Replication Guide

Version 1.0

Revision History

Date	Version	Description
28-04-2019	1.0	Initial version

Contents

Replication Overview

Replication enables you to asynchronously replicate a source disk to a destination disk located in a different (typically remote) cluster. This is also referred to as geo-replication.

Replication is enabled by creating a replication job on the source cluster. The first time a replication job runs, it will transfer all data from the source disk to the destination disk, later jobs instances will only transfer the changes/differences done since the previous replication instance.

The system uses snapshots to provide a point-in time view of the data during replication and for the computation of differences.

Note that any data present on the destination disk will be overwritten by the source disk data, so you should not write on the destination disk.

You can setup replication between Source and Destination Clusters by applying the following steps (Refer to Administrator Guide for more detail):

Steps at Destination Cluster

1. Enable Backup/Replication Nodes

You should assign one or more nodes to run the Backup/Replication service; you can select this service during the deployment wizard

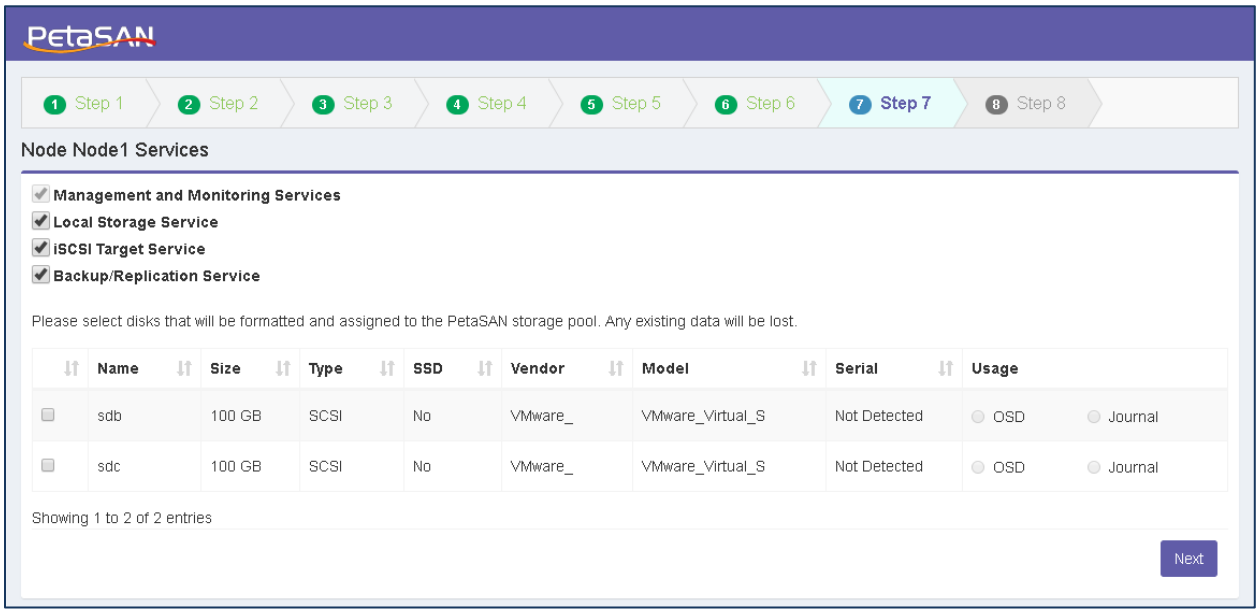

Or you can add the service after deployment using the Nodes List > Manage Roles Form.

2. Create Replication User

On the destination cluster, you need to create a replication user to be used by the system during replication and authorize this user on the pool(s) containing the destination disk(s) you wish to replicate. For a disk using EC pool, you need to give the user access to the replicated pool only.

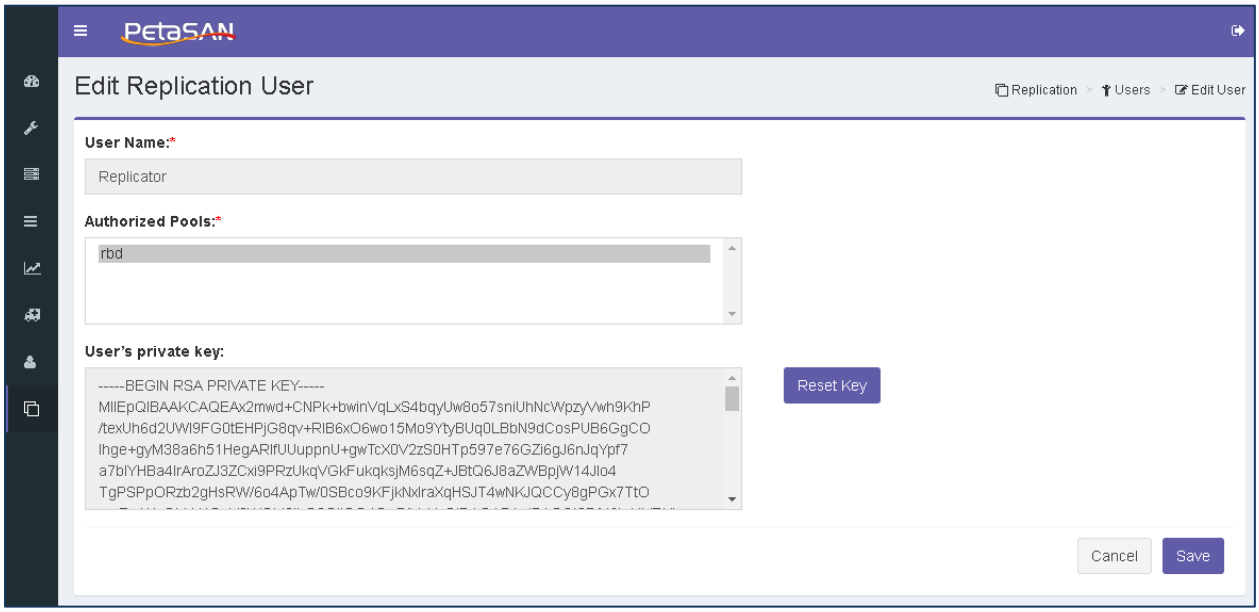

3. Create Destination Disk

For the destination disk, you can either create a new disk or use an existing one; you need to enable the replication option to flag the disk as a replication target. The disk needs to be the same size as the source.

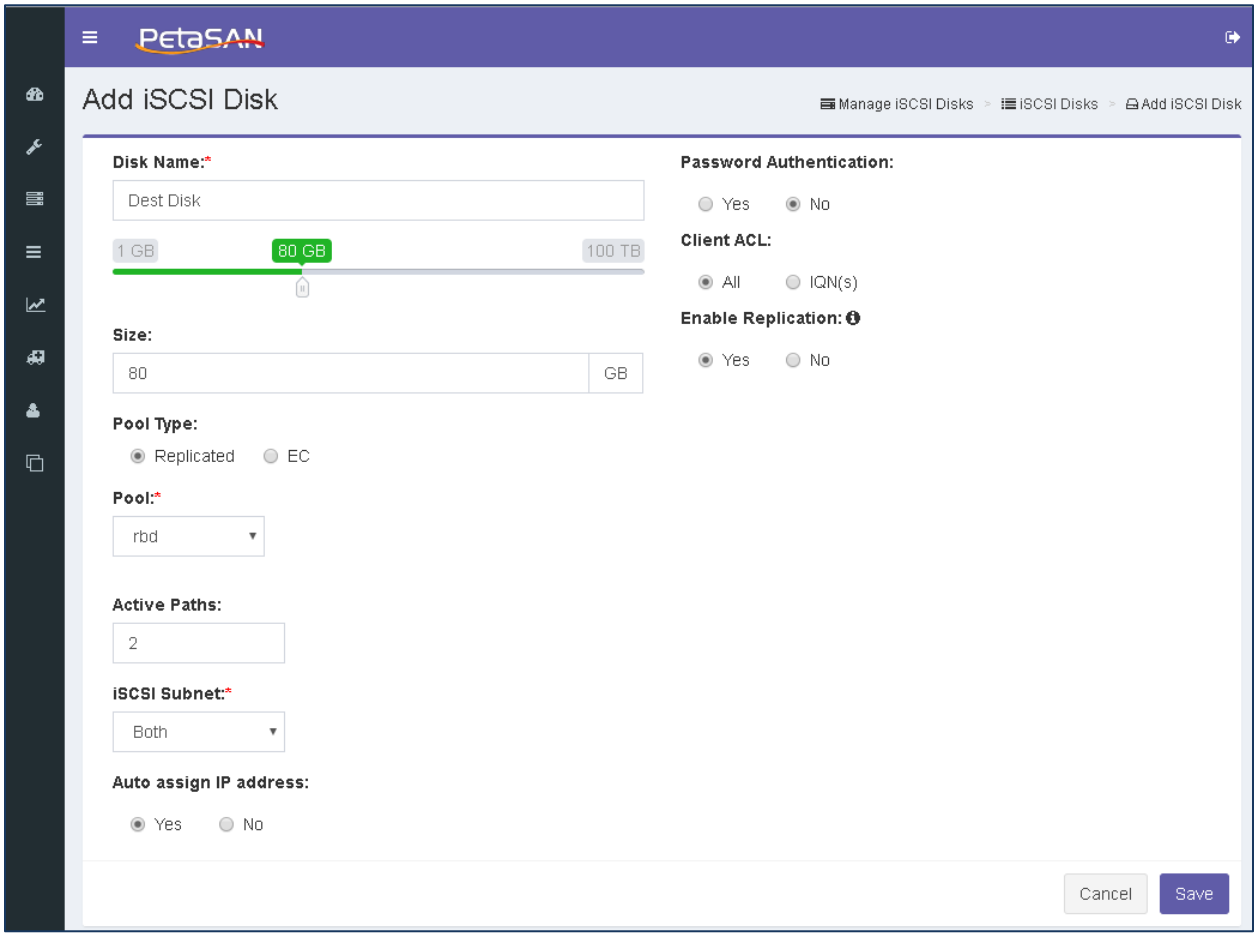

➢ *Notes:*

- *You should not write to the destination disk because data will be over-written when the next replication job is running, as the disk data will be synched from the source disk.*
- *System will stop the destination disk when its replication job runs.*

4. Provide information to Source Cluster Administrator

You need to provide the following information to the source cluster Administrator:

- Cluster Name
- Backup Node Remote IP
- Replication user name and private key
- Destination disk id

Steps at Source Cluster

1. Enable Backup/Replication Nodes

You should assign one or more nodes to run the Backup/Replication service; you can select the service during the deployment wizard

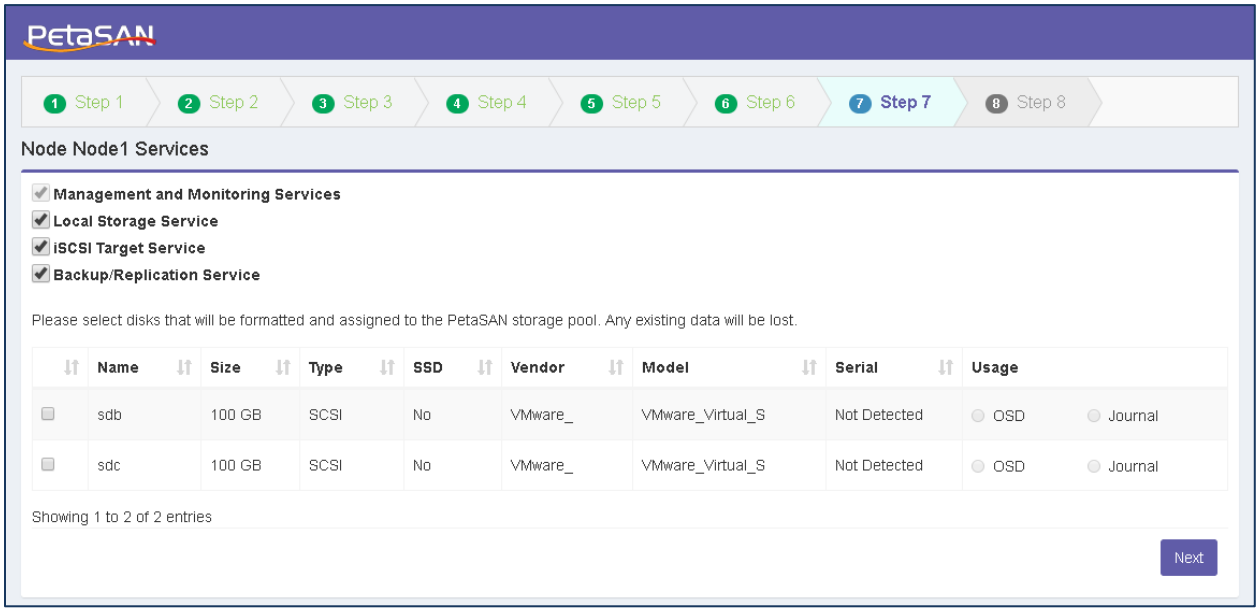

Or you can add the service after deployment using the Nodes List> Manage Roles Form.

2. Create Destination Cluster

You need to define a target destination cluster using the Cluster Name, Remote IP, Replication User Name and Private Key provided by the destination cluster administrator.

It is recommended you test the connection first before saving.

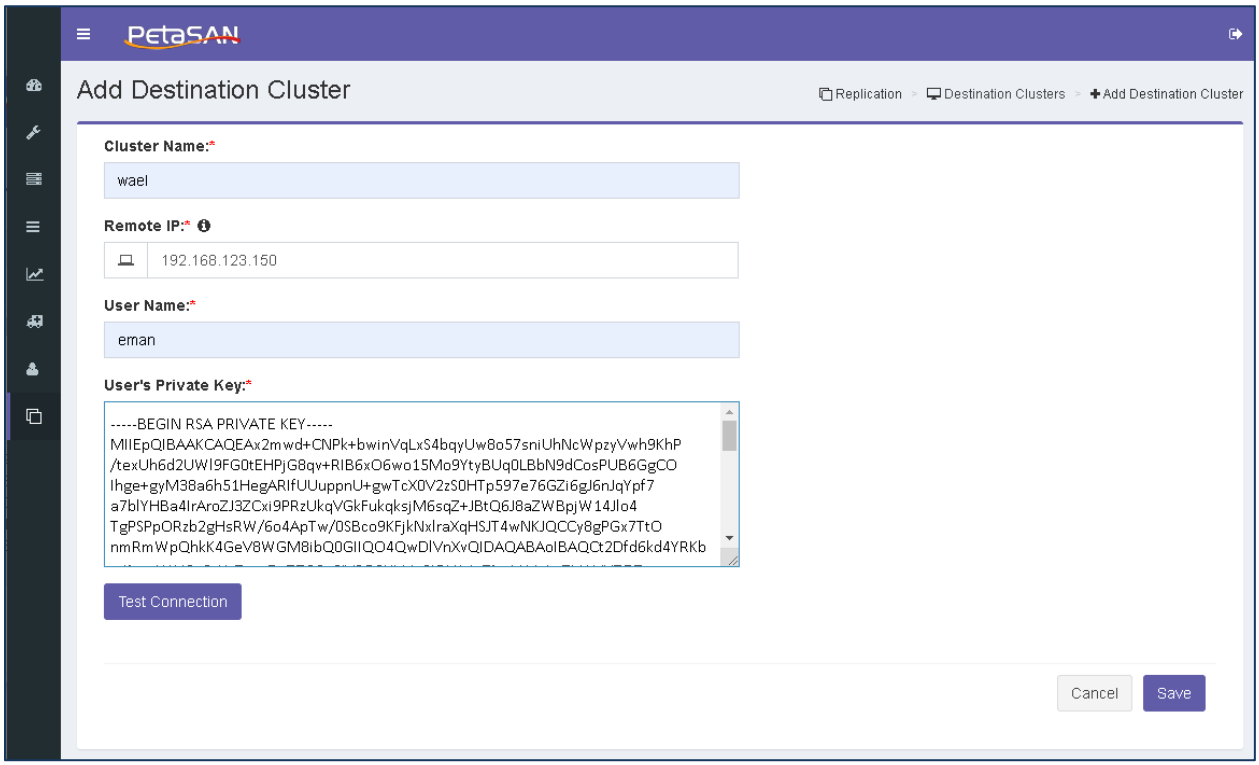

3. Create Replication Job

Create a replication job by entering the following:

- Select one of the backup/Replication nodes on the source cluster.
- Define the job schedule.
- Select the source disk and make sure it is the same size as the destination disk.
- Select the destination cluster you created.
- Select the destination disk id provided by the destination cluster administrator.
- You can choose to compress data during transfers by enabling compression and selecting the desired compression algorithm.
- In case you need to run custom scripts during a replication job, you can define external URLs to be called at specific stages of replication, such as prior to performing disk snapshots or after job completion. These could be used in more advanced setups to flush files, lock database tables or send email on job completion.

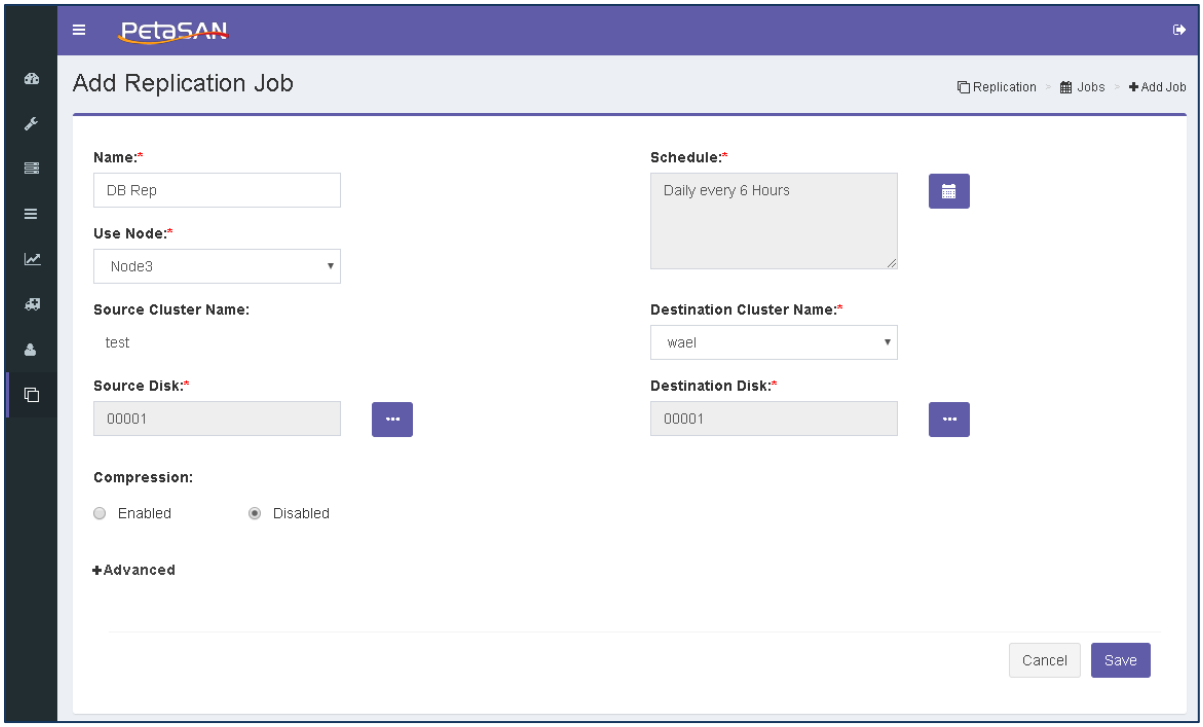

4. Monitor Active Jobs

At any time you can view the currently running jobs and monitor their progress using the Active Jobs List.

You can cancel a running job if needed using the cancel button.

If a replication job fails, you will receive an email notification (if you enabled notifications).

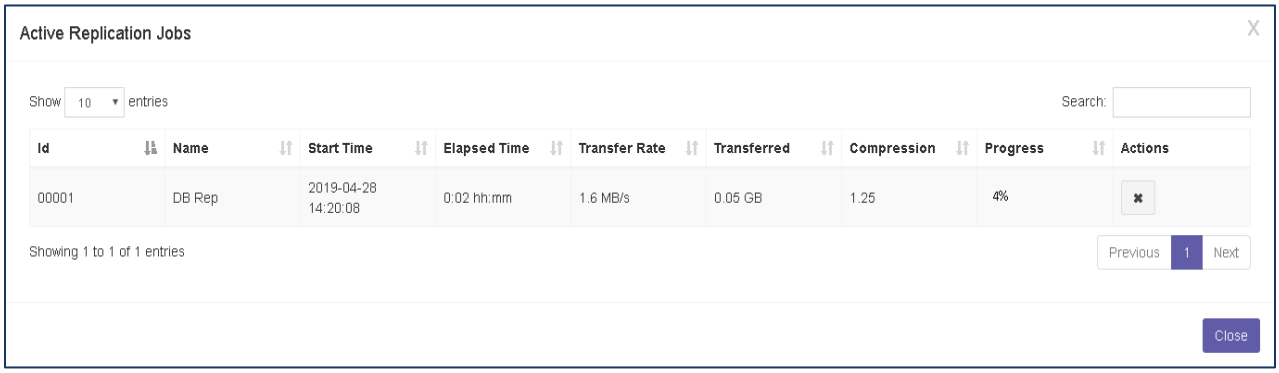

5. Manage Replication Jobs

You can stop/start replication jobs, modify schedules and view job logs.