



ONTAP® 9

SVM Root Volume Protection Express Guide

September 2016 | 215-11189-B0
doccomments@netapp.com

Updated for ONTAP 9.1

 **NetApp®**

Contents

Deciding whether to use this guide	4
SVM root volume protection workflow	5
Creating load-sharing mirror volumes	5
Creating load-sharing mirror relationships	7
Initializing the load-sharing mirror set	8
Updating a load-sharing mirror set	8
SVM root volume restore workflow	10
Restoring the SVM root volume by promoting a load-sharing mirror	10
Where to find additional information	13
Copyright information	14
Trademark information	15
How to send comments about documentation and receive update notifications	16
Index	17

Deciding whether to use the SVM Root Volume Protection Express Guide

This guide describes how cluster administrators can protect the root volume of a storage virtual machine (SVM) by creating load-sharing mirrors on every node of a cluster. This guide also describes how to recover from root volume failures or losses by restoring the SVM root volume from a load-sharing mirror.

Every SVM has a root volume that serves as the entry point to the namespace provided by that SVM. The root volume of any SVM is a FlexVol volume that resides at the top level of the namespace hierarchy and contains directories that are used as mount points, the paths where data volumes are junctioned into the namespace. These directories do not often change.

In the unlikely event that the root volume of the SVM is unavailable, NAS clients cannot access the namespace hierarchy and therefore cannot access data in the namespace. For this reason, it is a NetApp best practice to create a load-sharing mirror for the root volume on each node of the cluster so that the namespace directory information remains available in the event of a node outage or failover.

If the SVM root volume becomes temporarily unavailable, read access to the volume is provided through the load-sharing mirror volumes, and the namespace remains available. If the root volume is permanently unavailable, you can promote one of the load-sharing mirror volumes to enable write access.

You should use this guide if you want to protect the SVM root volume or restore the root volume from a load-sharing mirror volume in the following way:

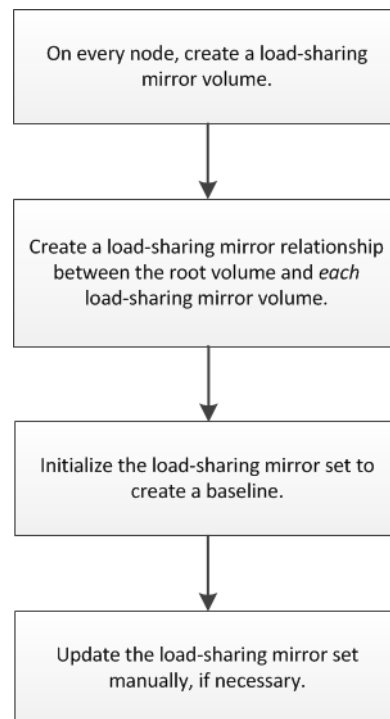
- You are a cluster administrator.
SVM administrators cannot create load-sharing mirrors.
- The root volume you want to protect is for an SVM that uses FlexVol volumes and not an Infinite Volume.
- The SVM has access to at least one data aggregate on each node of the cluster.
- You have NAS clients using CIFS or NFSv3.
SAN client connections (FC, FCoE, or iSCSI) do not depend on the SVM root volume.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.

If these assumptions are not correct for your situation or if you want more conceptual background information, you should see the following resources:

- [*NetApp Technical Report 4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP*](#)

SVM root volume protection workflow

To protect the storage virtual machine (SVM) namespace root volume, you can create a load-sharing mirror volume on every node in the cluster, including the node in which the root volume is located. Then you create a mirror relationship to each load-sharing mirror volume and initialize the set of load-sharing mirror volumes.




Creating load-sharing mirror volumes

For protecting the root volume of a storage virtual machine (SVM), you must create a FlexVol volume on every node of the cluster and designate it as a load-sharing mirror destination.

About this task

A group of load-sharing mirror destination volumes that replicate from the same source volume is referred to as a *load-sharing mirror set*.



When a client requests access to a volume configured with a set of load-sharing mirrors, Data ONTAP directs all read requests to the load-sharing mirror destination volumes. The set of load-sharing mirrors you create for the SVM root volume should therefore include a load-sharing mirror on the same node where the source volume resides.

Steps

1. Identify the data aggregates in the nodes of the cluster by using the `aggr show` command.

Example

The following command displays the list of data aggregates for each node in the ClusterA:

```
ClusterA::> aggr show -root false
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID Status
aggr_1	349.0GB	15.22GB	36%	online	2	ClusterA-01	raid_dp, normal
aggr_2	349.0GB	15.22GB	36%	online	3	ClusterA-02	raid_dp, normal
aggr_3	698.0GB	693.9GB	1%	online	8	ClusterA-03	raid_dp, normal
aggr_4	7.24GB	7.24GB	0%	online	3	ClusterA-04	raid_dp, normal

2. On each node, create a load-sharing mirror volume by using the `volume create` command with the `-type DP` parameter set to **DP**.

Note: The destination volume that you create must be the same size or greater than the SVM root volume.

Example

The following command creates a load-sharing mirror volume for the root volume `vs1_root` of the SVM `vs1` on each node in ClusterA. This follows the NetApp best practice naming convention of naming the load-sharing mirror volumes the same as the SVM name appended with a mirror designation, such as `_m1`, `_m2`, and so on:

```
ClusterA::> volume create -vserver vs1 -volume vs1_m1 -aggregate aggr_1 -size 1gb -state
online -type DP
[Job 4633] Job succeeded: Successful

ClusterA::> volume create -vserver vs1 -volume vs1_m2 -aggregate aggr_2 -size 1gb -state
online -type DP
[Job 4637] Job succeeded: Successful
```

```
ClusterA::> volume create -vserver vs1 -volume vs1_m3 -aggregate aggr_3 -size 1gb -state
online -type DP
[Job 4641] Job succeeded: Successful

ClusterA::> volume create -vserver vs1 -volume vs1_m4 -aggregate aggr_4 -size 1gb -state
online -type DP
[Job 4642] Job succeeded: Successful
```

3. Verify that the volume type is set to **DP** for the load-sharing mirror volumes by using the `volume show` command.

Example

The following command displays the properties of all the volumes in the SVM vs1:

```
ClusterA::>volume show -vserver vs1
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used%
vs1	vs1_root	aggr_1	online	RW	1GB	972.7MB	5%
vs1	vs1_m1	aggr_1	online	DP	1GB	1023MB	0%
vs1	vs1_m2	aggr_2	online	DP	1GB	1023MB	0%
vs1	vs1_m3	aggr_3	online	DP	1GB	1023MB	0%
vs1	vs1_m4	aggr_4	online	DP	1GB	1023MB	0%

Creating load-sharing mirror relationships

You must create a load-sharing mirror relationship between the root volume of the storage virtual machine (SVM) and *each* of the load-sharing mirrors created for the SVM root volume protection.

Steps

1. Create a load-sharing mirror relationship between the SVM root volume and each of the load-sharing mirrors created on every node of the cluster by using the `snapmirror create` command with the `-type` parameter set to **LS**.

Example

The following command creates load-sharing mirror relationships between the SVM root volume vs1_root and four load-sharing mirror volumes vs1_m1, vs1_m2, vs1_m3, and vs1_m4, with an hourly schedule. You do not have to use the `-schedule` parameter again because, because Data ONTAP automatically applies the same schedule to the set of all load-sharing mirrors that share the same source volume.

```
ClusterA::> snapmirror create -source-path vs1:vs1_root -destination-path vs1:vs1_m1
-type LS -schedule hourly
[Job 4623] Job is queued: snapmirror create for the relationship with destination[Job
4623] Job succeeded: SnapMirror: done

ClusterA::> snapmirror create -source-path vs1:vs1_root -destination-path vs1:vs1_m2
-type LS
[Job 4624] Job is queued: snapmirror create for the relationship with destination[Job
4624] Job succeeded: SnapMirror: done

ClusterA::> snapmirror create -source-path vs1:vs1_root -destination-path vs1:vs1_m3
-type LS
[Job 4625] Job is queued: snapmirror create for the relationship with destination[Job
4625] Job succeeded: SnapMirror: done

ClusterA::> snapmirror create -source-path vs1:vs1_root -destination-path vs1:vs1_m4
-type LS
[Job 4626] Job is queued: snapmirror create for the relationship with destination[Job
4626] Job succeeded: SnapMirror: done
```

The type attribute of the load-sharing mirror volumes changes from **DP** to **LS**.

2. Verify that the load-sharing mirror relationships are created and their mirror state is **Uninitialized** by using the `snapmirror show` command.

Example

The following command shows the status of the four load-sharing mirrors that were created for the SVM root volume vs1_root:

```
ClusterA::> snapmirror show -source-volume vs1_root
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Healthy	Last Updated
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m1	Uninitialized	Idle	-	true	-
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m2	Uninitialized	Idle	-	true	-
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m3	Uninitialized	Idle	-	true	-
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m4	Uninitialized	Idle	-	true	-

4 entries were displayed.

Initializing the load-sharing mirror set

You initialize the set of load-sharing mirrors to create a baseline of the root volume of the storage virtual machine (SVM) to the load-sharing mirror destination volumes.

Steps

1. Initialize all the load-sharing mirrors in the set by using the `snapmirror initialize-ls-set` command.

Example

The following command creates a baseline copy of the SVM root volume vs1_root to all of its load-sharing mirrors vs1_m1, vs1_m2, vs1_m3, and vs1_m4:

```
ClusterA::> snapmirror initialize-ls-set -source-path vs1:vs1_root
[Job 4628] Job is queued: snapmirror update-ls-set for source "ClusterA://vs1/vs1_root".
```

2. Verify that the load-sharing mirror relationships are in the **SnapshotMirrored** state by using the `snapmirror show` command.

Example

The following command shows the status of the four load-sharing mirrors that were created for the SVM root volume vs1_root:

```
ClusterA::> snapmirror show -source-volume vs1_root
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Healthy	Last Updated
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m1	SnapshotMirrored	Idle	-	true	-
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m2	SnapshotMirrored	Idle	-	true	-
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m3	SnapshotMirrored	Idle	-	true	-
ClusterA://vs1/vs1_root	LS	ClusterA://vs1/vs1_m4	SnapshotMirrored	Idle	-	true	-

4 entries were displayed.

Updating a load-sharing mirror set

You should update the set of load-sharing mirror volumes to make the changes in the root volume visible to all the clients before the next scheduled update. For example, when a new volume is mounted on the root volume of the SVM, you should update the set of load-sharing mirror volumes.

Step

1. Update all the load-sharing mirror volumes in the set by using the `snapmirror update-ls-set` command.

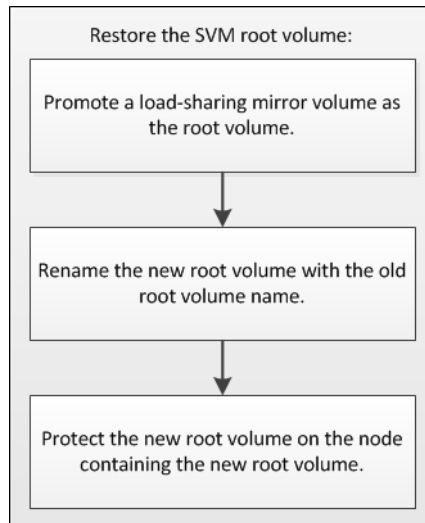
Example

The following command updates all of the load-sharing mirror volumes in the set created for the SVM root volume vs1_root:

```
ClusterA::> snapmirror update-ls-set -source-path vs1:vs1_root  
[Job 4628] Job is queued: snapmirror update-ls-set for source  
"ClusterA://vs1/vs1_root".
```

SVM root volume restore workflow

If the storage virtual machine (SVM) root volume becomes unavailable and you have protected it with a set of load-sharing mirrors, you can promote one of the mirrored volumes and then rename it to take the place of the original SVM source volume.



Restoring the SVM root volume by promoting a load-sharing mirror

If the storage virtual machine (SVM) root volume becomes unavailable, read access is provided through the load-sharing mirror volumes. To enable write access to the SVM root volume, you must promote one of the load-sharing mirror volumes and rename it with the original SVM root volume name.

About this task

Because promoting mirrors is an infrequent operation used only to recover from volume failures or losses, the command can be run only at the advanced privilege level and higher.

The promoted mirror volume has the same attributes as the source SVM root volume. When you recover the SVM root volume from a load-sharing mirror volume, the old source SVM root volume is deleted after the load-sharing mirror volume is promoted.

Steps

1. Promote one of the load-sharing mirror volumes:
 - a. Set the privilege level to **advanced** by using the `set` command.
 - b. Promote a load-sharing mirror as the new SVM root volume by using the `snapmirror promote` command.

Example

The following example promotes the load-sharing mirror `vs1_m2` as the new SVM root volume:

```

ClusterA::> set advanced
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
Do you want to continue? {y/n}: y

ClusterA::*> snapmirror promote -destination-path vs1:vs1_m2

Warning: Promote will delete the offline read-write volume
ClusterA://vs1/vs1_root and replace it with
ClusterA://vs1/vs1_m2. Because the volume is offline, it is
not possible to determine whether this promote will affect other
relationships associated with this source.
Do you want to continue? {y/n}: y
[Job 4629] Job is queued: snapmirror promote of destination "ClusterA://vs1/vs1_m2[Job 4629] Job succeeded: SnapMirror: done

ClusterA::*> volume show -vserver vs1 -volume vs1_m2 -field type

vserver      volume type
-----
vs1          vs1_m2  RW

ClusterA::*> snapmirror show -source-volume vs1_m2

Source      Destination      Mirror      Relationship      Total      Last
Path        Type Path          State         Status           Progress  Healthy Updated
-----
ClusterA://vs1/vs1_m2  LS ClusterA://vs1/vs1_m1  Snapmirrored  Idle           -         true   -
ClusterA://vs1/vs1_m2  LS ClusterA://vs1/vs1_m3  Snapmirrored  Idle           -         true   -
ClusterA://vs1/vs1_m2  LS ClusterA://vs1/vs1_m4  Snapmirrored  Idle           -         true   -

3 entries were displayed.

```

ONTAP makes the load-sharing mirror destination volume a read/write volume, redirects all load-sharing mirrors in the set to the new source volume, and deletes the original source SVM root volume if it is accessible.

Attention: The recovered SVM root volume might not have all of the data that the original root volume had because the load-sharing relationship has a scheduled, asynchronous update and the update might not have occurred recently.

2. From the admin privilege level, rename the new SVM root volume with the old root volume name by using the `volume rename` command.

Renaming the promoted SVM root volume maintains the original naming convention and is a NetApp best practice for NAS-enabled SVMs.

Example

The following example renames the new root volume `vs1_m2` with the old root volume name `vs1_root` from the admin privilege level:

```

ClusterA::*> set admin

ClusterA::> volume rename -vserver vs1 -volume vs1_m2 -newname vs1_root
[Job 4630] Job succeeded: Successful

```

3. Protect the new root volume of the SVM on the node containing the new root volume:
 - a. Identify the aggregate and node containing the new root volume of the SVM by using the `volume show` command.
 - b. Create a load-sharing mirror volume on the node containing the new root volume of the SVM by using the `volume create` command with the `-type` parameter set to **DP**.
 - c. Create a load-sharing mirror relationship between the new root volume of the SVM and the new load-sharing mirror volume by using the `snapmirror create` command with the `-type` parameter set to **LS**.
 - d. Initialize the load-sharing mirror relationship between the new root volume of the SVM and the new load-sharing mirror volume by using the `snapmirror initialize` command.

Example

The following example protects the new root volume `vs1_root` with the new load-sharing mirror volume `vs1_m2`:

```

ClusterA::> volume show -vserver vs1 -volume vs1_root -fields aggregate,node

vserver      volume      aggregate node
-----
vs1          vs1_root    aggr_2      ClusterA_02

ClusterA::> volume create -vserver vs1 -volume vs1_m2 -aggregate aggr_2 -size 1gb -state
online -type DP
[Job 4639] Job succeeded: Successful

ClusterA::> volume show -vserver vs1 -volume vs1_m2 -field type

vserver      volume type
-----
vs1          vs1_m2    DP

ClusterA::> snapmirror create -source-path vs1:vs1_root -destination-path vs1:vs1_m2 -
type LS
[Job 4643] Job is queued: snapmirror create for the relationship with destination[Job
4624] Job succeeded: SnapMirror: done

ClusterA::> snapmirror show -destination-path vs1:vs1_m2 -fields state
source-path      destination-path      state
-----
ClusterA://vs1/vs1_root    ClusterA://vs1/vs1_m2    Uninitialized

ClusterA::> snapmirror initialize -destination-path vs1:vs1_m2
[Job 4656] Job is queued: snapmirror initialize of destination "ClusterA://vs1/vs1_m2".

```

The new load-sharing mirror volume is now part of the load-sharing mirror set.

4. Verify that the new load-sharing mirror relationship is in the **SnapshotMirrored** state by using the `snapmirror show` command.

Example

The following example shows the status of the four load-sharing mirrors that were created for the SVM root volume `vs1_root`:

```

ClusterA::> snapmirror show -source-volume vs1_root

Source      Destination      Mirror      Relationship      Total      Last
Path        Type            Path        State             Status     Progress  Healthy  Updated
-----
ClusterA://vs1/vs1_root    LS    ClusterA://vs1/vs1_m1    SnapshotMirrored  Idle       -         true    -
ClusterA://vs1/vs1_root    LS    ClusterA://vs1/vs1_m2    SnapshotMirrored  Idle       -         true    -
ClusterA://vs1/vs1_root    LS    ClusterA://vs1/vs1_m3    SnapshotMirrored  Idle       -         true    -
ClusterA://vs1/vs1_root    LS    ClusterA://vs1/vs1_m4    SnapshotMirrored  Idle       -         true    -

4 entries were displayed.

```

Where to find additional information

After you successfully protect the root volume of the storage virtual machine (SVM), you can restore the SVM root volume if it becomes inaccessible. You can learn more about load-sharing mirror volumes and other methods of protecting the availability of your data resources.

- [*Data protection*](#)
Describes how to plan and manage disaster recovery and disk-to-disk backup of clustered systems.
- [*NetApp Technical Report 4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP*](#)
Describes information and best practices about configuring replication in clustered Data ONTAP.

Copyright information

Copyright © 1994–2017 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- about this guide
 - deciding whether to use the SVM Root Volume Protection Express Guide [4](#)

B

- backup vaults
 - where to get additional information about [13](#)
- baselines
 - creating for a set of load-sharing mirrors [8](#)

C

- comments
 - how to send feedback about documentation [16](#)

D

- data protection
 - where to get additional information about [13](#)
- documentation
 - how to receive automatic notification of changes to [16](#)
 - how to send feedback about [16](#)

E

- express guides
 - deciding whether to use SVM protection [4](#)

F

- feedback
 - how to send comments about documentation [16](#)
- FlexVol volumes
 - creating for load-sharing [5](#)
- flowcharts
 - SVM root volume promote workflow [10](#)
 - SVM root volume protection workflow [5](#)

I

- information
 - how to send feedback about improving documentation [16](#)

L

- load-sharing mirrors
 - creating a baseline for a set of [8](#)
 - creating destination volume for [5](#)
 - creating relationships for SVM root volume protection [7](#)
 - initializing a set of [8](#)
 - restoring the SVM root volume by promoting [10](#)

- updating relationship [8](#)
- where to get additional information about [13](#)

M

- mirrors
 - creating a baseline for a set of load-sharing [8](#)
 - creating load-sharing destination volume for [5](#)
 - creating load-sharing relationships for SVM root volume protection [7](#)
 - initializing a set of load-sharing [8](#)
 - restoring the SVM root volume by promoting load-sharing [10](#)
 - updating load-sharing relationship [8](#)

N

- namespace
 - SVM protection workflow [5](#)

P

- prerequisites
 - for using SVM protection express guide [4](#)
- protection
 - where to get additional information about data [13](#)

R

- relationships
 - creating load-sharing mirror [7](#)
- root volumes
 - creating load-sharing mirrors for protecting [5](#)
 - creating load-sharing relationships for protecting [7](#)
 - initializing a load-sharing mirror set to create a baseline of [8](#)
 - promote workflow for SVM [10](#)
 - protecting SVM, with load-sharing mirrors [4](#)
 - protection workflow for SVM [5](#)
 - restoring by promoting a load-sharing mirror [10](#)

S

- suggestions
 - how to send feedback about documentation [16](#)
- SVMs
 - creating load-sharing mirror relationships for root volume protection [7](#)
 - protecting with load-sharing mirrors [4](#)
 - restoring the root volume by promoting a load-sharing mirror [10](#)
 - root volume promote workflow [10](#)
 - root volume protection workflow [5](#)

T

- Twitter

how to receive automatic notification of documentation changes [16](#)

U

updating
load-sharing mirror set [8](#)

V

volumes

creating FlexVol, for load-sharing [5](#)
root, promote workflow for SVM [10](#)
root, protection workflow for SVM [5](#)

W

workflows
SVM root volume promote flowchart [10](#)
SVM root volume protection flowchart [5](#)