



Configure High Availability

SnapCenter Software 6.0

NetApp
November 06, 2025

This PDF was generated from https://docs.netapp.com/us-en/snapcenter-60/install/concept_configure_snapcenter_servers_for_high_availability_using_f5.html on November 06, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- Configure High Availability 1
 - Configure SnapCenter Servers for High Availability..... 1
 - High availability for the SnapCenter MySQL repository 5

Configure High Availability

Configure SnapCenter Servers for High Availability


To support High Availability (HA) in SnapCenter running either on Windows or on Linux, you can install the F5 load balancer. F5 enables the SnapCenter Server to support active-passive configurations in up to two hosts that are in the same location. To use F5 Load Balancer in SnapCenter, you should configure the SnapCenter Servers and configure F5 load balancer.

You can also configure Network Load Balancing (NLB) to set up SnapCenter High Availability. You should manually configure NLB outside of SnapCenter installation for high availability.

For cloud environment, you can configure high availability using either Amazon Web Services (AWS) Elastic Load Balancing (ELB) and Azure load balancer.

Configure high availability using F5

The workflow image lists the steps to configure SnapCenter Servers for high availability using F5 load balancer. For detailed instruction, refer to [How to configure SnapCenter Servers for high availability using F5 Load Balancer](#).



You must be a member of the Local Administrators group on the SnapCenter Servers (in addition to being assigned to the SnapCenterAdmin role) to use the following cmdlets for adding and removing F5 clusters:

- Add-SmServerCluster
- Add-SmServer
- Remove-SmServerCluster

For more information, refer to [SnapCenter Software Cmdlet Reference Guide](#).

Additional information

- After you install and configure SnapCenter for high availability, edit the SnapCenter desktop shortcut to point to the F5 cluster IP.
- If a failover occurs between SnapCenter Servers and if there is also an existing SnapCenter session, you must close the browser and log on to SnapCenter again.
- In load balancer setup (NLB or F5), if you add a host that is partially resolved by the NLB or F5 host and if the SnapCenter host is not able to reach out to this host, then the SnapCenter host page switches between hosts down and running state frequently. To resolve this issue, you should ensure that both the SnapCenter hosts are able to resolve the host in NLB or F5 host.

- SnapCenter commands for MFA settings should be executed on all the hosts. Relying party configuration should be done in the Active Directory Federation Services (AD FS) server using F5 cluster details. The host level SnapCenter UI access will be blocked after MFA is enabled.
- During failover, the audit log settings will not reflect on the second host. Hence, you should manually repeat the audit log settings on F5 passive host when it becomes active.

Configure high availability using Network Load Balancing (NLB)

You can configure Network Load Balancing (NLB) to set up SnapCenter High Availability. You should manually configure NLB outside of SnapCenter installation for high availability.

For information about how to configure Network Load Balancing (NLB) with SnapCenter refer to [How to configure NLB with SnapCenter](#).

Configure high availability using AWS Elastic Load Balancing (ELB)

You can configure high availability SnapCenter environment in Amazon Web Services (AWS) by setting up two SnapCenter servers in separate availability zones (AZs) and configuring them for automatic failover. The architecture includes virtual private IP addresses, routing tables, and synchronization between active and standby MySQL databases.

Steps

1. Configure virtual private overlay IP in AWS. For information, refer to [Configure virtual private overlay IP](#).
2. Prepare your Windows host
 - a. Force IPv4 being prioritized above IPv6:
 - Location: HKLM\SYSTEM\CurrentControlSet\Services\Tcpip6\Parameters
 - Key: DisabledComponents
 - Type: REG_DWORD
 - Value: 0x20
 - b. Ensure that the fully qualified domain names can be resolved via DNS or via local host configuration to the IPv4 addresses.
 - c. Ensure that you do not have a system proxy configured.
 - d. Ensure that the administrator password is same on both the Windows Server when using a setup without an Active Directory and the servers are not in one domain.
 - e. Add virtual IP on both Windows Servers.
3. Create the SnapCenter cluster.
 - a. Start Powershell and connect to SnapCenter. `Open-SmConnection`
 - b. Create the cluster. `Add-SmServerCluster -ClusterName <cluster_name> -ClusterIP <cluster_ip> -PrimarySCServerIP <primary_ip> -Verbose -Credential administrator`
 - c. Add the secondary server. `Add-SmServer -ServerName <server_name> -ServerIP <server_ip> -CleanUpSecondaryServer -Verbose -Credential administrator`
 - d. Get the high availability details. `Get-SmServerConfig`
4. Create the Lambda function to adjust the routing table in case the virtual private IP endpoint becomes unavailable, monitored by AWS CloudWatch. For information, refer to [Create a Lambda function](#).

5. Create a monitor in CloudWatch to monitor the availability of the SnapCenter endpoint. An alarm is configured to trigger a Lambda function if the endpoint is unreachable. The Lambda function adjusts the routing table to redirect traffic to the active SnapCenter server. For information, refer to [Create synthetic canaries](#).
6. Implement workflow using a step function as an alternative to CloudWatch monitoring, providing smaller failover times. The workflow includes a Lambda probe function to test the SnapCenter URL, a DynamoDB table for storing failure counts, and the Step Function itself.
 - a. Use a lambda function for probing the SnapCenter URL. For information, refer to [Create Lambda function](#).
 - b. Create a DynamoDB table for storing the failure count between two Step Function iterations. For information, refer to [Get started with DynamoDB table](#).
 - c. Create the Step Function. For information, refer to [Step Function documentation](#).
 - d. Test a single step.
 - e. Test the complete function.
 - f. Create IAM Role and adjust permissions to be allowed to execute Lambda function.
 - g. Create schedule to trigger Step Function. For information, refer to [Using Amazon EventBridge Scheduler to start a Step Functions](#).

Configure high availability using Azure load balancer

You can configure high availability SnapCenter environment using Azure load balancer.

Steps

1. Create virtual machines in a scale set using Azure portal. The Azure virtual machine scale set allows you to create and manage a group of load balanced virtual machines. The number of virtual machine instances can automatically increase or decrease in response to demand or a defined schedule. For information, refer to [Create virtual machines in a scale set using Azure portal](#).
2. After configuring the virtual machines, log into each virtual machine in VM set and install SnapCenter Server in both the nodes.
3. Create the cluster in host 1. `Add-SmServerCluster -ClusterName <cluster_name> -ClusterIP <specify the load balancer front end virtual ip> -PrimarySCServerIP <ip address> -Verbose -Credential <credentials>`
4. Add the secondary server. `Add-SmServer -ServerName <name of node2> -ServerIP <ip address of node2> -Verbose -Credential <credentials>`
5. Obtain the high availability details. `Get-SmServerConfig`
6. If required, rebuild the secondary host. `Set-SmRepositoryConfig -RebuildSlave -Verbose`
7. Failover to the second host. `Set-SmRepositoryConfig ActiveMaster <name of node2> -Verbose`

== Switch from NLB to F5 for high availability

You can change your SnapCenter HA configuration from Network Load Balancing (NLB) to use F5 Load Balancer.

Steps

1. Configure SnapCenter Servers for high availability using F5. [Learn more](#).

2. On the SnapCenter Server host, launch PowerShell.
3. Start a session by using the `Open-SmConnection` cmdlet, and then enter your credentials.
4. Update the SnapCenter Server to point to the F5 cluster IP address using the `Update-SmServerCluster` cmdlet.

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](#).

High availability for the SnapCenter MySQL repository

MySQL replication is a feature of MySQL Server that enables you to replicate data from one MySQL database server (master) to another MySQL database server (slave). SnapCenter supports MySQL replication for high availability only on two Network Load Balancing-enabled (NLB-enabled) nodes.

SnapCenter performs read or write operations on the master repository and routes its connection to the slave repository when there is a failure on the master repository. The slave repository then becomes the master repository. SnapCenter also supports reverse replication, which is enabled only during failover.

If you want to use the MySQL high availability (HA) feature, you must configure Network Load Balancer (NLB) on the first node. The MySQL repository is installed on this node as part of the installation. While installing SnapCenter on the second node, you must join to the F5 of the first node and create a copy of the MySQL repository on the second node.

SnapCenter provides the *Get-SmRepositoryConfig* and *Set-SmRepositoryConfig* PowerShell cmdlets to manage MySQL replication.

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](#).

You must be aware of the limitations related to the MySQL HA feature:

- NLB and MySQL HA are not supported beyond two nodes.
- Switching from a SnapCenter standalone installation to an NLB installation or vice versa and switching from a MySQL standalone setup to MySQL HA are not supported.
- Automatic failover is not supported if the slave repository data is not synchronized with the master repository data.

You can initiate a forced failover by using the *Set-SmRepositoryConfig* cmdlet.

- When failover is initiated, jobs that are running might fail.

If failover happens because MySQL Server or SnapCenter Server is down, then any jobs that are running might fail. After failing over to the second node, all subsequent jobs run successfully.

For information about configuring high availability, see [How to configure NLB and ARR with SnapCenter](#).

Copyright information

Copyright © 2025 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.