



Clone SQL Server database resources

SnapCenter Software 4.6

NetApp
June 11, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapcenter-46/protect-scsql/reference_clone_sql_server_database_resources.html on June 11, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Clone SQL Server database resources 1
 - Clone workflow 1
 - Clone from a SQL Server database backup 1
 - Clone backups using PowerShell cmdlets 4
 - Perform Clone Lifecycle 8
 - Monitor SQL database clone operations 10
 - Cancel SQL resource clone operations 11
 - Split a clone 11


Clone SQL Server database resources

Clone workflow

You must perform several tasks using SnapCenter Server before cloning database resources from a backup. Database cloning is the process of creating a point-in-time copy of a production database or its backup set. You can clone databases to test functionality that has to be implemented using the current database structure and content during application development cycles, to use the data extraction and manipulation tools when populating data warehouses, or to recover data that was mistakenly deleted or changed.

A database cloning operation generates reports based on the job IDs.

The following workflow shows the sequence in which you must perform the cloning operations:



You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, recovery, verify, and clone operations. For detailed information about PowerShell cmdlets, use the SnapCenter cmdlet help or see the [SnapCenter Software Cmdlet Reference Guide](#)

Find more information

[Clone from a SQL Server database backup](#)

[Perform Clone Lifecycle](#)

[Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value](#)

Clone from a SQL Server database backup

You can use SnapCenter to clone a SQL Server database backup. If you want to access or restore an older version of the data, you can clone database backups on demand.

What you will need

- You should have prepared for data protection by completing tasks such as adding hosts, identifying resources, and creating storage system connections.
- You should have backed up databases or resource groups.
- The protection type such as mirror, vault, or mirror-vault for data LUN and log LUN should be same to discover secondary locators during cloning to an alternate host using log backups.
- If the mounted clone drive cannot be found during a SnapCenter clone operation, you should change the CloneRetryTimeout parameter of SnapCenter Server to 300.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).

About this task

- While cloning to a standalone database instance, ensure that the mount point path exists and it is a dedicated disk.
- While cloning to a Failover Cluster Instance (FCI), ensure that the mount points exists, it is a shared disk, and the path and the FCI should belong to the same SQL resource group.
- Ensure that there is only one vFC or FC initiator attached to each host. This is because, SnapCenter supports only one initiator per host.
- If the source database or the target instance is on a cluster shared volume (csv), then the cloned database will be on the csv.



For virtual environments (VMDK/RDM), ensure that the mount point is a dedicated disk.


Steps

1. In the left navigation pane, click **Resources**, and then select **SnapCenter Plug-in for SQL Server** from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.



Cloning of a backup of an instance is not supported.

Steps

1. Select the database or resource group.
2. From the Manage Copies view page, select the backup either from primary or secondary (mirrored or vaulted) storage system.
3. Select the backup, and then click .
4. In the Clone Options page, perform the following actions:

For this field...	Do this...
Clone server	Choose a host on which the clone should be created.

For this field...	Do this...
Clone instance	<p>Choose a clone instance to which you want to clone the database backup.</p> <p>This SQL instance must be located in the specified clone server.</p>
Clone suffix	<p>Enter a suffix that will be appended to the clone file name to identify that the database is a clone.</p> <p>For example, <i>db1_clone</i>. If you are cloning to the same location as the original database, you must provide a suffix to differentiate the cloned database from the original database. Otherwise, the operation fails.</p>
Auto assign mount point or Auto assign volume mount point under path	<p>Choose whether to automatically assign a mount point or a volume mount point under a path.</p> <p>Auto assign volume mount point under path: The mount point under a path allows you to provide a specific directory. The mount points will be created within that directory. Before you choose this option, you must ensure that the directory is empty. If there is a database in the directory, the database will be in an invalid state after the mount operation.</p>

5. In the Logs page, select one of the following options:

For this field...	Do this...
None	Choose this option when you want to clone only the full backup without any logs.
All log backups	Choose this option to clone all the available log backups dated after the full backup.
By log backups until	Choose this option to clone the database based on the backup logs that were created up to the backup log with the selected date.
By specific date until	<p>Specify the date and time after which the transaction logs are not applied to the cloned database.</p> <p>This point-in-time clone halts the clone of the transaction log entries that were recorded after the specified date and time.</p>

6. In the Script page, enter the script timeout, path, and the arguments of the prescript or postscript that

should be run before or after the clone operation, respectively.

For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

The default script timeout is 60 seconds.

7. 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. If you want to attach the report of the clone operation performed, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using either the GUI or the PowerShell command `Set-SmSmtServer`.

For EMS, you can refer to [Manage EMS data collection](#)

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

After you finish

After the clone is created, you should never rename it.

Find more information

[Back up SQL Server database, or instance, or availability group](#)

[Clone backups using PowerShell cmdlets](#)

[Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value](#)

[The failover cluster instance database clone fails](#)

Clone 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.

Steps

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

```
Open-SmConnection -SMSbaseurl https://snapctr.demo.netapp.com:8146
```

2. List the backups that can be cloned by using the `Get-SmBackup` or `Get-SmResourceGroup` cmdlet.

This example displays information about all available backups:

```
C:\PS>PS C:\> 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	

This example displays information about a specified resource group, its resources, and associated policies:

```
PS C:\> Get-SmResourceGroup -ListResources -ListPolicies
```

```
Description :  
CreationTime : 8/4/2015 3:44:05 PM  
ModificationTime : 8/4/2015 3:44:05 PM  
EnableEmail : False  
EmailSMTPServer :  
EmailFrom :  
EmailTo :  
EmailSubject :  
EnableSysLog : False  
ProtectionGroupType : Backup  
EnableAsupOnFailure : False  
Policies : {FinancePolicy}  
HostResourceMapping : {}  
Configuration : SMCoreContracts.SmCloneConfiguration  
LastBackupStatus :  
VerificationServer :  
EmailBody :  
EmailNotificationPreference : Never  
VerificationServerInfo : SMCoreContracts.SmVerificationServerInfo  
SchedulerSQLInstance :  
CustomText :  
CustomSnapshotFormat :  
SearchResources : False  
ByPassCredential : False  
IsCustomSnapshot :  
MaintenanceStatus : Production  
PluginProtectionGroupTypes : {SMSQL}  
Name : Payrolldataset  
Type : Group  
Id : 1  
Host :
```

```
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
ApplySnapvaultUpdate : False
ApplyRetention : False
RetentionCount : 0
RetentionDays : 0
ApplySnapMirrorUpdate : False
SnapVaultLabel :
MirrorVaultUpdateRetryCount : 7
AppPolicies : {}
Description : FinancePolicy
PreScriptPath :
PreScriptArguments :
PostScriptPath :
PostScriptArguments :
ScriptTimeout : 60000
DateModified : 8/4/2015 3:43:30 PM
DateCreated : 8/4/2015 3:43:30 PM
Schedule : SMCoreContracts.SmSchedule
PolicyType : Backup
PluginPolicyType : SMSQL
Name : FinancePolicy
Type :
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
clab-a13-13.sddev.lab.netapp.com
DatabaseGUID :
SQLInstance : clab-a13-13
DbStatus : AutoClosed
DbAccess : eUndefined
IsSystemDb : False
IsSimpleRecoveryMode : False
IsSelectable : True
SqlDbFileGroups : {}
SqlDbLogFiles : {}
AppFileStorageGroups : {}
```



```

LogDirectory :
AgName :
Version :
VolumeGroupIndex : -1
IsSecondary : False
Name : TEST
Type : SQL Database
Id : clab-a13-13\TEST
Host : clab-a13-13.sddev.mycompany.com
UserName :
Passphrase :
Deleted : False
Auth : SMCOREContracts.SmAuth
IsClone : False

```

3. Initiate a clone operation from an existing backup by using the `New-SmClone` cmdlet.

This example creates a clone from a specified backup with all logs:

```

PS C:\> New-SmClone
-BackupName payroll_dataset_vise-f3_08-05-2015_15.28.28.9774
-Resources @{"Host"="vise-f3.sddev.mycompany.com";
"Type"="SQL Database";"Names"="vise-f3\SQLExpress\payroll"}
-CloneToInstance vise-f3\sqlexpress -AutoAssignMountPoint
-Suffix _clonefrombackup
-LogRestoreType All -Policy clonefromprimary_ondemand

PS C:> New-SmBackup -ResourceGroupName PayrollDataset -Policy
FinancePolicy

```

This example creates a clone to a specified Microsoft SQL Server instance:

```

PS C:\> New-SmClone
-BackupName "BackupDS1_NY-VM-SC-SQL_12-08-2015_09.00.24.8367"
-Resources @{"host"="ny-vm-sc-sql";"Type"="SQL Database";
"Names"="ny-vm-sc-sql\AdventureWorks2012_data"}
-AppPluginCode SMSQL -CloneToInstance "ny-vm-sc-sql"
-Suffix _CLPOSH -AssignMountPointUnderPath "C:\SCMounts"

```

4. View the status of the clone job by using the `Get-SmCloneReport` cmdlet.

This example displays a clone report for the specified job ID:

```
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}
```

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](#).

Perform Clone Lifecycle

Using SnapCenter, you can create clones from a resource group or database. You can either perform on-demand clone or you can schedule recurring clone operations of a resource group or database. If you clone a backup periodically, you can use the clone to develop applications, populate data, or recover data.

SnapCenter enables you to schedule multiple clone operations to run simultaneously across multiple servers.

What you will need

- While cloning to a standalone database instance, ensure that the mount point path exists and it is a dedicated disk.
- While cloning to a Failover Cluster Instance (FCI), ensure that the mount points exists, it is a shared disk, and the path and the FCI should belong to the same SQL resource group.
- If the source database or the target instance is on a cluster shared volume (csv), then the cloned database will be on the csv.



For virtual environments (VMDK/RDM), ensure that the mount point is a dedicated disk.

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 either **Database** or **Resource Group** from the **View** list.
3. Select the resource group or database, and then click **Clone Lifecycle**.
4. In the Options page, perform the following actions:

For this field...	Do this...
Clone job name	Specify the clone life cycle job name that helps in monitoring and modifying the clone life cycle job.
Clone server	Choose the host on which the clone should be placed.
Clone instance	Choose the clone instance to which you want to clone the database. This SQL instance must be located in the specified clone server.
Clone suffix	Enter a suffix that will be appended to the clone database to identify that it is a clone. Each SQL instance that is used to create a clone resource group must have a unique database name. For example, if the clone resource group contains a source database “db1” from an SQL instance “inst1”, and if “db1” is cloned to “inst1”, then the clone database name should be “db1clone”. “clone” is a mandatory user-defined suffix because the database is cloned to the same instance. If “db1” is cloned to the SQL instance “inst2”, then the clone database name can remain “db1” (the suffix is optional) because the database is cloned to a different instance.
Auto assign mount point or Auto assign volume mount point under path	Choose whether to automatically assign a mount point or volume mount point under a path. Choosing to auto assign a volume mount point under a path enables you to provide a specific directory. The mount points will be created within that directory. Before you choose this option, you must ensure that the directory is empty. If there is a database in the directory, the database will be in an invalid state after the mount operation.

5. In the Location page, select a storage location to create a clone.
6. In the Script page, enter the path and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.

For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

The default script timeout is 60 seconds.

7. In the Schedule page, perform one of the following actions:
 - Select **Run now** if you want to execute the clone job immediately.
 - Select **Configure schedule** when you want to determine how frequently the clone operation should occur, when the clone schedule should start, on which day the clone operation should occur, when the schedule should expire, and whether the clones have to be deleted after the schedule expires.
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. If you want to attach the report of the clone operation performed, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using either the GUI or the PowerShell command `Set-SmSmtServer`.

For EMS, you can refer to [Manage EMS data collection](#)

9. Review the summary, and then click **Finish**.

You should monitor the cloning process using the **Monitor > Jobs** page.

Monitor SQL database 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**.

Cancel SQL resource clone operations

You can cancel clone operations that are queued.


You should be logged in as the SnapCenter Admin or job owner to cancel clone operations.

About this task

- You can cancel a queued clone operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running clone operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued clone operations.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued clone operations of other members while using that role.

Step

Perform one of the following actions:

From the...	Action
Monitor page	<ul style="list-style-type: none"> a. In the left navigation pane, click Monitor > Jobs. b. Select the operation, and click Cancel Job.
Activity pane	<ul style="list-style-type: none"> a. After initiating the clone operation, click  on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click Cancel Job.

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.

Find more information

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

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.