



Storage virtualization

ONTAP 9

NetApp
March 08, 2024

Table of Contents

- Storage virtualization 1
 - Storage virtualization overview 1
 - SVM use cases 2
 - Cluster and SVM administration 3
 - System Manager insights 3
 - Namespaces and junction points 7

Storage virtualization

Storage virtualization overview

You use *storage virtual machines (SVMs)* to serve data to clients and hosts. Like a virtual machine running on a hypervisor, an SVM is a logical entity that abstracts physical resources. Data accessed through the SVM is not bound to a location in storage. Network access to the SVM is not bound to a physical port.



SVMs were formerly called "vservers." The ONTAP command line interface still uses the term "vserver".


An SVM serves data to clients and hosts from one or more volumes, through one or more network *logical interfaces (LIFs)*. Volumes can be assigned to any data aggregate in the cluster. LIFs can be hosted by any physical or logical port. Both volumes and LIFs can be moved without disrupting data service, whether you are performing hardware upgrades, adding nodes, balancing performance, or optimizing capacity across aggregates.

The same SVM can have a LIF for NAS traffic and a LIF for SAN traffic. Clients and hosts need only the address of the LIF (IP address for NFS, SMB, or iSCSI; WWPN for FC) to access the SVM. LIFs keep their addresses as they move. Ports can host multiple LIFs. Each SVM has its own security, administration, and namespace.

In addition to data SVMs, ONTAP deploys special SVMs for administration:

- An *admin SVM* is created when the cluster is set up.
- A *node SVM* is created when a node joins a new or existing cluster.
- A *system SVM* is automatically created for cluster-level communications in an IPspace.

You cannot use these SVMs to serve data. There are also special LIFs for traffic within and between clusters, and for cluster and node management.



Data accessed through an SVM is not bound to a physical storage location. You can move a volume without disrupting data service.

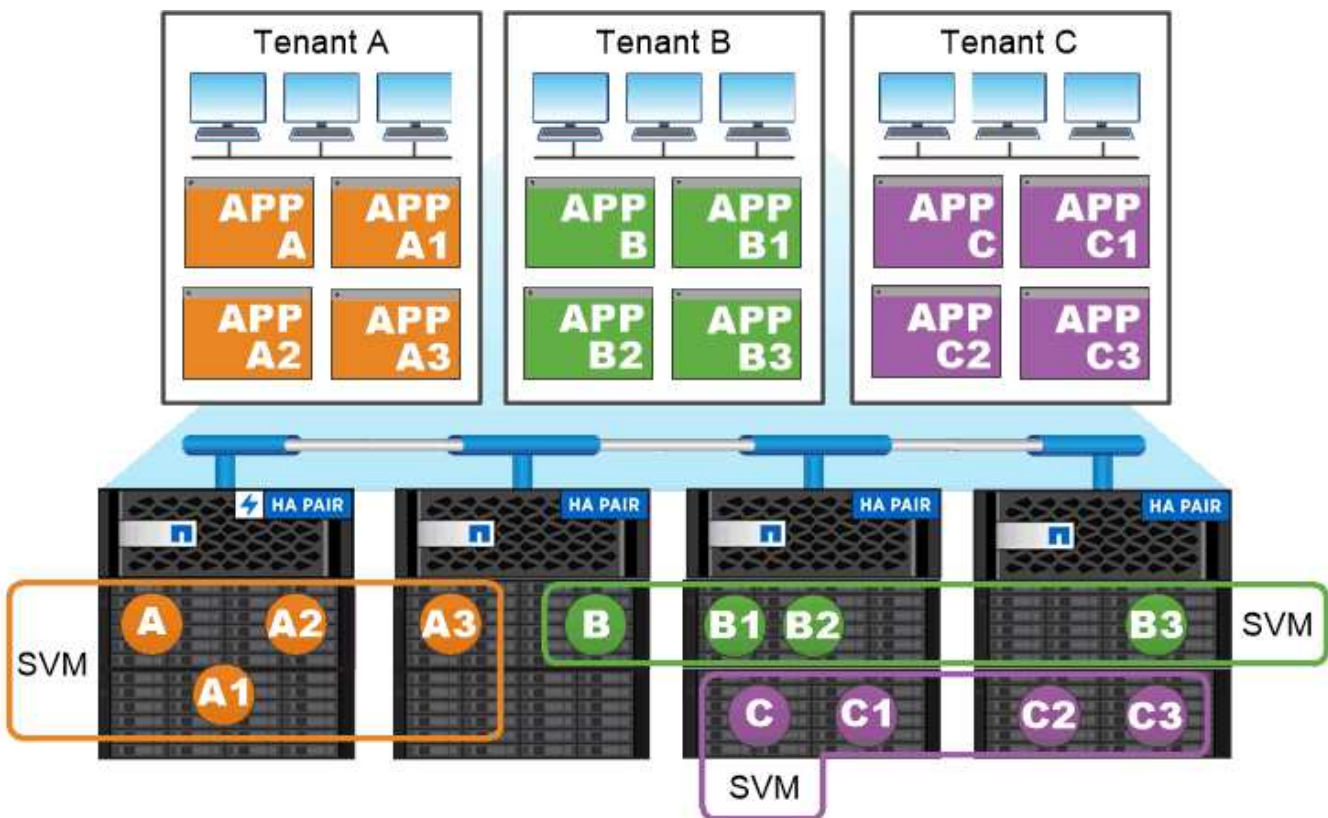
Why ONTAP is like middleware

The logical objects ONTAP uses for storage management tasks serve the familiar goals of a well-designed middleware package: shielding the administrator from low-level implementation details and insulating the configuration from changes in physical characteristics like nodes and ports. The basic idea is that the administrator should be able to move volumes and LIFs easily, reconfiguring a few fields rather than the entire storage infrastructure.

SVM use cases

Service providers use SVMs in secure multitenancy arrangements to isolate each tenant's data, to provide each tenant with its own authentication and administration, and to simplify chargeback. You can assign multiple LIFs to the same SVM to satisfy different customer needs, and you can use QoS to protect against tenant workloads “bullying” the workloads of other tenants.

Administrators use SVMs for similar purposes in the enterprise. You might want to segregate data from different departments, or keep storage volumes accessed by hosts in one SVM and user share volumes in another. Some administrators put iSCSI/FC LUNs and NFS datastores in one SVM and SMB shares in another.



Service providers use SVMs in multitenant environments to isolate tenant data and simplify chargeback.

Cluster and SVM administration

A *cluster administrator* accesses the admin SVM for the cluster. The admin SVM and a cluster administrator with the reserved name `admin` are automatically created when the cluster is set up.

A cluster administrator with the default `admin` role can administer the entire cluster and its resources. The cluster administrator can create additional cluster administrators with different roles as needed.

An *SVM administrator* accesses a data SVM. The cluster administrator creates data SVMs and SVM administrators as needed.

SVM administrators are assigned the `vsadmin` role by default. The cluster administrator can assign different roles to SVM administrators as needed.

Role-Based Access Control (RBAC)

The *role* assigned to an administrator determines the commands to which the administrator has access. You assign the role when you create the account for the administrator. You can assign a different role or define custom roles as needed.

System Manager insights

Beginning with ONTAP 9.11.1, System Manager displays *insights* that help you optimize the performance and security of your system.



To view, customize, and respond to insights, refer to [Gain insights to help optimize your system](#)

Capacity insights

System Manager can display the following insights in response to capacity conditions in your system:

Insight	Severity	Condition	Fixes
Local tiers are lacking space	Remediate risks	One or more local tiers are more than 95% full and quickly growing. Existing workloads might be unable to grow, or in extreme cases, existing workloads might run out of space and fail.	Recommended fix: Perform one of following options. <ul style="list-style-type: none">• Clear the volume recovery queue.• Enable thin provisioning on thick provisioned volumes to free up trapped storage.• Move volumes to another local tier.• Delete unneeded Snapshot copies.• Delete unneeded directories or files in the volumes.• Enable Fabric Pool to tier the data to the cloud.

Applications are lacking space	Needs attention	One or more volumes are more than 95% full, but they do not have autogrow enabled.	<p>Recommended: Enable autogrow up to 150% of current capacity.</p> <p>Other options:</p> <ul style="list-style-type: none"> • Reclaim space by deleting Snapshot copies. • Resize the volumes. • Delete directories or files.
FlexGroup volume's capacity is imbalanced	Optimize storage	The size of the constituent volumes of one or more FlexGroup volumes has grown unevenly over time, leading to an imbalance in capacity usage. If the constituent volumes become full, write failures could occur.	Recommended: Rebalance the FlexGroup volumes.
Storage VMs are running out of capacity	Optimize storage	One or more storage VMs are near their maximum capacity. You will not be able to provision more space for new or existing volumes if the storage VMs reach maximum capacity.	Recommended: If possible, increase the maximum capacity limit of the storage VM.

Security insights

System Manager can display the following insights in response to conditions that might jeopardize the security of your data or your system.

Insight	Severity	Condition	Fixes
Volumes are still in anti-ransomware learning mode	Needs attention	One or more volumes have been in the anti-ransomware learning mode for 90 days.	Recommended: Enable the anti-ransomware active mode for those volumes.

Automatic deletion of Snapshot copies is enabled on volumes	Needs attention	Snapshot auto-deletion is enabled on one or more volumes.	Recommended: Disable the automatic deletion of Snapshot copies. Otherwise, in case of a ransomware attack, data recovery for these volumes might not be possible.
Volumes don't have Snapshot policies	Needs attention	One or more volumes don't have an adequate Snapshot policy attached to them.	Recommended: Attach a Snapshot policy to volumes that don't have one. Otherwise, in case of a ransomware attack, data recovery for these volumes might not be possible.
Native FPolicy is not configured	Best practice	FPolicy is not configured on one or more NAS storage VMs.	Recommended: Configure the FPolicy in NAS storage VMs to control the file extensions that are allowed or not allowed to be written on the volumes in the cluster. This helps prevent attacks from ransomware with known file extensions.
Telnet is enabled	Best practice	Secure Shell (SSH) should be used for secure remote access.	Recommended: Disable Telnet and use SSH for secure remote access.
Too few NTP servers are configured	Best practice	The number of servers configured for NTP is less than 3.	Recommended: Associate at least three NTP servers with the cluster. Otherwise, problems can occur with the synchronization of the cluster time.
Remote Shell (RSH) is enabled	Best practice	Secure Shell (SSH) should be used for secure remote access.	Recommended: Disable RSH and use SSH for secure remote access.
Login banner isn't configured	Best practice	Login messages are not configured either for the cluster, for the storage VM, or for both.	Recommended: Setup the login banners for the cluster and the storage VM and enable their use.
AutoSupport is using a nonsecure protocol	Best practice	AutoSupport is not configured to communicate via HTTPS.	Recommended: It is strongly recommended to use HTTPS as the default transport protocol to send AutoSupport messages to technical support.

Default admin user is not locked	Best practice	Nobody has logged in using a default administrative account (admin or diag), and these accounts are not locked.	Recommended: Lock default administrative accounts when they are not being used.
Secure Shell (SSH) is using nonsecure ciphers	Best practice	The current configuration uses nonsecure CBC ciphers.	Recommended: You should allow only secure ciphers on your web server to protect secure communication with your visitors. Remove ciphers that have names containing "cbc", such as "ais128-cbc", "aes192-cbc", "aes256-cbc", and "3des-cbc".
Global FIPS 140-2 compliance is disabled	Best practice	Global FIPS 140-2 compliance is disabled on the cluster.	Recommended: For security reasons, you should enable Global FIPS 140-2 compliant cryptography to ensure ONTAP can safely communicate with external clients or server clients.
Volumes aren't being monitored for ransomware attacks	Needs attention	Anti-ransomware is disabled on one or more volumes.	Recommended: Enable anti-ransomware on the volumes. Otherwise, you might not notice when volumes are being threatened or under attack.
Storage VMs aren't configured for anti-ransomware	Best practice	One or more storage VMs aren't configured for anti-ransomware protection.	Recommended: Enable anti-ransomware on the storage VMs. Otherwise, you might not notice when storage VMs are being threatened or under attack.

Configuration insights

System Manager can display the following insights in response to concerns about the configuration of your system.

Insight	Severity	Condition	Fixes
Cluster isn't configured for notifications	Best practice	Email, webhooks, or an SNMP trap host is not configured to let you receive notifications about problems with the cluster.	Recommended: Configure notifications for the cluster.

Cluster isn't configured for automatic updates.	Best practice	The cluster hasn't been configured to receive automatic updates for the latest disk qualification package, disk firmware, shelf firmware, and SP/BMC firmware files when they are available.	Recommended: Enable this feature.
Cluster firmware isn't up-to-date	Best practice	Your system doesn't have the latest update to the firmware which could have improvements, security patches, or new features that help secure the cluster for better performance.	Recommended: Update the ONTAP firmware.

Namespaces and junction points

A NAS *namespace* is a logical grouping of volumes joined together at *junction points* to create a single file system hierarchy. A client with sufficient permissions can access files in the namespace without specifying the location of the files in storage. Junctioned volumes can reside anywhere in the cluster.


Rather than mounting every volume containing a file of interest, NAS clients mount an NFS *export* or access an SMB *share*. The export or share represents the entire namespace or an intermediate location within the namespace. The client accesses only the volumes mounted below its access point.

You can add volumes to the namespace as needed. You can create junction points directly below a parent volume junction or on a directory within a volume. A path to a volume junction for a volume named "vol3" might be /vol1/vol2/vol3, or /vol1/dir2/vol3, or even /dir1/dir2/vol3. The path is called the *junction path*.

Every SVM has a unique namespace. The SVM root volume is the entry point to the namespace hierarchy.



To ensure that data remains available in the event of a node outage or failover, you should create a *load-sharing mirror* copy for the SVM root volume.



A namespace is a logical grouping of volumes joined together at junction points to create a single file system hierarchy.

Example

The following example creates a volume named “home4” located on SVM vs1 that has a junction path /eng/home:

```
cluster1::> volume create -vserver vs1 -volume home4 -aggregate aggr1
-size 1g -junction-path /eng/home
[Job 1642] Job succeeded: Successful
```

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.