



SnapCenter Plug-in for PostgreSQL

SnapCenter Software 6.0

NetApp
July 23, 2024

Table of Contents

- SnapCenter Plug-in for PostgreSQL 1
 - SnapCenter Plug-in for PostgreSQL overview 1
 - What you can do using the SnapCenter Plug-in for PostgreSQL 1
 - SnapCenter Plug-in for PostgreSQL features 1
 - Storage types supported by SnapCenter Plug-in for PostgreSQL 2
 - Minimum ONTAP privileges required for PostgreSQL plug-in 3
 - Prepare storage systems for SnapMirror and SnapVault replication for PostgreSQL 5
 - Backup strategy for PostgreSQL 5
 - Restore and recovery strategy for PostgreSQL 8

SnapCenter Plug-in for PostgreSQL

SnapCenter Plug-in for PostgreSQL overview

The SnapCenter Plug-in for PostgreSQL cluster is a host-side component of the NetApp SnapCenter software that enables application-aware data protection management of PostgreSQL clusters. The Plug-in for PostgreSQL cluster automates the backup, restore, and cloning of PostgreSQL clusters in your SnapCenter environment.

SnapCenter supports single cluster and multi cluster PostgreSQL setups. You can use the Plug-in for PostgreSQL Clusters in both Linux and Windows environments. In Windows environments, PostgreSQL will be supported as manual resource.

When the Plug-in for PostgreSQL cluster is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume. You can also use the plug-in with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance.

SnapCenter Plug-in for PostgreSQL supports NFS and SAN on ONTAP and Azure NetApp File storage layouts.

VMDK or virtual storage layout is supported.

What you can do using the SnapCenter Plug-in for PostgreSQL

When you install the Plug-in for PostgreSQL cluster in your environment, you can use SnapCenter to back up, restore, and clone PostgreSQL clusters and their resources. You can also perform tasks supporting those operations.

- Add clusters.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

SnapCenter Plug-in for PostgreSQL features

SnapCenter integrates with the plug-in application and with NetApp technologies on the storage system. To work with the Plug-in for PostgreSQL Cluster, you use the SnapCenter graphical user interface.

- **Unified graphical user interface**

The SnapCenter interface provides standardization and consistency across plug-ins and environments.

The SnapCenter interface enables you to complete consistent backup, restore, and clone operations across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- **Automated central administration**

You can schedule backup operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp snapshot copy technology**

SnapCenter uses NetApp snapshot technology with the Plug-in for PostgreSQL cluster to back up resources.

Using the Plug-in for PostgreSQL also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation

You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology

A FlexClone license is required on the storage system where you want to create the clone.

- Support for the consistency group (CG) snapshot feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts

In a single operation, snapshots are consolidated when resources in a single host share the same volume.

- Capability to create snapshots using external commands.
- Support for Linux LVM on XFS file system.

Storage types supported by SnapCenter Plug-in for PostgreSQL

SnapCenter supports a wide range of storage types on both physical machines and virtual machines (VMs). You must verify the support for your storage type before installing SnapCenter Plug-in for PostgreSQL.

| Machine | Storage type |
|------------------------------|-----------------------|
| Physical and virtual servers | FC-connected LUNs |
| Physical server | iSCSI-connected LUNs |
| Physical and virtual servers | NFS-connected volumes |

Minimum ONTAP privileges required for PostgreSQL plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 8.3.0 and later
 - event generate-autosupport-log
 - job history show
 - job stop
 - lun
 - lun create
 - lun create
 - lun create
 - lun delete
 - lun igroup add
 - lun igroup create
 - lun igroup delete
 - lun igroup rename
 - lun igroup rename
 - lun igroup show
 - lun mapping add-reporting-nodes
 - lun mapping create
 - lun mapping delete
 - lun mapping remove-reporting-nodes
 - lun mapping show
 - lun modify
 - lun move-in-volume
 - lun offline
 - lun online
 - lun persistent-reservation clear
 - lun resize
 - lun serial
 - lun show
 - snapmirror policy add-rule
 - snapmirror policy modify-rule
 - snapmirror policy remove-rule
 - snapmirror policy show
 - snapmirror restore

- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify
- volume qtree create
- volume qtree delete
- volume qtree modify
- volume qtree show
- volume restrict
- volume show
- volume snapshot create
- volume snapshot delete
- volume snapshot modify
- volume snapshot modify-snaplock-expiry-time
- volume snapshot rename
- volume snapshot restore
- volume snapshot restore-file
- volume snapshot show
- volume unmount
- vservers cifs
- vservers cifs share create
- vservers cifs share delete
- vservers cifs shadowcopy show
- vservers cifs share show

- vserver cifs show
- vserver export-policy
- vserver export-policy create
- vserver export-policy delete
- vserver export-policy rule create
- vserver export-policy rule show
- vserver export-policy show
- vserver iscsi
- vserver iscsi connection show
- vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
 - network interface
 - network interface show
 - vserver

Prepare storage systems for SnapMirror and SnapVault replication for PostgreSQL

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.



If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.



SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary > Mirror > Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the [ONTAP documentation](#).

Backup strategy for PostgreSQL

Define a backup strategy for PostgreSQL

Defining a backup strategy before you create your backup jobs helps you to have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent snapshots of the cluster.
5. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.
6. Determine the retention period for the snapshots on the source storage system and the SnapMirror destination.
7. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Automatic discovery of resources on Linux host

Resources are PostgreSQL clusters and instances on the Linux host that are managed by SnapCenter. After installing the SnapCenter Plug-in for PostgreSQL plug-in, the PostgreSQL clusters from all the instances on that Linux host are automatically discovered and displayed in the Resources page.

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports snapshot copy-based backup type for PostgreSQL clusters.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp snapshot technology to create online, read-only copies of the volumes on which the PostgreSQL clusters reside.

How SnapCenter Plug-in for PostgreSQL uses consistency group snapshots

You can use the plug-in to create consistency group snapshots for resource groups. A consistency group is a container that can house multiple volumes so that you can

manage them as one entity. A consistency group is simultaneous snapshots of multiple volumes, providing consistent copies of a group of volumes.

You can also specify the wait time for the storage controller to consistently group snapshots. The available wait time options are **Urgent**, **Medium**, and **Relaxed**. You can also enable or disable Write Anywhere File Layout (WAFL) sync during consistent group snapshot operation. WAFL sync improves the performance of a consistency group snapshot.

How SnapCenter manages housekeeping of data backups

SnapCenter manages the housekeeping of data backups on the storage system and file system levels.

The snapshots on the primary or secondary storage and their corresponding entries in the PostgreSQL catalog are deleted based on the retention settings.

Considerations for determining backup schedules for PostgreSQL

The most critical factor in determining a backup schedule is the rate of change for the resource. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)

Backup frequency, also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

- Backup schedules (exactly when backups are to be performed)

Backup schedules are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m.

Number of backup jobs needed for PostgreSQL

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions for Plug-in for PostgreSQL clusters

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

resourcegroupname_hostname_timestamp

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

Restore and recovery strategy for PostgreSQL

Define a restore and recovery strategy for PostgreSQL resources

You must define a strategy before you restore and recover your cluster so that you can perform restore and recovery operations successfully.



Only manual recovery of cluster is supported.

Steps

1. Determine the restore strategies supported for manually added PostgreSQL resources
2. Determine the restore strategies supported for auto discovered PostgreSQL clusters
3. Decide the type of recovery operations that you want to perform.

Types of restore strategies supported for manually added PostgreSQL resources

You must define a strategy before you can successfully perform restore operations using SnapCenter.



You cannot recover manually added PostgreSQL resources.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource



If the resource contains volumes or qtrees, the snapshots taken after the snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

NOTE: Plug-in for PostgreSQL creates a backup_label and tablespace_map in

`/<OS_temp_folder>/postgresql_sc_recovery<Restore_JobId>/_` folder to help recover manually .

Type of restore strategy supported for automatically discovered PostgreSQL

You must define a strategy before you can successfully perform restore operations using SnapCenter.

Complete resource restore is the restore strategy supported for automatically discovered PostgreSQL clusters. This restores all the volumes, qtrees, and LUNs of a resource.

Types of restore operations for auto discovered PostgreSQL

SnapCenter Plug-in for PostgreSQL supports Single File SnapRestore, and connect-and-copy restore types for automatically discovered PostgreSQL clusters.

Single File SnapRestore is performed in NFS environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup selected is from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Single File SnapRestore is performed in SAN environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup is selected from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Types of recovery operations supported for PostgreSQL clusters

SnapCenter enables you to perform different types of recovery operations for PostgreSQL clusters.

- Recover the cluster up to the most recent state
- Recover the cluster up to a specific point in time

You must specify the date and time for recovery.

SnapCenter also provides the No recovery option for PostgreSQL clusters.

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.