

SLVM Online Volume Re-configuration



| | |
|--|----|
| Abstract | 2 |
| Overview | 2 |
| Commands..... | 4 |
| vgchange(1M)..... | 4 |
| Changing the mode to exclusive from shared | 4 |
| Changing the mode to shared from exclusive | 4 |
| Procedures | 5 |
| Making a configuration change in an active shared volume group | 5 |
| Determining whether SLVM SNOR is available on the system | 8 |
| SLVM SNOR Messages | 9 |
| vgchange(1M)..... | 9 |
| Required software and patch list..... | 11 |

Abstract

The SLVM Single Node Online volume Re-configuration (SNOR) feature allows changing the configuration of an active shared volume group in a cluster.

SLVM (Shared Logical Volume Manager) is a mechanism that permits multiple systems in a ServiceGuard cluster to share (read/write) disk resources in the form of volume groups.

Using new options in LVM commands, SLVM SNOR allows the system administrator to change the configuration of a shared volume group, and of logical and physical volumes in that volume group, while keeping it active in a single node. Using this procedure, applications on at least one node will be available during the volume group re-configuration.

This document will first provide an overview of the SNOR operation, detail the command changes, and describe the step-by-step procedure and the possible error messages.

Overview

Prior to SLVM SNOR consider a ServiceGuard cluster set up on three nodes. Data is shared in the cluster using common storage. SLVM is used to create and manage volume groups shared across all nodes of the cluster. Applications access the volume group for data read/write operations from disk. To change a shared VG configuration, like creation or modification of logical volumes/physical volumes or a modification in the volume group (VG) attributes, pre-SNOR it is required to:

1. Shutdown the applications using the VG.
2. Deactivate the VG on all nodes.
3. Activate the VG in exclusive mode on one of the nodes.
4. Change the configuration on the VG.
5. Deactivate the VG.
6. If required, generate the map file with the changed configuration, copy it over and re-import it to all the other nodes of the cluster.
7. Activate the VG in shared mode on all the cluster nodes.
8. Restart the applications using the VG.

With the SNOR online volume re-configuration functionality in place, the configuration changes to the VG can be achieved without having to bring down the applications that access the VG. A new command line option to the *vgchange* command (-x) is introduced, which in conjunction with the -a option, allows a single node to switch an active shared VG between shared and exclusive modes.

With the SNOR functionality installed, to change the VG configuration, it is required that:

1. Applications using the VG are scaled down to one node. These applications are still running on that one node
2. The VG is deactivated on all but one node of the cluster.
3. The activation mode of the VG is changed to exclusive in a single node.
4. Change the configuration on the VG.
5. If required, generate the map file with the changed configuration, copy it over and re-import it to all the other nodes of the cluster.
6. Activate the VG in shared mode on the single node and then on all the other cluster nodes.

Application and availability considerations

On the cluster node where the VG stays active, user I/Os will be uninterrupted during and after the activation mode changes from shared to exclusive and back to shared.

On the other cluster nodes where the VG is deactivated, the application must be scaled down to the single cluster node, causing a loss of redundancy and possibly performance degradation for that period.

Commands

vgchange(1M)

A new option (-x) was added to the *vgchange* command to change the activation mode of a shared volume group.

Changing the mode to exclusive from shared

```
vgchange -a e -x vg_name
```

The *vgchange* command provides a new "-x" option that must be used together with the existing "-a e" option to change the activation mode of a shared VG to exclusive mode. Upon success, the VG will be in the same state than if it had been activated in exclusive mode from inactive.

The operation will fail if the volume group is activated on any other node than the local node or if current operations on the volume group do not allow an activation change, e.g. loss of runtime quorum (more than 50% of the last active set of physical volumes lost).

Changing the mode to shared from exclusive

```
vgchange -a s -x vg_name
```

The *vgchange* command provides a new "-x" option that must be used together with the existing "-a s" option to change the activation mode of an active VG from exclusive to shared mode. Upon success, the VG will be in the same state than if it had been activated in shared mode from inactive.

The operation will fail if the volume group is currently inactive or activated in shared mode on the local node, or if current operations on the volume group do not allow an activation change, e.g. loss of runtime quorum.

Procedures

Making a configuration change in an active shared volume group

1. Identify the shared volume group on which a configuration change is required. Call it *vg_shared*.
2. Identify one node of the cluster which is running an application (say SGeRAC) using the shared volume group. Call it *node1*. The applications using the volume group, *vg_shared*, on this node will remain unaffected during the procedure. The cluster application, SGeRAC, needs to be scaled down to the single cluster node *node1*.
3. Deactivate the volume group on all other nodes of the cluster, except *node1*, using the "-n" option to the *vgchange* command.

vgchange -a n vg_shared

Ensure the volume group, *vg_shared*, is now active only on a single cluster node, *node1* by using the *vgdisplay* command on all cluster nodes. The "Status" should show *available* on a single node only.

4. Change the activation mode to exclusive on *node1*

vgchange -a e -x vg_shared

Note: Ensure that none of the mirrored logical volumes in this volume group have Consistency Recovery set to MWC (refer to *lvdisplay(1M)*). Changing the activation mode back to shared will not be allowed in that case, since Mirror Write Cache consistency recovery (MWC) is not valid in volume groups activated in shared mode.

5. Make the desired configuration change on the volume group.

On *node1*, run the required command to change the configuration. For example,

lvextend -m 2 /dev/vg_shared/lvol1

Warning: If *lvsplit* is run on the logical volumes of a volume group in exclusive mode, then changing the mode to shared would result in discarding the bit map used for the subsequent fast *lvmerge* operation. In the absence of the bit map, an *lvmerge* operation causes the entire logical volume to be resynchronized (refer to *lvsplit(1M)*, *lvmerge(1M)*).

6. Export the changes to other cluster nodes if required.

If the configuration change required the creation or deletion of a new logical or physical volume (i.e. any of the following commands were used: *lvcreate*(1M), *lvreduce*(1M), *vgextend*(1M), *vgreduce*(1M), *lvsplit*(1M), *lvmerge*(1M)), then the following sequence of steps is required.

Warning: If one of the above commands was used to change the configuration of the VG, and the mapfile, reflecting the change of configuration, is not imported on all the other nodes, then unexpected errors could occur. For example, errors would be reported if:

- a) A *vgchange* tries to reattach PVs which were removed using *vgreduce*.
- b) A logical volume which was reduced via *lvreduce* is opened.

6.1 From *node1*, export the mapfile for *vg_shared*.

```
vgexport -s -p -m /tmp/vg_shared.map vg_shared
```

6.2 Copy this mapfile */tmp/vg_shared* to all the other nodes of the cluster.

6.3 On the other cluster nodes, export *vg_shared* and re-import it using the new map file.

```
ls -l /dev/vg_shared/group
```

```
crw-rw-rw-  1 root      sys          64 0x050000 Nov 16 15:27 /dev/vg_shared/group
```

Note the minor number (0x050000 in the above example)

```
vgexport vg_shared
```

```
mkdir /dev/vg_shared
```

```
mknod /dev/vg_shared/group c 64 0x050000 (minor number should match with above)
```

```
vgimport -m /tmp/vg_shared.map -s vg_shared
```

Note: The *vgimport*(1M)/*vgexport*(1M) sequence will not preserve the order of physical volumes in the */etc/lvmtab* file. If the ordering is significant due to the presence of active-passive devices, or if the volume group has been configured to maximize throughput by ordering the paths accordingly, the ordering would need to be repeated. In this case it may be better to use the *vgexport* and *vgimport* with the *-f* flag instead and manually edit the file if device paths differ.

7. Change the activation mode back to shared in all the cluster nodes

On *node1*, change the mode back to shared

```
vgchange -a s -x vg_shared
```

On the other cluster nodes, activate *vg_shared* in shared mode

```
vgchange -a s vg_shared
```

8. As a last step, we recommend using *vgcfgbackup* on all nodes to backup the changes made to the volume group.

```
vgcfgbackup vg_shared
```

Determining whether SLVM SNOR is available on the system

SLVM SNOR is available in patch form for HP-UX 11i v1 and v2 and as standard in future HP-UX releases. Patches for LVM commands and kernel components, and the supported ServiceGuard version need to be installed on the system to enable the feature. It will not be available if patches for any of the three components are missing on the system.

- For HP-UX 11i Version 1:

LVM Patches PHKL_33390 and PHCO_33310 (or superseding patches)
ServiceGuard patch PHSS_33834 on SG11.16.

- For HP-UX 11i Version 2:

LVM Patches PHKL_33312 and PHCO_33309 (or superseding patches)
SG 11.17 release or ServiceGuard patch PHSS_33835 on SG11.16.

SLVM SNOR Messages

vgchange(1M)

1. Usage: `vgchange`
{-a Availability -q Quorum [-l] [-p] [-s] [VolumeGroupName...]}
{[-S shareable] -c cluster VolumeGroupName}
{[-P parallel_nomwc_resync_count]}
"x": Illegal option.

This error is reported when the "-x" option to *vgchange* is used on a system that does not have the SNOR LVM Commands patch.

Install the appropriate LVM command patch (see the Required software and patch list).

2. Activation of volume group `"/dev/vg_shared"` denied by another node in the cluster. Request on this system conflicts with Activation Mode on remote system.

This error is reported when the "-x" option to *vgchange* is used on a system that does not have the SNOR ServiceGuard component installed.

Install the ServiceGuard patch as specified in the Required software and patch list.

3. The HP-UX kernel running on this system does not provide this feature.

This error is reported when the "-x" option to *vgchange* is used on a system that does not have SLVM SNOR (LVM kernel component) installed.

Install the appropriate LVM kernel patch (see the Required software and patch list).

4. Activation mode of volume group `"vg_shared"` cannot be changed at this time. Please try later.

This error is reported when the "-x" option to *vgchange* is used on a volume group which has either lost quorum, or is undergoing server recovery or sparing activity.

5. Provide names of active volume group(s) for which reactivation in a different mode is required.

This error is reported when the "-x" option to *vgchange* is used without specifying the volume group name.

6. Volume group `"/dev/vg_shared"` is already active in requested mode.

This error is reported when the "-x" option to *vgchange* is used on a volume group which is already active in the requested mode.

7. To change activation mode, use `lvchange(1M)` to disable the Mirror Write Cache flag for all Logical Volumes in "vg_shared".

This error is reported when the "-x" option to `vgchange` is used, with mirrored logical volumes which have mirror consistency enabled but do not have the Mirror Write Cache flag off, and the requested activation mode is shared (refer `lvchange(1M)`). The Mirror Write Cache is never used when a volume group is activated shared.

8. The volume group `"/dev/vg_shared"` is not active on this system.
Cannot perform requested change.

This error is reported when the "-x" option to `vgchange` is used on a volume group which is not currently active.

9. Activation mode requested for the volume group `"/dev/vg_shared"` conflicts with configured mode.

This error is reported when the "-x" option to `vgchange` is used on a volume group which is not configured in shared mode.

Required software and patch list

- For HP-UX 11i Version 1:

LVM Patches PHKL_33390 and PHCO_33310 (or superseding patches)
ServiceGuard patch PHSS_33834 on SG11.16.

- For HP-UX 11i Version 2:

LVM Patches PHKL_33312 and PHCO_33309 (or superseding patches)
SG 11.17 release or ServiceGuard patch PHSS_33835 on SG11.16.