



Clone PostgreSQL resource backups

SnapCenter Software 6.0

NetApp
July 23, 2024

Table of Contents

- Clone PostgreSQL resource backups 1
 - Clone workflow 1
 - Clone a PostgreSQL backup 1
 - Clone PostgreSQL cluster backups using PowerShell cmdlets 3
- Monitor PostgreSQL clone operations 5
- Split a clone 6
- Delete or split PostgreSQL cluster clones after upgrading SnapCenter 7

Clone PostgreSQL resource backups


Clone workflow

The clone workflow includes performing the clone operation and monitoring the operation.

About this task

- You can clone on the source PostgreSQL server.
- You might clone resource backups for the following reasons:
 - To test functionality that has to be implemented using the current resource structure and content during application development cycles
 - For data extraction and manipulation tools when populating data warehouses
 - To recover data that was mistakenly deleted or changed

The following workflow shows the sequence in which you must perform the clone operation:



You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

Clone a PostgreSQL backup

You can use SnapCenter to clone a backup. You can clone from primary or secondary backup.

Before you begin

- You should have backed up the resources or resource group.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- For pre clone or post clone commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:

For Windows: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands_list.txt*

For Linux: */var/opt/snapcenter/scc/allowed_commands_list.txt*



If the commands do not exist in the command list, then the operation will fail.

About this task

- For information about clone split operation limitations, see [ONTAP 9 Logical Storage Management Guide](#).
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.


The resources are displayed along with information such as type, host, associated resource groups and policies, and status.

3. Select the resource or resource group.

You must select a resource if you select a resource group.

The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.

5. Select the data backup from the table, and then click .

6. In the Location page, perform the following actions:

For this field...	Do this...
Clone server	Choose a host on which the clone should be created.
Target Port	Enter the target PostgreSQL target port to clone from the existing backups.
NFS Export IP Address	Enter IP addresses or the host names on which the cloned volumes will be exported. This is applicable only to NFS storage type resource.
Capacity Pool Max. Throughput (MiB/s)	Enter the maximum throughput of a capacity pool. This is applicable only for ANF storage type resource.

7. In the Scripts page, perform the following steps:



The scripts are run on the plug-in host.

- a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.
 - Pre clone command: delete existing clusters with the same name
 - Post clone command: verify a cluster or start a cluster.
- b. Enter the mount command to mount a file system to a host.

Mount command for a volume or qtree on a Linux machine:

Example for NFS:

```
mount VSERVER_DATA_IP:%VOLUME_NAME_Clone /mnt
```



For pre and post commands for quiesce, Snapshot, and unquiesce operations, you should check if the commands exist in the command list available on the plug-in host from the `/opt/snapcenter/snapcenter/scc/allowed_commands.config` path for Linux and `C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands_list.txt` for Windows.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email.

9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor > Jobs**.

Clone PostgreSQL cluster backups using PowerShell cmdlets

The clone workflow includes planning, performing the clone operation, and monitoring the operation.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help command_name`. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the `Open-SmConnection` cmdlet.

```
PS C:\> Open-SmConnection
```

2. Retrieve the backups to perform the clone operation by using the `Get-SmBackup` cmdlet.

This example shows that two backups are available for cloning:

```
C:\PS> Get-SmBackup
```

BackupId	BackupName
BackupTime	BackupType
-----	-----
1	Payroll Dataset_vise-f6_08... 8/4/2015
11:02:32 AM	Full Backup
2	Payroll Dataset_vise-f6_08... 8/4/2015
11:23:17 AM	

3. Initiate a clone operation from an existing backup and specify the NFS export IP addresses on which the cloned volumes are exported.

This example shows that the backup to be cloned has an NFSExportIPs address of 10.32.212.14:

For PostgreSQL cluster:

```
PS C:\> New-SmClone -AppPluginCode PostgreSQL -BackupName "
scpostgresl01_ openenglab_netapp_com_PostgreSQL_postgres_5432_06-26-
2024_00_33_41_1570" -Resources @{"Host"="
10.32.212.13";"Uid"="postgres_5432"} -port 2345 -CloneToHost
10.32.212.14
```



If NFSExportIPs is not specified, the default is exported to the clone target host.

4. Verify that the backups were cloned successfully by using the Get-SmCloneReport cmdlet to view the clone job details.

You can view details such as clone ID, start date and time, end date and time.

```

PS C:\> Get-SmCloneReport -JobId 186

SmCloneId           : 1
SmJobId             : 186
StartDateTime       : 8/3/2015 2:43:02 PM
EndDateTime         : 8/3/2015 2:44:08 PM
Duration            : 00:01:06.6760000
Status              : Completed
ProtectionGroupName : Draper
SmProtectionGroupId : 4
PolicyName          : OnDemand_Clone
SmPolicyId          : 4
BackupPolicyName    : OnDemand_Full_Log
SmBackupPolicyId    : 1
CloneHostName       : SCSPR0054212005.mycompany.com
CloneHostId         : 4
CloneName           : Draper__clone__08-03-2015_14.43.53
SourceResources     : {Don, Betty, Bobby, Sally}
ClonedResources     : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
Sally_DRAPER}
SmJobError          :







```

Monitor PostgreSQL clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.


About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.

3. In the **Jobs** page, perform the following steps:
 - a. Click  to filter the list so that only clone operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Clone**.
 - d. From the **Status** drop-down list, select the clone status.
 - e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about clone split operation limitations, see [ONTAP 9 Logical Storage Management Guide](#).
- Ensure that the volume or aggregate on the storage system is online.


Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select Database from the View list.
For file systems	Select Path from the View list.

3. Select the appropriate resource from the list.

The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.

6. Monitor the operation progress by clicking **Monitor > Jobs**.

The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

For example:

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```

The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

[SnapCenter clone or verification fails with aggregate does not exist](#)

Delete or split PostgreSQL cluster clones after upgrading SnapCenter

After upgrading to SnapCenter 4.3, you will no longer see the clones. You can delete the clone or split the clones from the Topology page of the resource from which the clones were created.



About this task

If you want to locate the storage footprint of the hidden clones, run the following command: `Get-SmClone -ListStorageFootprint`

Steps

1. Delete the backups of the cloned resources by using the `remove-smbbackup` cmdlet.
2. Delete the resource group of the cloned resources by using the `remove-smresourcegroup` cmdlet.
3. Remove the protection of the cloned resource by using the `remove-smprotectresource` cmdlet.
4. Select the parent resource from the Resources page.

The resource topology page is displayed.

5. From the Manage Copies view, select the clones either from the primary or secondary (mirrored or replicated) storage systems.
6. Select the clones, and then click  to delete clones or click  to split the clones.
7. Click **OK**.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.