in the Subnet field to specify the subnet that will be used by the node listed in the Type column, or accept the default to allow the host to create a new subnet for the node.

Click Next to continue.

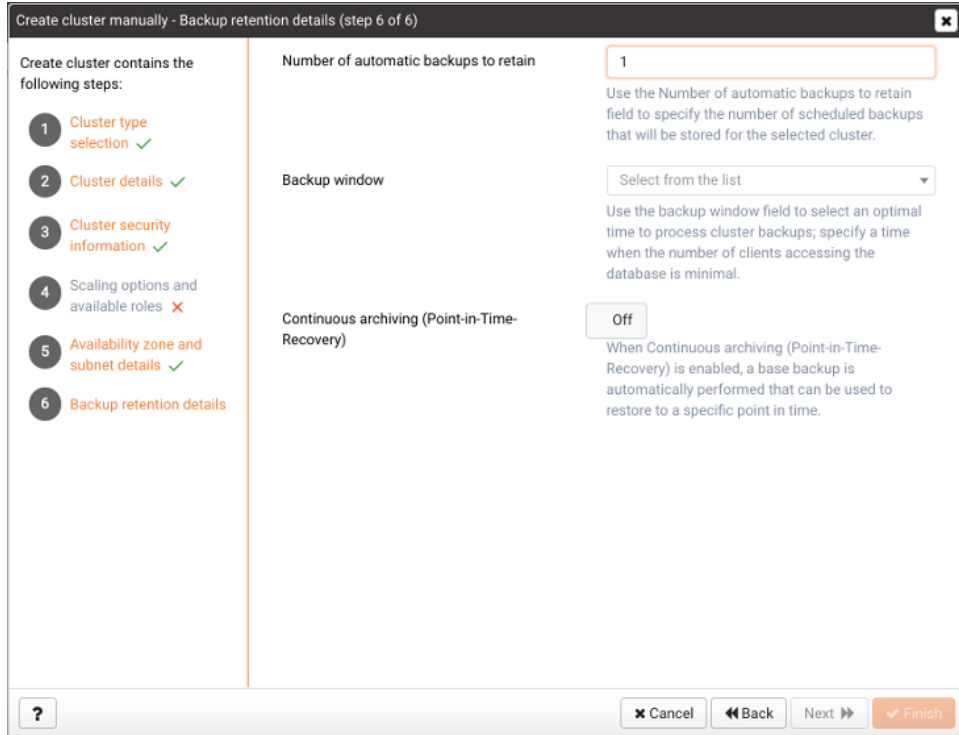


Figure 2.32 - The Backup retention details dialog.

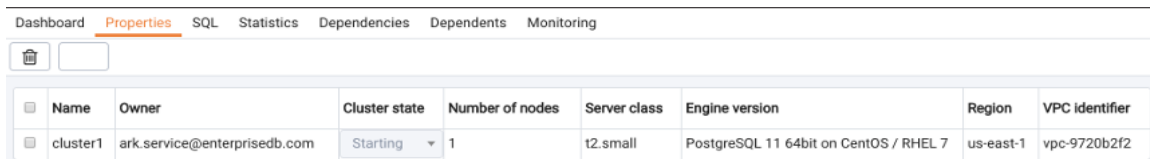
Use fields on the Backup retention details dialog to specify your backup preferences for the cluster:

- Use the Number of automatic backups to retain field to specify the number of backups that will be stored. When the specified number of backups is reached, the server will delete the oldest backup to make room for a new backup.
- Use the Backup window drop-down listbox to select a time that it is convenient to backup the server (you may wish to schedule backups to occur when the CPU load is the lightest).
- Use the Continuous archiving (Point-in-Time Recovery) slider to enable point-in-time recovery for the cluster. When enabled, a base backup is automatically performed that can to be used to restore to a specific point in time. All subsequent automatic scheduled backups will also support point-in-time recovery. Note that if you deselect this option, the cluster (and subsequent automatic backups) will be re-configured to not include support for point-in-time recovery.

Use the Back button to review or modify previous selections, or select Finish to create the defined cluster. Choose Cancel to exit without preserving your selections.

2.4.1.1 Reviewing a List of Existing Clusters

To review a list of existing clusters that reside on a specific cluster host, highlight the `Clusters` node beneath the cluster name in the PEM client tree control, and navigate to the `Properties` tab.



Name	Owner	Cluster state	Number of nodes	Server class	Engine version	Region	VPC identifier
cluster1	ark.service@enterprisedb.com	Starting	1	t2.small	PostgreSQL 11 64bit on CentOS / RHEL 7	us-east-1	vpc-9720b2f2

Figure 2.33 - The clusters list on the Properties tab.

The `Properties` tab displays:

- The name of the cluster in the `Name` column.
- The name of the cluster owner in the `Owner` column.
- The current state of the cluster in the `Cluster state` column.
- The number of cluster nodes in the `Number of nodes` column.
- The server class that is being used by the cluster in the `Server class` column.
- The engine version that is deployed on the cluster in the `Engine version` column.
- The region in which the cluster is deployed in the `Region` column.
- The VPC identifier of the VPC that is in use by the cluster in the `VPC identifier` column.

2.4.2 Managing a Cluster

To view statistical information and log entries for a cluster, select the cluster name in the PEM Browser tree control and navigate to the `Dashboard` tab.

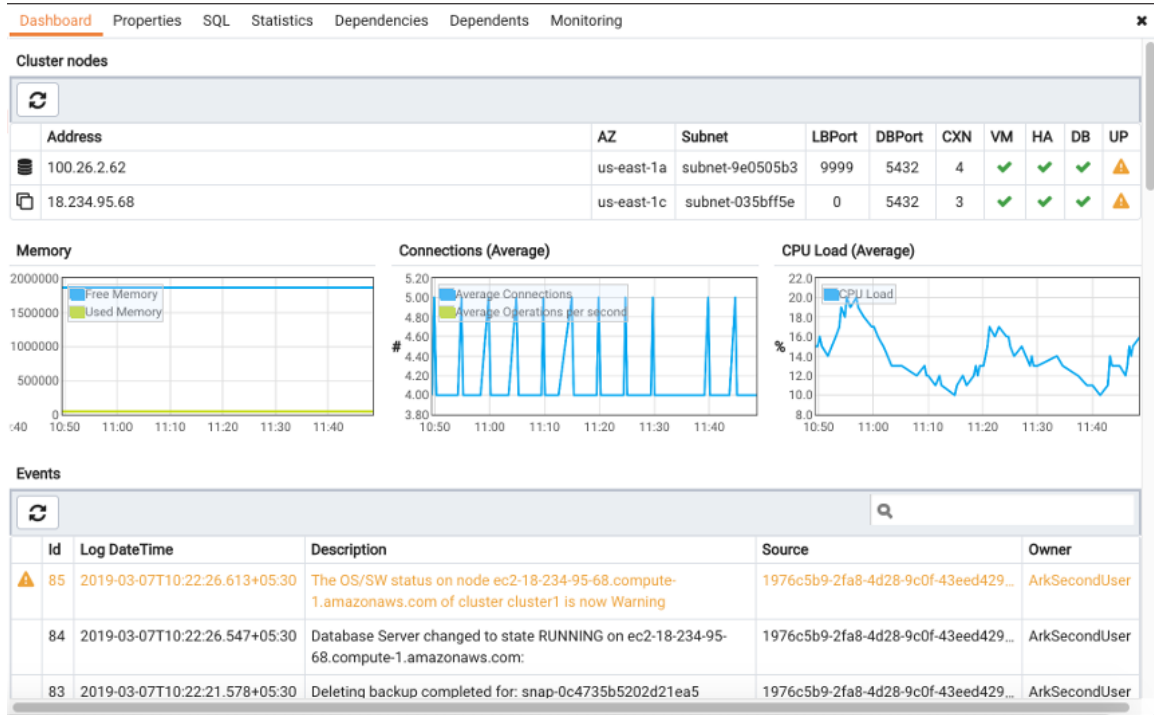


Figure 2.34 - The Cluster Dashboard tab.

The `Cluster nodes` table displays each node of the cluster, and the current state of the node.

- Use the refresh icon in the upper-left corner of the panel header to refresh the displayed content; the most-recent entries in the cluster log file are displayed nearest to the top of the panel.
- Use the search box in the upper-right corner of the panel header to search the table content for key words or strings; separate string values with a pipe character (`|`) to include entries that match more than one search term.

Icons in the left most column indicate the role of the node within the cluster; a disc icon indicates that the node is a master node within the scenario, and a copy symbol indicates that a node is currently acting as a replica.

The following columns display connection information about each node:

- The `Address` column displays the IP address of the node.
- The `AZ` column displays the Availability Zone in which the node resides.
- The `Subnet` column displays the subnet ID on which the node resides.
- The `LBPoRT` column displays the port number to which a client application should connect to utilize load balancing.
- The `DBPoRT` column displays the default listener port for the Advanced Server or PostgreSQL server.
- The `CXN` column displays the current number of connections to the node.

The following columns display the health of the node:

- The `VM` column displays the state of the virtual machine on which the cluster resides.
- The `HA` column displays the state of the high-availability cluster.
- The `DB` column displays the state of the database server.
- The `UP` column displays the current status of the packages installed on the cluster. Periodically, the cluster manager performs a check to see if the packages are up to date.

Indicators display the health of the node:

- A green checkmark indicates that an object is healthy.
- A yellow alert symbol calls attention to an object that requires attention.
- A red error symbol signifies that an object is not available.
- An hourglass signals that the cluster is processing a request.
- A question mark indicates that the state of the resource is unknown.

Graphs display usage information about the cluster:

- The memory usage of the cluster.
- The average number of connections made to the cluster.

- The average CPU Load of the cluster.

The **Events** panel displays events that are logged for the current cluster.

- Use the refresh icon in the upper-left corner of the panel header to refresh the displayed content; the most-recent entries in the cluster log file are displayed nearest to the top of the panel.
- Use the search box in the upper-right corner of the panel header to search the log file content for key words or strings; please note that the search is not context-sensitive.

You can use options on the cluster's context menu to manage the cluster. To access the context menu, right-click on the cluster name in the PEM Browser tree control. Then, to manage the cluster, select from the menu options:

Option	Action
Refresh	Select Refresh to update the definition of the server or server objects that reside on the console.
Delete/Drop	Select Delete/Drop to delete the cluster. PEM will confirm that you wish to delete the cluster before removing it.
Download SSH Key	Select Download SSH Key to download the SSH key for the cluster.
Backup...	Select Backup... to take a backup of the cluster.
Scale Up	Select Scale Up to add replicas or storage to the selected cluster.
Scale Down	Select Scale Down to remove replicas from the selected cluster.
Scale Machine Type	Select Scale Machine Type to change the size of the virtual machine on which the cluster resides.
Upgrade OS/Software	Select Upgrade OS/Software to perform a yum update on each node of the cluster.
Properties...	Select Properties... to review or modify cluster properties.

To view information about the cluster configuration, select the cluster name in the PEM Browser tree control and navigate to the `Properties` tab.

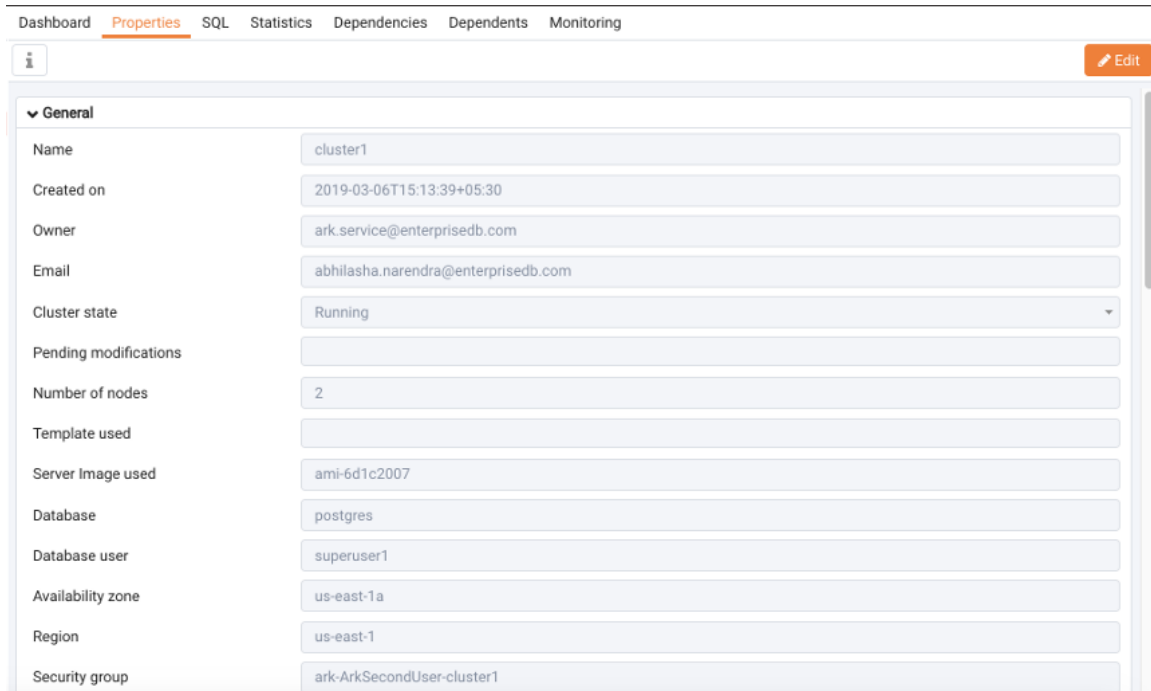


Figure 2.35 - The Cluster Properties tab.

Panels on the `Properties` tab display deployment details for the cluster.

- The `General` panel displays deployment details such as owner information, creation date, and current cluster state.
- The `Storage` panel displays information about the storage allocated for the cluster, the amount of storage used, and the amount of storage still available.
- The `Network` panel displays information about the network addresses and ports in use by the cluster.
- The `Maintenance` panel displays information about the monitoring and healing policies enforced on the cluster.
- The `Auto Scale` panel displays information about the scaling policies currently in use by the cluster; if automatic scaling is enabled, the cluster will be scaled up when the thresholds specified in the `% of storage size used` or `# of server connections` fields are reached.
- The `Backup` panel displays the backup schedule and retention policy enforced on the cluster.

2.4.2.1 Connecting to a Cluster Instance

To download the cluster's SSH key, select `Download SSH Key` from the console's context menu. An informational popup will provide information about connecting to the cluster instance:

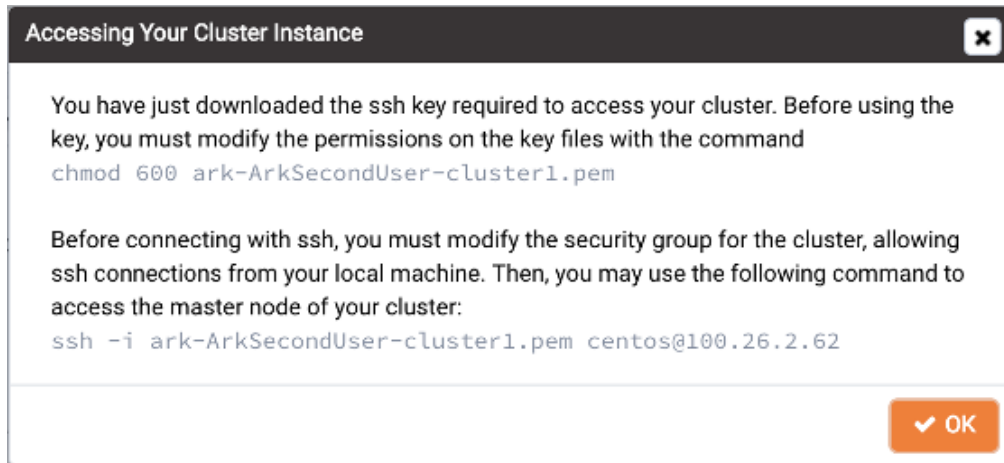


Figure 2.36 - Downloading an SSH Key.

Before using SSH to connect to an instance, an administrative user must modify the security group used by the cluster to allow SSH connections. For more information, please see the EDB Ark Administrative User's Guide, available via the Ark server's Dashboard tab.

2.4.2.2 Taking a Cluster Backup

To take a backup of the currently selected cluster, select `Backup...` from the console's context menu or the PEM web interface's `Object` menu. A popup will open, prompting you to provide a comment so you can identify the backup later should you need to use the backup to restore or clone the cluster:

Figure 2.37 - Taking a cluster backup.

Provide a comment in the `Notes for the backup` field, and click `Backup` to start the backup. An informational popup will confirm that the backup process has started.

Managing Backups

To view a list of cluster backups, select the `Backups` node in the PEM Browser tree control and navigate to the `Properties` tab.

ID	Cluster	Notes	Engine version	Capacity (in GB)	Backup type	Start time
snap-0d5c51c540f378bde	cluster1	Scheduled backup: Thu Mar 07 02:03...	PostgreSQL 11 64bit on CentOS / RH...	2	Automatic	2019-03-07T02
snap-03325db58551526d0	cluster1		PostgreSQL 11 64bit on CentOS / RH...	2	Manual	2019-03-07T10

Figure 2.38 - The Backups Properties tab.

The `Properties` tab displays detailed information about completed backups:

- The `ID` column displays the backup identifier.
- The `Cluster` column displays the name of the cluster that was backed up.
- The `Notes` column displays the comment provided when the backup was taken.
- The `Engine version` column displays the engine version of the cluster.

- The `Capacity (in GB)` column displays the names of the groups/roles/tenants that have access to the template.
- The `Backup type` column displays `Automatic` if the backup is a scheduled backup, and `Manual` if the backup was started by a user.
- The `Start time` column displays the time at which the backup was started.
- The `End time` column displays time at which the backup completed.

You can use options on the backup's context menu to manage the backup. To access the context menu, right-click on the backup identifier in the PEM Browser tree control. Then, select from the menu options:

Option	Action
Refresh	Select Refresh to update the backup information.
Delete/Drop	Select Delete/Drop to delete the backup. PEM will confirm that you wish to delete the backup before removing it.

2.4.2.3 Scaling a Cluster

You can use the `Scale Up` and `Scale Down` menu options to add replicas or storage to your cluster, or to delete unneeded replica nodes. Access the scaling options on either the cluster's context menu or on the PEM web interface's `Object` menu.

Type	Availability zone	Subnet
Master	us-east-1a	subnet-9e0505b3
Replica	Auto Availability Zone	Create New Subnet

Figure 2.39 - Adding replicas to a cluster.

To add one or more replicas to a cluster, select `Scale Up`, then `Replicas`.

When the `Scale Up` dialog opens, specify the number of replicas that you would like to add to the cluster. If applicable, select a subnet and an availability zone for each replica, and click `Ok`.

The server will confirm that you wish to add the specified number of replicas; click `Yes` to continue. An informational popup will confirm when the replica(s) are being added.

To remove one or more replicas from a cluster, select `Scale Down`.

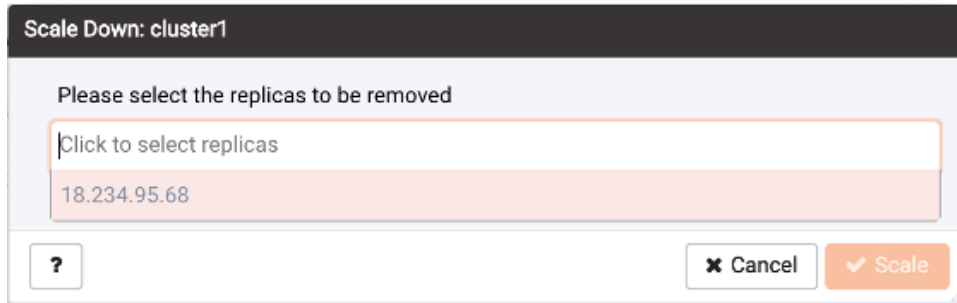


Figure 2.40 - Removing replicas from a cluster.

When the `Scale Down` dialog opens, click in the field and select a replica address from the list of replicas. To select more than one replica for removal, click in the field again, and select an additional address. When you've finished selecting replicas for removal, click `OK`. The server will confirm that you wish to remove the replicas; click `OK` to continue. An informational popup will confirm when the replica(s) are being removed.

To add more storage space to a cluster, select `Scale Up`, then `Storage`.




Figure 2.41 - Adding storage space to a cluster.

When the `Scale Up` dialog opens, specify the amount of storage (in Gigabytes) that you would like to add to the cluster, and click `OK`. The server will confirm that you wish to add the specified storage; click `Yes` to continue. An informational popup will confirm when the storage is being added.

2.4.2.4 Scaling a Machine Type

To modify the instance type on which a cluster resides, select `Scale Machine Type` from the console's context menu or from the PEM web interface's `Object` menu.

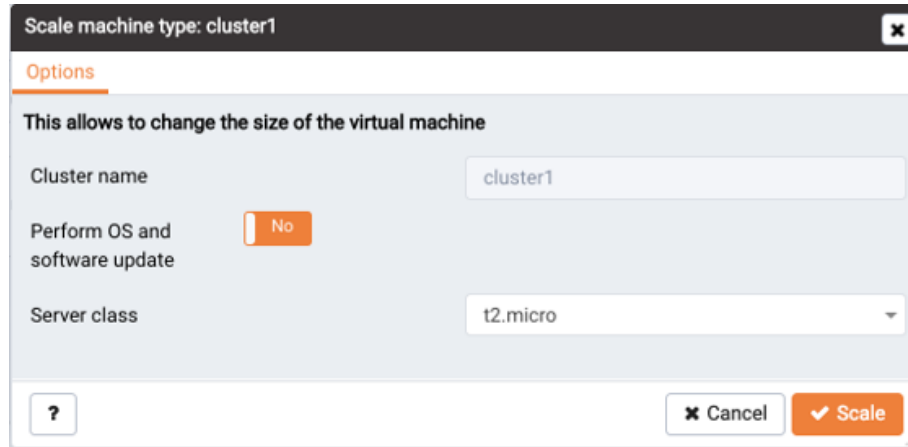


Figure 2.42 - Scaling a machine type.

Use fields on the `Scale machine type` dialog to specify your scaling preferences:

- The `Cluster name` field displays the name of the cluster; this field is not modifiable.
- Move the `Perform OS and software update` slider to `Yes` to perform a software update with the cluster is provisioned; this option can be time consuming.
- Use the `Server class` drop-down listbox to select the new server class for the cluster.

When you've specified your preferences, click `Scale` to continue. The server will ask you to confirm your changes; click `Continue` to start the scaling process. An informational popup confirms when the scaling begins.

2.4.2.5 Upgrading a Cluster

You can use the `Upgrade OS/Software` option to invoke a `yum update` command on each node of your cluster, updating any installed packages to the most recent version available through your specified repositories. To perform an upgrade, select `Upgrade OS/Software` from the console's context menu or from the PEM web interface's `Object` menu.

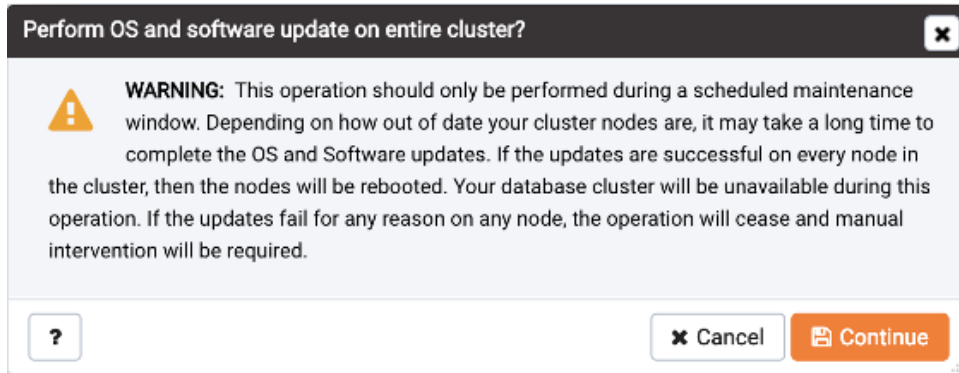


Figure 2.43 - Upgrading a Cluster.

Before starting an upgrade, PEM will prompt you to confirm that you wish to `Continue`. Please note that a software update can take some time to complete; we recommend scheduling software updates during a maintenance window. After performing a `yum update`, the cluster nodes will be rebooted (initiating any kernel updates required).