
PostgreSQL DBA's Guide to Actifio Copy Data Management

Copyright, Trademarks, and other Legal Matter

Copyright © 2009 - 2018 Actifio, Inc. All rights reserved.

Actifio[®], AnyIT[®], Dedup Async[®], OnVault[®], Enterprise Data-as-a-Service[®], FlashScan[®], AppFlash DEVOPS Platform[®], Copy Data Cloud[®], and VDP[®] are registered trademarks of Actifio, Inc.

Actifio Sky[™], Actifio One[™], and Virtual Data Pipeline[™] are trademarks of Actifio, Inc.

All other brands, product names, goods and/or services mentioned herein are trademarks or property of their respective owners.

Actifio, Inc., is a provider of data protection and availability products. Actifio's technology is used in products sold by the company and products and services sold and offered by its commercial partners. The current list of Actifio patents is available online at: <http://www.actifio.com/patents/>

Actifio believes the information in this publication is accurate as of its publication date. Actifio reserves the right to make changes to information published in this document, including without limitation specifications and product descriptions, at any time and without notice. This document supersedes and replaces all information supplied prior to the publication hereof.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." ACTIFIO, INC. MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

This software and the associated documentation are proprietary and confidential to Actifio. Use, copying, and distribution of any Actifio software described in this publication requires an applicable software license. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

Actifio strives to produce quality documentation and welcomes your feedback. Please send comments and suggestions to docs@actifio.com.

Preface

The information presented in this guide is intended for users who are familiar with basic Actifio processes and procedures as described in ***Getting Started with Actifio Copy Data Management*** and who are qualified to administer PostgreSQL databases.

Your Actifio appliance's Documentation Library contains detailed, step-by-step, application-specific instructions on how to protect and access your data. Each guide is in PDF format and may be viewed online, downloaded, or printed on demand. The following guides will be of particular interest:

- ***Connecting Hosts to Actifio Appliances***
- ***Virtualizing and Protecting Copy Data with the Application Manager***
- ***Accessing and Recovering Copy Data with the Application Manager***
- ***Restoring Copy Data with the Application Manager***

The ActifioNOW Customer Portal

During the configuration and initialization of your Actifio appliance your Actifio representative provided you with a user name and password for the ActifioNOW customer portal.

From the customer portal you can obtain detailed reports about your Actifio appliance as well as search the portal's knowledge base for answers to specific questions.

To log into the ActifioNOW customer portal:

1. Go to: <https://now.actifio.com>
2. When prompted, enter the user name and password provided by your Actifio representative.

Actifio Support Centers

To contact an Actifio support representative, you can:

- Send email to: support@actifio.com
- Call:

From anywhere: +1.315.261.7501

US Toll-Free: +1.855.392.6810

Australia: 0011 800-16165656

Germany: 00 800-16165656

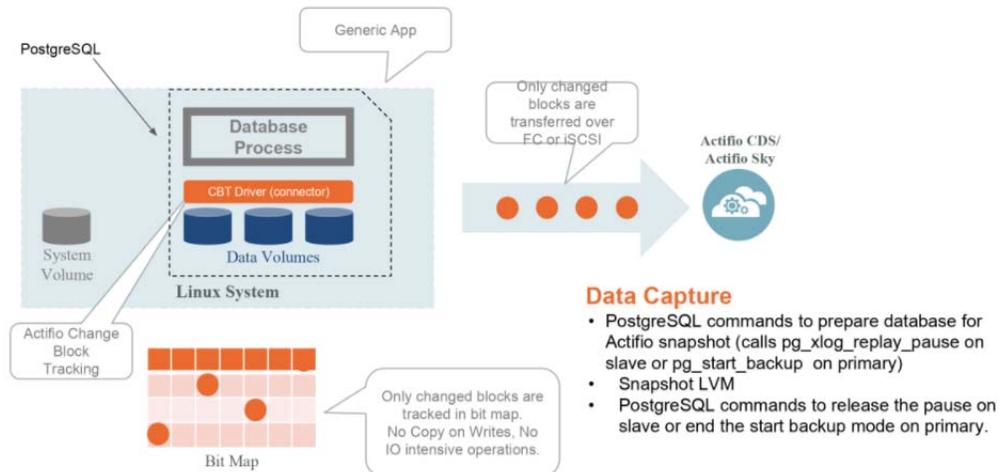
New Zealand: 00 800-16165656

UK: 0 800-0155019

Contents

Preface	v
The ActifioNOW Customer Portal.....	v
Actifio Support Centers	v
Chapter 1 - Introducing Actifio Copy Data Management for PostgreSQL Databases	1
Chapter 2 - Preparing the PostgreSQL Database	3
Chapter 3 - Protecting a PostgreSQL Database	5
Protecting a PostgreSQL Database	5
Protecting PostgreSQL Database Logs.....	7
Chapter 4 - Accessing and Recovering a PostgreSQL Database	9
Mount and Refresh Target PostgreSQL Database as a Virtual Application (Same Name or Different Name)	9
Logs Roll Forward after Mount and Refresh of Target PostgreSQL Database	12
Mount and Refresh the Target PostgreSQL Database as a Virtual Slave Application	15

1 Introducing Actifio Copy Data Management for PostgreSQL Databases



PostgreSQL with Actifio Out-of-Band Linux CBT

Operation	Benefit
Backup	Manual and/or scheduled online backups (incremental-forever full database or log backups)
Recovery	Recovery of a database to its most recent state, a point in time data backup or log backup Recovery to the original host or to an alternate host
Test/Dev Copy	Multiple point-in-time copies and instant Test/Dev refresh Accelerate and automate Test/Dev provisioning Migration of PostgreSQL from a physical to a virtualized environment
Backup Catalog view	Actifio Desktop

PostgreSQL DB Snapshots

The PostgreSQL database storage snapshot is taken by the Actifio Connector. If the database is slave, then the slave process is halted (using `pg_xlog_replay_pause`) during the snapshot period. If the database is primary, then a read lock is placed on the database using (`pg_start_backup`) during snapshot and locks are released after snapshot is completed.

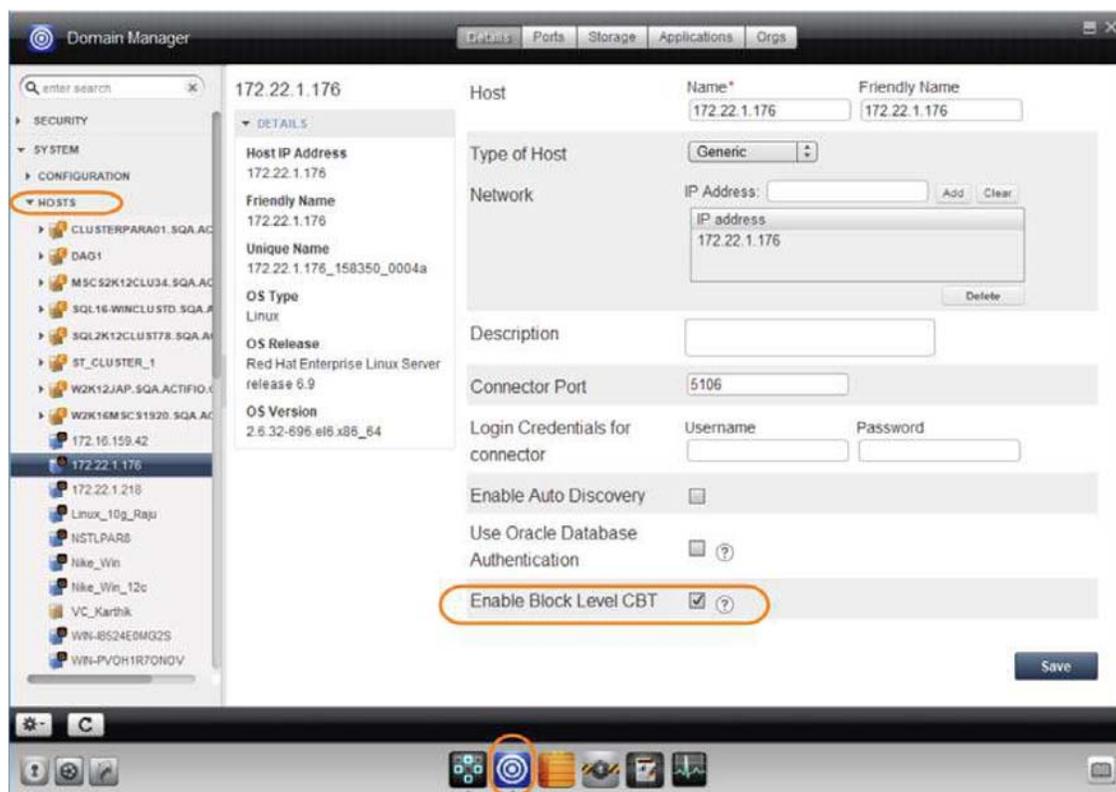
PostgreSQL Configurations Support

PostgreSQL Primary and Slave configuration and protection can be set from primary or from the slave. Protection can be set at instance level (containing all user databases) or can be set at individual database (one or multiple). This can be achieved by setting up the right parameters under the `act_PostgreSQL.conf` config file (see [Protecting a PostgreSQL Database](#) on page 5).

2 Preparing the PostgreSQL Database

Prerequisites

- The PostgreSQL database must be residing under LVM and it must not be the boot volume. To get the database data path, run `ps -ef | grep -i postgres`
- For best results, the LVM volume from which the PostgreSQL volumes are provisioned should have at least 20% free space.
- Install the Actifio Connector on the PostgreSQL server host (see **Connecting Hosts to Actifio Appliances** in ActifioNOW and in the Doc Library on your Actifio appliance.
- Enable Linux change block tracking on the database server from the Actifio Desktop > Domain Manager > Host > Details:



Enabling Linux CBT for the PostgreSQL Server

3 Protecting a PostgreSQL Database

This section includes:

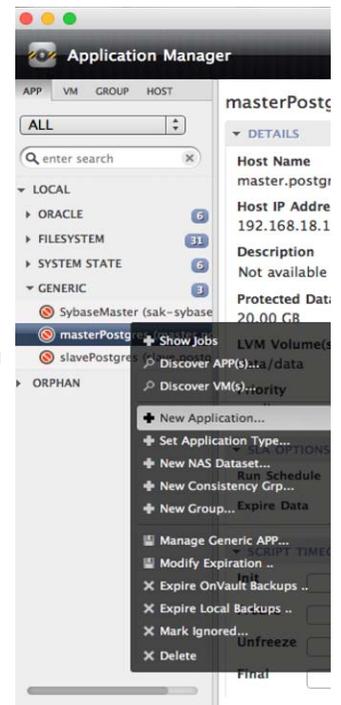
[Protecting a PostgreSQL Database](#) on page 5

[Protecting PostgreSQL Database Logs](#) on page 8

Protecting a PostgreSQL Database

To protect a PostgreSQL database:

1. From the Actifio Desktop > Application Manager Service Menu > New Application, create a new application to protect the data and log volumes of the PostgreSQL database. This will open the Add an Application window.
2. From Add an Application:
 - a. Select **Out-of-Band** as an application type and enter an Application Name to identify this database in the Actifio Desktop.
 - b. From the Select Host drop down, select the PostgreSQL database server.
 - c. Under Select logical volume, use the green + to add the data and log volume to Selected volumes on the right.
 - d. Under Generic LVM Script Name, put `genlvmscript.sh`.
 - e. Click **Add** to add this application as generic application.
 - f. To check, look at Actifio Desktop > Application Manager > Generic.

A screenshot of the 'Manage Application' form in the Actifio interface. The form contains the following fields and options:

- Select Band Type***: Radio buttons for 'In-Band' and 'Out-of-Band' (selected).
- Enter Application Name***: Text input field containing 'masterPostgres'.
- Select Host***: Dropdown menu showing 'master.postgres'.
- Volume Group Slack Space**: Input field showing '20 %'.
- Select Logical Volumes(s)***: A table with two columns: 'Available Volumes(1)' and 'Selected Volumes(1)'.

Available Volumes(1)	Selected Volumes(1)
archive/archiveLog 10.0 GB +	data/data -
- Generic LVM Script Name**: Text input field containing 'genlvmscript.sh'.

3. Set up the scripts. Log into the database server as root and cd to /act: `cd /act`. Create an /act/scripts directory if it does not already exist:

```
mkdir scripts
cd /act/scripts
```

Copy the provided script /act/act_scripts/postgresql/outofband to /act/scripts. This contains 4 scripts:

- o genlvmscript.sh
- o freeze.xxx
- o thaw.xxx
- o act_PostgreSQL.conf

Provide 755 permission: `chmod 755 *`

4. Get the application id for the created generic application. From the Application Manager, mouse over and get the application id, for example 5619.

Note: You can also get this application id from the appliance command line; run `udsinfolapplication`.

5. Modify freeze, thaw script extension from "xxx" to the application id from step 4 above (for example: 5619). From the command line:

```
mv freeze.xxx freeze.5619
mv thaw.xxx thaw.5619
```

6. Edit and modify the act_PostgreSQL.conf script for the input parameters:

```
vi act_PostgreSQL.conf
```

Replace the line:

```
PG_HOME='<protected postgres database home directory path>'
PORT='<protected postgres port>'
OSUSER='<protected postgres OS user>'
PG_LVM_MOUNT='<protected postgres data LVM Mount point. Multiple LVM mount points can
be provided with comma separated>'
DBNAME='<all or postgres database name to be protected (Multiple database names can be
provided with comma separated)>'
```

With:

```
PG_HOME='/home/postgres/postgresql_home'
PORT='5432'
OSUSER='postgres'
PG_LVM_MOUNT=' /data,/data2'
DBNAME='actdb,data2" #,testdb1,acttest'
# uncomment the parameter below and provide the value to protect log backup.
#LOGBKPLC='<Postgres database log backup location>'
#BKP_LOG_RETENTION='<# of days to retain the log backup under LOGBKPLC>'
```

Where:

- PG_HOME='<protected postgres database home directory path>' (PG_HOME can be retrieved by using the command `ps -ef grep postgres`)
- PORT='<protected postgres database port number>' (PORT number on which the database is running can be retrieved from postgresql.conf under \$PG_DATA_PATH directory)
- OSUSER='<protected postgres database OS user>' (OSUSER can be retrieved by using the command `ps -ef grep postgres`)

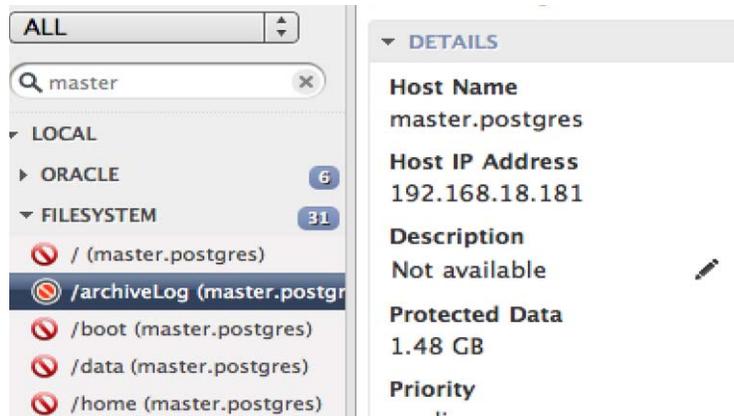
- `PG_LVM_MOUNT='<protected postgres data LVM Mount point. Multiple LVM mount points can be provided if comma-separated>'`
(LVM Mount can be checked using `'ps -ef | grep postgres'` where `-D` will show the data path.)
 - `DBNAME= '<all or postgres database name to be protected (multiple database names can be provided if comma-separated)>'`
 - `#LOGBKPLOC='<Postgres database log backup location>'`
 - `#BKP_LOG_RETENTION='<# of days to retain the log backup under LOGBKPLOC>'`
7. Save the file.
 8. To protect, select an SLA template from Template drop-down list and then select a resource profile from Profile drop-down list. You can run the snapshot job immediately as an on-demand job, or wait for the scheduler to run the job during the time period specified in the template.

Protecting PostgreSQL Database Logs

Note: Database log protection support is from the Primary database only.

To protect PostgreSQL database logs:

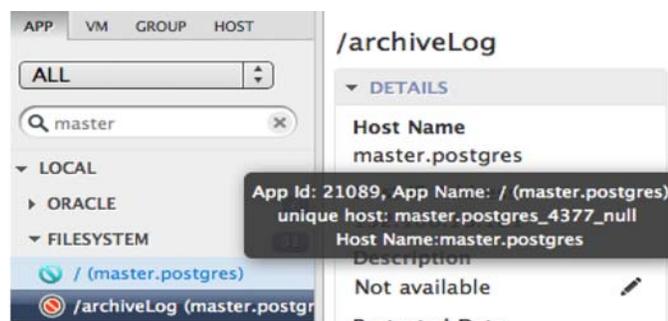
1. Discover the file system (log backup path from [Step 2](#) above) by running Discover App from the Actifio Desktop (if not discovered already).
2. The log backup mount path will appear in the Actifio Desktop under App > Filesystem:



Note: You can restrict the log backup path using "Start Paths" under the Application Advanced Settings.

3. Set up the scripts:
 - a. Log into the database server as root.
 - b. cd to /act (cd /act)
 - c. Create an /act/scripts directory if not there:

```
mkdir scripts
cd /act/scripts
```
 - d. Copy the Actifio-provided scripts in /act/act_scripts/postgresql/outofband/log to /act/scripts.
 - e. This contains 2 files:
 - o appid.xxx
 - o thaw.xxx
 - f. provide 755 permission: `chmod 755 *`
4. From the Application Manager, mouse over and get the application id for the created filesystem application.



Note: Another way to get the application id is to log into the appliance and run the `udsinfo lsapplication` command.

5. Modify the `appid.xxx` and `thaw.xxx` script extension from “xxx” to the appid from [Step 4](#) (ex: 21089)

- a. From the command line:

```
mv thaw.xxx thaw.21089
mv appid.xxx appid.21089
```

6. Edit and modify `act_PostgreSQL.conf` script for input parameters:

```
Vi act_PostgreSQL.conf
#uncomment the parameter below and provide the value
#LOGBKPLC='<Postgres database log backup location>'
#BKP_LOG_RETENTION='<# of days to retain the log backup under LOGBKPLC>'
```

With:

```
LOGBKPLC=/archiveLog/pgdata_xlog
BKP_LOG_RETENTION=2
```

Where:

LOGBKPLC: PostgreSQL archive log backup location

BKP_LOG_RETENTION: <# of days to retain the log backup under LOGBKPLC>

7. To protect:
 - a. Select an SLA template from Template drop-down list.
 - b. Select a resource profile from Profile drop-down list.
 - c. You can run the snapshot job immediately as an on-demand job, or wait for the scheduler to run the job during the time period specified in the template.

4 Accessing, Recovering, or Restoring a PostgreSQL Database

Use cases:

- [Mount and Refresh a Target PostgreSQL Database as a Virtual Application to a Database Backup Point in Time, and/or Roll Forward the Log to a Specific Point in Time](#): To present and refresh a read-write virtual copy of the PostgreSQL database on a new target to a scheduled database backup point in time and or Mount and Refresh a target PostgreSQL database with roll-forward of log.
- [Mount and Refresh the Target PostgreSQL Database as a Virtual Slave Application](#): To add a virtual copy as the secondary node to the primary PostgreSQL database.

Mount and Refresh a Target PostgreSQL Database as a Virtual Application to a Database Backup Point in Time, and/or Roll Forward the Log to a Specific Point in Time

To present and refresh as a virtual copy of the PostgreSQL database from source to any target with the same name as the primary or with a different name:

1. On the target node, set up the scripts:
 - a. Log into the target database node as root.
 - b. cd to /act (cd /act).
 - c. Create /act/scripts directory if not there:

```
mkdir scripts
cd /act/scripts
```
 - d. Copy Actifio provided script from /act/act_scripts/postgresql/outofband/clone to target node /act/scripts:
 - o act_Pre_Target_Master.sh
 - o act_Conf_Target_Master.conf
 - o act_Post_Target_Master.sh
 - o act_Post_Target_Master_PITR.sh
 - e. Edit and update the parameter of the configuration file act_Conf_Target_MASTER.conf for target clone operation:

```
TARGET_MASTER_NODE= <This is target Postgres master node>
TARGET_MASTER_PORT=<This is target Postgres master port>
TARGET_POSTGRES_HOME= <This is target Postgres home>
TARGET_POSTGRES_DATA_PATH= <This is target Postgres data path>
TARGET_MASTER_MOUNTPOINT= <This is database snapshot mountpoint on target>
#TARGET_DB_RENAME="<(Optional). Uncomment and provide List of source and target
databases to be renamed in the format <sourceDB:targetdb,..>. Note that the DB
name is case sensitive. For Example:
"<sourceDB1>:<TargetDB1>,<sourceDB2>:<TargetDB2>">
```

```
# uncomment and specify the value below if archivelog backup image is mounted and
being rollforward
#TARGET_MASTER_ARC_MOUNTPOINT='<Archivelog Mount Directory from Actifio mounted
Image of Target Postgres Master>'
#TARGET_MASTER_ARC_LOCATION='<(optional) Archive log location of Target Postgres
Master>'
#RECOVERY_TIME='<recovery time for roll-forward of Target Postgres Master>'
```

For example:

```
TARGET_MASTER_NODE="192.168.18.183"
TARGET_MASTER_PORT="5499"
TARGET_POSTGRES_HOME="/home/postgres/postgresql_home"
TARGET_POSTGRES_DATA_PATH=/postgresmaster/data/pgdata
TARGET_MASTER_MOUNTPOINT="/postgresmaster/data"
```

Note: To Rename target database uncomment and provide the value in the format below:

```
#TARGET_DB_RENAME="acttest:acttestX,data1:data_1,data2:data22,testdb1:test,noDb:
yesDb"
```

f. Save the file.

Now that the scripts are set up, the rest of the steps are performed in the Actifio Desktop.

2. Open the Actifio Desktop to the Application Manager and select the snapshot image.
3. From the gear icon dropdown menu, select **Mount**. The Mount window opens.



4. Provide a label as needed. This is optional.
5. In the Select Host drop down list, select the PostgreSQL target database nodes.
6. Provide a mount point name.



7. Click **Mount**. This will mount the data volume to the target server as /postgresmaster.

Note: If you have protected the database logs, and if you want to roll forward the logs, continue to [Logs Roll Forward after Mount and Refresh of Target PostgreSQL Database](#).

8. Log into the target database server as root.
9. Change the directory to **/act/scripts** on the target host and run the script **act_Post_Target_Master.sh**. This will bring up PostgreSQL services and database online.

To run this script from command line, edit **act_Post_Target_Primary.sh** and uncomment these two lines:

```
#export ACT_JOBTYPE="mount"  
#export ACT_PHASE="post"
```

Note: If running from workflow using these scripts as pre/post, make sure these two lines are commented.

```
./act_Post_Target_Master.sh
```

Logs Roll Forward after Mount and Refresh of Target PostgreSQL Database

For a point in time mount and refresh of a target database, first follow up to Step 7 of the procedure in [Mount and Refresh the Target PostgreSQL Database as a Virtual Slave Application](#) on page 17, and then:

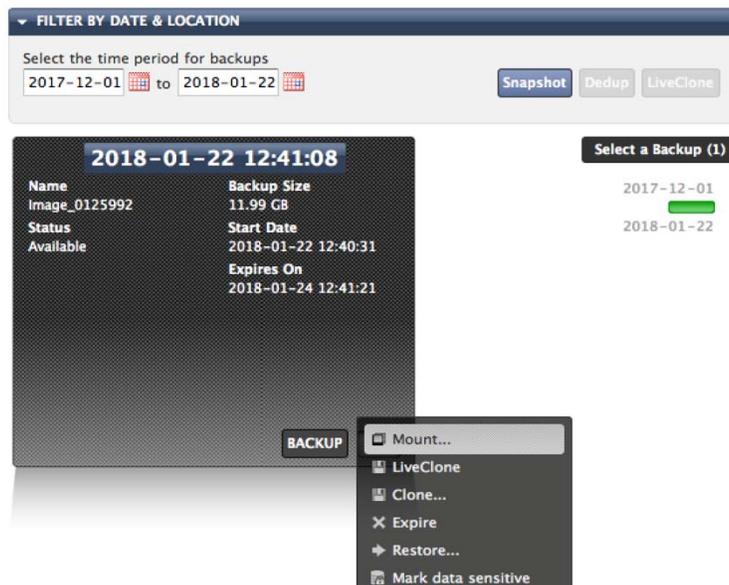
1. On the target node, set up the scripts:
 - a. Log into the target database node as root.
 - b. cd to /act/scripts.

```
cd /act/scripts
```
2. Edit, uncomment, and update the parameter of the configuration file `act_Conf_Target_Master.conf`:

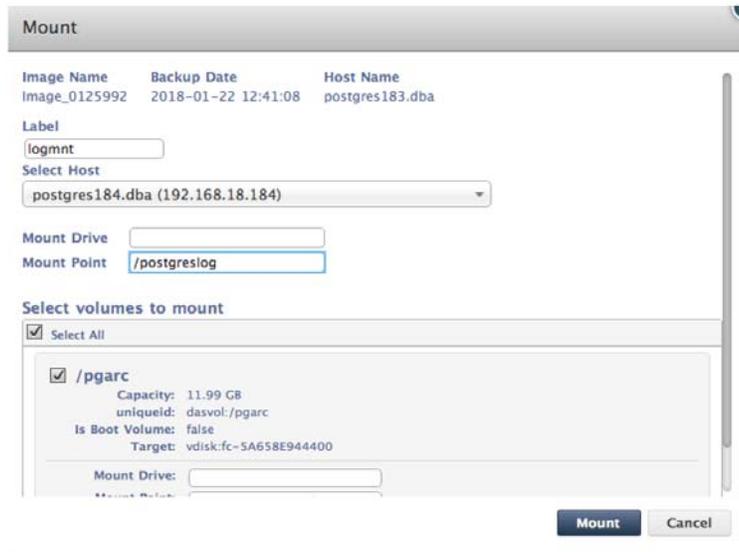
```
# If the archive is protected and the target database is rollforward then the below parameter is required
TARGET_MASTER_ARC_LOCATION= <This is archive log snapshot mountpoint on target>
TARGET_MASTER_ARC_MOUNTPOINT= <Optional parameter. This is archive log mountpoint>
RECOVERY_TIME= <Optional parameter. This is point in time recovery in format 'yyyy-mm-dd hh24:mi:ss'>
```

With:

```
TARGET_MASTER_ARC_MOUNTPOINT="/postgreslog"
TARGET_MASTER_ARC_LOCATION="/postgreslog/pgdata_xlog"
RECOVERY_TIME="2018-02-15 09:09:35.152288-05"
```
3. Save the file.
4. Open the Actifio Desktop to the **Application Manager** and from the protected log backup file system application select the Snapshot Image.
5. From the gear icon dropdown menu, select **Mount**. The Mount window opens:



6. Provide a label as needed. This is optional.
7. In the **Select Host** drop down list, select the PostgreSQL target database node.
8. Provide a mount point name.



9. Click **Mount**. This will mount the data volume to the target server as /postgreslog.
10. Change the directory to /act/scripts on the target host and run the script act_Post_Target_Primary.sh. This will bring up PostgreSQL services and database online with rollforward of log to the specified point in time.

Note: To run this script from command line, edit act_Post_Target_Primary.sh and uncomment these two lines:

```
#export ACT_JOBTYPE="mount"
#export ACT_PHASE="post"

./act_Post_Target_Master.sh
```

To Unmount and Delete the Image

To unmount and delete any previously mounted image:

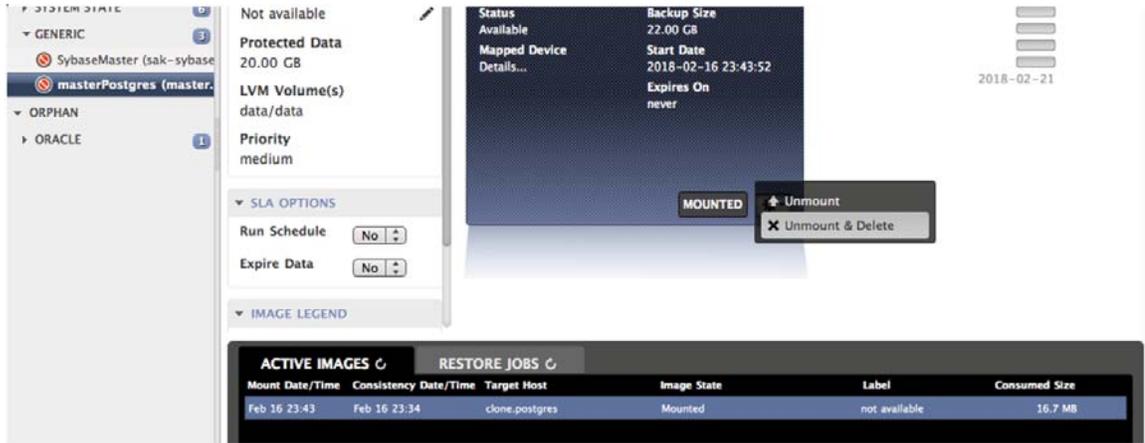
1. Log into the target database server as root.
2. `cd /act/scripts`
3. Run the pre script to stop and clean up any previously mounted image.
4. To run this script from command line, edit **act_Pre_Target_Master.sh** and uncomment these two lines:

```
#export ACT_JOBTYPE="unmount"
#export ACT_PHASE="pre"
```

Note: If running from workflow using these scripts as pre/post, make sure these two lines are commented.

5. Run the act_Pre_AddToMaster_Slave.sh script:


```
./act_Pre_AddToMaster_Slave.sh
```
6. Open the Actifio Desktop to the Application Manager.
7. From the protected application, click on the **Restore** tab and select the mounted image from the ACTIVE IMAGES window.
8. Click the gear wheel and select **Unmount & Delete**.



Unmounting and Deleting the PostgreSQL Image

- The Unmount window opens. **Submit** the job.



Confirming the Mounted Image Deletion

Mount and Refresh the Target PostgreSQL Database as a Virtual Slave Application

To present and refresh as a virtual copy of the PostgreSQL database from source to any target as Slave of the source Primary database:

1. On Target node set up the scripts:
 - a. Login to target database node as root.
 - b. cd to /act (cd /act)
 - c. Create an /act/scripts directory if it does not exist:

```
mkdir scripts
cd /act/scripts
```

- d. Copy the Actifio-provided scripts from /act/act_scripts/PostgreSQL/outofband/clone to target node /act/scripts folder.
 - o act_Pre_AddToMaster_Slave.sh
 - o act_Conf_AddToMaster_Slave.conf
 - o act_Post_AddToMaster_Slave.sh
- e. Edit and update the configuration file act_Conf_AddToMaster_Slave.conf for these parameters:

```
TARGET_SLAVE_NODE='<Host IP of Target Postgres Slave>'
TARGET_SLAVE_PORT='<Port of Target Postgres Slave>'
TARGET_POSTGRES_HOME='<Home directory of Target Postgres Slave Home>'
POSTGRES_MASTER_NODE='<Host IP of Source Postgres>'
POSTGRES_MASTER_PORT='<Port of Source Postgres>'
POSTGRES_MASTER_DATA_PATH='<Data Path directory of Source Postgres>'
POSTGRES_MASTER_HOME='<Home directory of Source Postgres>'
POSTGRES_MASTER_OS_USER='<OS user of Source Postgres>'
TARGET_SLAVE_MOUNTPOINT='<Mountpoint of Target Postgres Slave>'
TARGET_POSTGRES_DATA_PATH='<Data directory of Target Postgres Slave>'
```

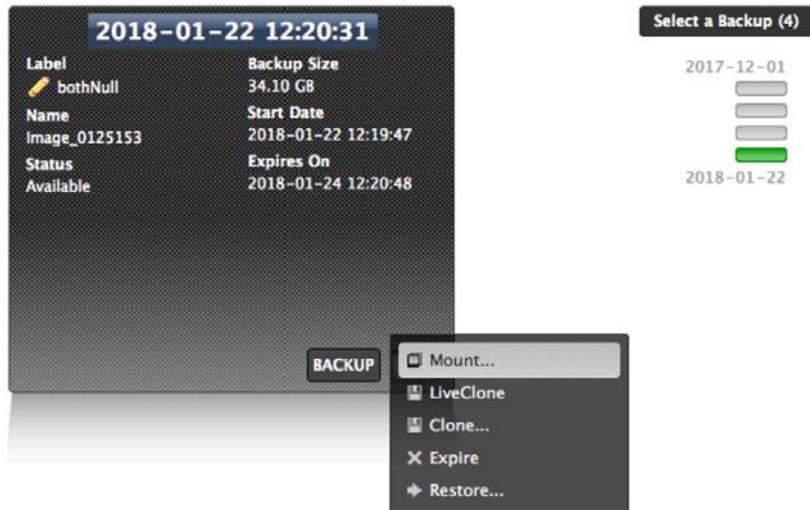
Example:

PORT number on which the master database is running can be retrieved from postgresql.conf under \$PG_DATA_PATH directory on the master database host.

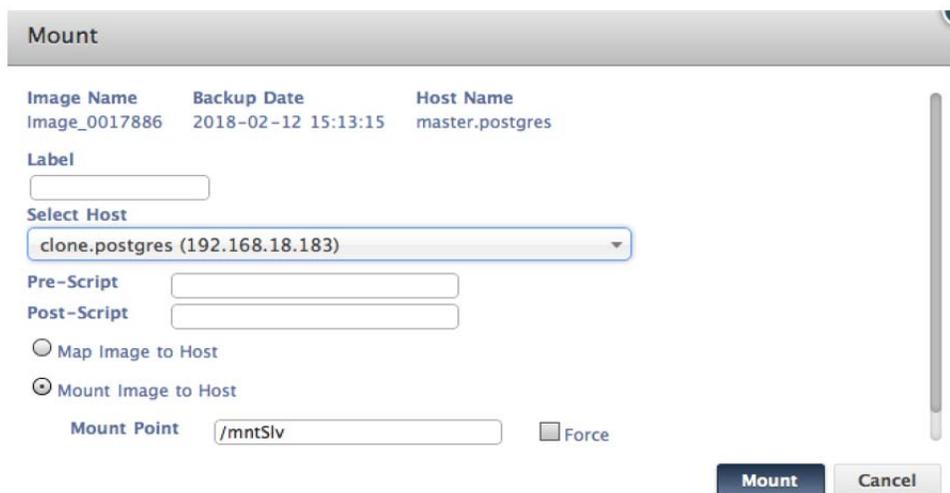
PG_HOME and OSUSER can be retrieved by using the command 'ps -ef grep postgres' on master database host.

```
TARGET_SLAVE_NODE="192.168.18.183"
TARGET_SLAVE_PORT="5491"
TARGET_POSTGRES_HOME="/home/postgres/postgresql_home"
POSTGRES_MASTER_NODE="192.168.18.181"
POSTGRES_MASTER_PORT="5432"
POSTGRES_MASTER_DATA_PATH="/data/pgdata"
POSTGRES_MASTER_HOME="/home/postgres/postgresql_home"
POSTGRES_MASTER_OS_USER="postgres"
TARGET_SLAVE_MOUNTPOINT="/mntSlv/data"
TARGET_POSTGRES_DATA_PATH="/mntSlv/data/pgdata"
```

2. Save the file.
3. Open the Actifio Desktop to the Application Manager and select the snapshot image.
4. From the gear icon dropdown menu, select **Mount**. The Mount window opens.



5. Provide a label as needed. This is optional.
6. In the Select Host drop down list, select the PostgreSQL target database node.
7. Provide a mount point name.



8. Click **Mount**. This will mount the data volume to the target server as /mntSlv.
9. On the target host, log in as root and change the directory to /act/scripts and execute the script `act_Post_AddToMaster_Slave.sh`.

To run this script from command line, open the **act_Post_AddToMaster_Slave.sh** script and uncomment these two lines:

```
#export ACT_JOBTYPE="mount"
#export ACT_PHASE="post"

./act_Post_AddToMaster_Slave.sh
```

Note: If running from workflow using these script as pre/post, make sure these two lines are commented

This will bring up the PostgreSQL services and convert the database to slave.

How to Unmount and Delete the Image

1. Log into the target server as root.
2. `cd /act/scripts`
3. Run the pre script to stop and clean up any previously mounted image:

To run this script from command line, edit **act_Pre_AddToMaster_Slave.sh** and uncomment these two lines:

```
#export ACT_JOBTYPE="unmount"  
#export ACT_PHASE="pre"
```

Note: If running from workflow using these script as pre/post, make sure these two lines are commented.

4. Run the `act_Pre_Target_Slave.sh` script: **xxxAddToMaster?**
`./act_Pre_AddToMaster_Slave.sh`
5. Open the Actifio Desktop to the Application Manager.
6. From the target, click on restore tab and select the mounted image from the Active Image panel at the bottom.
7. Click the gear wheel and select **Unmount & Delete**.



8. The Unmount window opens. **Submit** the job.



