

Using the vgversion Command to Perform LVM Volume Group Version Migration



Abstract.....	2
Audience.....	2
Overview.....	2
Prerequisites.....	3
Metadata Space Requirements for Version 2.x Volume Groups	3
Unsupported Features in Version 2.x Volume Groups	3
How to Use vgversion	4
Volume Group Migration Flow	5
Volume Group Migration from Version 1.0 to Version 2.x	6
Change in the major number of the LVM device special files	6
Change in the minor number of the LVM special device special files.....	6
Changes to /etc/lvmtab file.....	6
Volume Group Migration from Version 2.0 to Version 2.1	7
Volume Group Migration from Version 2.1 to Version 2.0	7
Recovering from Failures.....	7
Reverting Back to Version 1.0.....	8
Examples.....	9
Migrating a Version 1.0 Volume Group to a Version 2.x Volume Group	9
Verify that all prerequisites are met.....	9
Perform the migration.....	10
Activate the volume group	10
Verify the volume group	10
Migrating a Version 2.0 Volume Group to Version 2.1	11
vgversion Status Messages.....	12
For more information.....	15
Call to action	15

Abstract

Beginning with the HP-UX 11i v3 September 2008 Update release, Logical Volume Manager (LVM) supports three versions of volume groups:

- Version 2.1 (supported on HP-UX 11i v3 September 2008 Update and later)
- Version 2.0 (supported on HP-UX 11i v3 March 2008 Update and later)
- Version 1.0 (supported on all HP-UX 11i v3 releases)

In this document, 2.x refers to volume group Version 2.0 and later.

Version 2.0 or later volume groups enable LVM to increase many limits (for example, the size of logical volumes and physical volumes) that are constrained in Version 1.0 volume groups. To see a comparison of the limits for volume group versions 1.0, 2.0, and 2.1, use the `lvmdm` command (for more information, see `lvmdm(1M)`).

The HP-UX 11i v3 March 2009 Update release provides a new command, `vgversion`, to take advantage of the improvements in volume groups Version 2.0 and later. The `vgversion` command enables in-place volume group migration from Version 1.0 to Version 2.0 or 2.1. The `vgversion` command also supports migration of a volume group from Version 2.0 to Version 2.1 and from Version 2.1 to Version 2.0. This document explains how to use the `vgversion` command to perform in-place volume group migration.

Audience

This document is intended for system administrators, operators, and customers who want to migrate Version 1.0 volume groups to Version 2.x groups, and for those who want to migrate Version 2.x volume groups to Version 2.y volume groups. It is assumed that you have a basic knowledge of LVM.

Overview

The `vgversion` command lets you migrate the following versions of volume groups:

- An existing Version 1.0 volume group to any supported Version 2.x volume group
- An existing Version 2.0 volume group to a Version 2.1 volume group
- An existing Version 2.1 volume group to a Version 2.0 volume group

Currently, volume group version migration is supported only in offline mode; you must deactivate the volume group to perform the operation. The volume group version migration occurs with the data in place.

To migrate the volume group to a new version without the `vgversion` command, you must back up your data, recreate the volume group with the new volume group version, and then copy the data back.

HP recommends using `vgversion` to take advantage of:

- The new features in the newer volume group versions.
- The extended limits supported by the newer volume group versions.

For more information on the new features and the limits supported by various volume group versions, see the [LVM Version 2.0 Volume Groups in HP-UX 11i v3](#) white paper.

Prerequisites

To migrate a volume group from one version to another, you must meet the following prerequisites:

- The volume group must be deactivated during migration.
- The volume group must be cluster unaware (`vgchange -c n`) before changes are made. After the volume group migration, you can then make the volume group cluster aware (`vgchange -c y`).
- No physical volume in the volume group is configured as a cluster lock disk for a high availability cluster.
- All physical volumes belonging to the volume group must be accessible.

For more information, see the individual volume group migration sections.

Metadata Space Requirements for Version 2.x Volume Groups

When compared to a Version 1.0 volume group, a Version 2.x volume group can have more physical volumes and logical volumes. The maximum size supported for logical volumes and physical volumes is increased in Version 2.x volume groups. As a result, a Version 2.x volume group needs more space for metadata than a Version 1.0 volume group. The metadata of a Version 1.0 volume group is stored at the beginning of the physical volume. The major portion of the Version 2.x volume group metadata is stored at the bottom of the physical volume, providing the flexibility of dynamically expanding the physical volume to accommodate more space for the metadata.

The additional space required for the Version 2.x volume group can be provided by using either or both of the following approaches:

- Dynamically extend the physical volume from the storage end and run the `vgversion` command later. The `vgversion` command automatically considers the additional space at the end of the disk to keep the metadata.
- Let the `vgversion` command automatically convert the unused free extents from the end of the physical volume into metadata space.

The free space at the end of the disk is given preference over the unused free extents towards the end of the physical volume.

If the disk capacity increased from the storage side is more than what is required to keep the metadata, the extra space is left unused. To convert the extra space into user extents, use the `vgmodify` command. Alternatively, you can reduce this extra space from storage end by using dynamic LUN contraction.

Unsupported Features in Version 2.x Volume Groups

The following features are supported in Version 1.0 volume groups, but are not supported by Version 2.x volume groups:

- Bootable physical volumes
- Spare physical volumes
- Bad block relocation

You cannot migrate a bootable Version 1.0 volume group to a Version 2.x volume group. If you do not need a bootable volume group, use `vgmodify -B n` and convert all bootable physical volumes to nonbootable physical volumes. Then, retry volume group migration.

If a Version 1.0 volume group has a physical volume that is marked as a spare physical volume, the volume group cannot be migrated. You must change all spare physical volumes to regular physical volumes using `pvchange -z n`.

If any of the physical volumes belonging to a Version 1.0 volume group have relocated bad blocks, then using `vgversion` fails. You cannot migrate this volume group Version 2.x.

A physical volume belonging to a Version 2.x volume group must not have any relocated bad blocks. If any physical volume has relocated bad blocks, you must replace that physical volume with a new disk. For more information on how to replace a disk, see the [When Good Disks Go Bad: Dealing with Disk Failures under LVM](#) white paper.

By default, bad block relocation for all logical volumes of a Version 2.x volume group is set to `NONE`. For more information, see the `-r` option in `lvchange(1M)`. When a Version 1.0 volume group is migrated to Version 2.x, the bad block relocation option for all logical volumes is automatically switched to `NONE`.

How to Use `vgversion`

To use the `vgversion` command, follow these steps:

1. Back up the volume group configuration using `vgcfgbackup -f`. If needed, you can use this command to recover the initial or original volume group configuration. Save the configuration backup file in a known location other than the default location; otherwise, it might be overwritten by a later configuration change.
2. Review the migration operation to check if the migration can go through successfully or if the volume group needs any configuration changes before the migration. The `vgversion -r` (review option) command indicates volume group features that are not supported in the final volume group version and prerequisites that are not fulfilled. You can perform the review operation when the volume group is active. For example:

```
# vgversion -V 2.0 -r -v vg01      # Review the migration operation
```

3. If the existing volume group has any features or functionality that is not supported in the final volume group version, disable these features before migrating the volume group to the new version.
4. Make sure that all the prerequisites are met. The `vgversion` command reports failure if any of the prerequisites are not met.
5. Make sure that all the physical volumes have enough space to write the metadata of the final volume group version. A Version 2.x volume group needs more space for metadata compared to a Version 1.0 volume group. In the review option, `vgversion` reports if any of the physical volumes do not have enough space for metadata. If any physical volume does not have enough space for metadata, follow the steps in the [Metadata Space Requirements for Version 2.x Volume Groups](#) section.
6. After making any changes, enter the `vgversion -r` command again to be sure there are no failures.

```
# vgversion -V 2.0 -r vg01        # Note -r, review mode
```

7. During a maintenance window, deactivate the volume group, make it cluster unaware (if necessary), run `vgversion` to apply the new settings, and activate the volume group.

```
# vgchange -a n vg01              # Deactivate the volume group
# vgversion -V 2.0 vg01           # Migrate the volume group
```

```
# vgchange -a y vg01
```

```
# Activate the volume group
```

Volume Group Migration Flow

The following steps describe the flow of the volume group migration. The flow helps to take necessary action in case of any unexpected failures.

1. The `vgversion` command ensures that all physical volumes are present, then activates the volume group. It then takes a configuration backup of the volume group and saves that as `/etc/lvmconf/vgversion_vg_name/vg_name_initial VG version.conf`. You can use this file to undo the migration. Later, it deactivates the volume group. The following example shows the `vgversion` command output:

```
Performing "vgchange -a r -l -p -s vg01" to collect data
Activated volume group
Volume group "vg1" has been successfully activated.
Old Volume Group configuration for "/dev/vg01" has been saved in
"/etc/lvmconf/vgversion_vg01/vg01_1.0.conf"
Deactivating Volume Group "vg01"
Volume group "vg1" has been successfully deactivated.
```

2. The `vgversion` command creates the configuration backup file corresponding to the final volume group version and saves that as `/etc/lvmconf/vgversion_vg_name/vg_name_final VG version.conf`. The following is the output for this step:

```
New Volume Group configuration for "/dev/vg01" has been saved in
"/etc/lvmconf/vgversion_vg01/vg01_2.1.conf"
```

Note: At this point, the volume group is not modified. Even if the `vgversion` command is killed for any reason, there is no need for any recovery.

3. The `vgversion` command removes the volume group entries from `/etc/lvmtab` in case of Version 1.0 volume groups or from `/etc/lvmtab_p` in case of Version 2.x volume groups. The following is the output for this step:

```
Removing the Volume Group /dev/vg01 from /etc/lvmtab      # Version 1.0 volume groups
Removing the Volume Group /dev/vg01 from /etc/lvmtab_p    # Version 2.x volume groups
```

4. The `vgversion` command removes the `/dev/vg_name` directory and its contents. They are recreated later with new major numbers and minor numbers.
5. The `vgversion` command applies the configuration to all the physical volumes in the reverse order in which they are listed in the `/etc/lvmtab` or `/etc/lvmtab_p` file. If `vgversion` fails to apply the configuration on any of the physical volumes, it undoes all previous changes. The following is the output for this step:

```
Applying the configuration to all Physical Volumes from
"/etc/lvmconf/vgversion_vg1/vg1_2.1.conf"
Volume Group configuration has been restored to /dev/rdisk/disk378
```

Note: If the `vgversion` command is killed while applying the configuration to the disks, the volume group might not be found in either the `/etc/lvmtab` or `/etc/lvmtab_p` file. It is possible that some of the disks have the new configuration and some disks have the old

configuration. HP recommends you run the restore script to recover the volume group. For more information, see [Reverting Back to Version 1.0](#).

6. The `vgversion` command recreates the `/dev/vg_name` directory and device special files of the logical volumes.
7. The `vgversion` command adds the volume group entries to the `/etc/lvmtab_p` file. This completes the volume group version migration.

Volume Group Migration from Version 1.0 to Version 2.x

Using the `vgversion` command, you can convert a Version 1.0 volume group to any supported Version 2.x volume group, but a Version 2.x volume group cannot be migrated to a Version 1.0 volume group. See the Recovering from Failures section for information on how to undo the changes that were made by running the `vgversion` command during the migration process.

If the volume group is deactivated, `vgversion` activates it in both the review option (`-r`) and during the migration operation (`-v`). The `vgversion` command terminates if the volume group is active and the review option is not used (this is necessary as `vgversion` must be in full control of the volume group; it must be able to choose the options used to activate it, and have the ability to deactivate as necessary). Then, `vgversion` calls `vgcfgbackup` to create the configuration backup file corresponding to the initial configuration of the volume group. You can use this configuration backup file to undo the changes performed by the `vgversion` command.

Change in the major number of the LVdevice special files

All the device special files (DSFs) that belong to a Version 1.0 volume group have the major number 64. All the DSFs that belong to a Version 2.x volume group have the major number 128. When a volume group is migrated from Version 1.0 to Version 2.x, all DSFs under the `/dev/vg_name` volume group directory are deleted and recreated with the new major number; the minor numbers remain the same. The names of the DSFs are not altered during volume group migration.

Change in the minor number of the LVspecial device special files

As the major number is changed, there might already be a Version 2.x volume group with the same minor number or index number. In this case, the first unused minor number is used. When the minor numbers do not clash, the same minor number is used. For example, assume that the current setup is as follows:

```
# ll /dev/*/group
crw-r----- 1 root sys 64 0x000000 Jan 6 23:58 /dev/vg00/group
crw-r----- 1 root sys 64 0x010000 Jan 11 20:34 /dev/vg01/group
crw-r----- 1 root sys 128 0x000000 Jan 12 00:19 /dev/vg02/group
crw-r----- 1 root sys 128 0x001000 Jan 12 00:19 /dev/vg03/group
```

While converting `vg01` from Version 1.0 to Version 2.x, minor number 2 is automatically assigned as minor numbers 0 and 1 are in use by `vg02` and `vg03`, respectively.

After conversion the output is as follows:

```
# ll /dev/*/group
crw-r----- 1 root sys 64 0x000000 Jan 6 23:58 /dev/vg00/group
crw-r----- 1 root sys 128 0x002000 Jan 12 22:38 /dev/vg01/group
crw-r----- 1 root sys 128 0x000000 Jan 12 00:19 /dev/vg02/group
crw-r----- 1 root sys 128 0x001000 Jan 12 00:19 /dev/vg03/group
```

Changes to `/etc/lvmtab` file

Configuration information about LVM Version 1.0 volume groups is stored in the `/etc/lvmtab` file; Version 2.x volume groups information is stored in the `/etc/lvmtab_p` file. When a volume group

is migrated from Version 1.0 to Version 2.x, the volume group entry is moved from `/etc/lvmtab` to `/etc/lvmtab_p`.

Volume Group Migration from Version 2.0 to Version 2.1

When a Version 2.0 volume group is migrated to Version 2.1, the device numbers of the device special files do not change. The space required for the metadata for a Version 2.0 volume group is same as a Version 2.1 volume group.

Volume Group Migration from Version 2.1 to Version 2.0

The maximum number of logical volumes supported by a Version 2.0 volume group is 511, and the maximum number of physical volumes supported is 511. The migration of a Version 2.1 volume group to Version 2.0 fails if it has more than 511 logical volumes or more than 511 physical volumes.

If the `lv` minor number of any of the logical volumes of a Version 2.x volume group is more than 511, `vgversion` fails the migration operation even though the total number of logical volumes in the volume group is less than 511 (the maximum number of logical volumes supported by a Version 2.0 volume group). For example:

```
# ll /dev/vg01/lvol*
brw-r----- 1 root  sys  128 0x000101 Jan 18 10:53 /dev/vg01/lvol1
brw-r----- 1 root  sys  128 0x000302 Jan 18 10:54 /dev/vg01/lvol2
```

In the previous output, the `lv` minor number of `lvol1` is `0x101` or 257 and the `lv` number of `lvol2` is `0x302` or 770. Migration of volume group `lvol2` to Version 2.0 will fail because the `lv` minor number is greater than 511.

If the `pvkey` of any physical volume belonging to a Version 2.1 volume group is greater than 511, `vgversion` fails in the migration to Version 2.0. Because the maximum supported physical volumes in a Version 2.0 volume group is 511, the `pvkey` of a physical volume can only range between 0 and 510. It is possible to have a physical volume with a `pvkey` greater than 510 and the total number physical volumes in the volume group less than 511. This happens when more than 511 physical volumes are added to the volume group and later few physical volumes with a `pvkey` less than 510 are reduced from the volume group. There is no direct command to see the `pvkey` of a physical volume. Currently, `lvdisplay -k -v` option shows `pvkey` instead of physical volume path.

Recovering from Failures

This section describes how to recover the volume group in case the `vgversion` command is interrupted.

To simplify the recovery, `vgversion` creates a `vgversion_vg_name_restore` restore script in the `/etc/lvmconf/vgversion_vgname` directory. This shell script is not used by `vgversion`, and is provided for manual restoration of all physical volumes in case of failures or interruptions. HP recommends you use the recover script in case of any failures.

If the `vgversion` command is interrupted while migrating a Version 1.0 volume group to Version 2.1, run the restore script by providing the configuration backup file of the final volume group version configuration. The configuration backup file is `/etc/lvmconf/vgversion_vg_name/vg_name.VG final VG version.conf`

For example:

```
# /etc/lvmconf/vgversion_vg01/vgversion_vg01_restore vg01_2.1.conf
Volume Group configuration has been restored to /dev/rdisk/disk349
```

```

Volume Group configuration has been restored to /dev/rdisk/disk348
The Volume Group already exists in lvmtab/lvmtab_p
Exporting the Volume group
vgexport: Volume group "/dev/vg01" has been successfully removed.
vgimport: Volume group "/dev/vg01" has been successfully created.
Warning: A backup of this volume group may not exist on this machine.
Please remember to take a backup using the vgcfgbackup command after activating
the volume group.
Completed restore of VG /dev/vg01

```

In case of a migration failure, remember the following:

- The review operation of the `vgversion` command does not create or change any configuration backup files.
- During the migration operation, the `vgversion` command creates the latest configuration backup file of the volume group with the current configuration; this can be used to recover the original volume group configuration. Before applying the new configuration on the disks `vgversion` creates the configuration backup file corresponding to the new configuration. This file can be used to update the new configuration on those disks to which the new configuration could not be applied during the `vgversion` operation.
- The configuration files are stored in the `/etc/lvmconf/vgversion_vg_name` directory.
- The `vgversion` command creates two mapfiles named `mapfile_vgname_VG initial VG version` and `mapfile_vgname_VG final VG version` corresponding to initial volume group configuration and final volume group configuration, respectively. These mapfiles contain details of the logical volume names, and are equivalent to those created using the `vgexport -m -f vg_name` command. The restore script uses these mapfiles.
- The `vgversion` command creates a disk file named `vg_name_disks`, containing the list of physical volumes belonging to the volume group. The restore script uses this file.

Reverting Back to Version 1.0

After you convert a Version 1.0 volume group to a Version 2.x volume group, if you need to revert back to the original configuration, follow these steps:

8. Deactivate the volume group, if it is active.
9. Run the restore script by providing the Version 1.0 configuration file as an argument. The configuration files are stored in the `/etc/lvmconf/vgversion_vgname` directory.

Note: Any changes made to the volume group after converting it to Version 2.x are lost if the Version 1.0 configuration is restored on the physical volumes.

The restore script uses `vgcfgrestore` to apply the configuration to all the physical disks. If any of the physical disks are unavailable, the restore script fails. You must rerun it to complete the operation. If the volume group is found in any of the LVM `/etc/lvmtab` or `/etc/lvmtab_p` configuration files, the restore script runs `vgexport vg_name` to export the volume group. Later, the restore script runs `vgimport -m mapfile_vgname_1.0 -f vg_name_disks vg_name` to import the volume group. For example:

```

# /etc/lvmconf/vgversion_vg01/vgversion_vg01_restore_vg01_1.0.conf
vgcfgrestore: Volume group "/dev/vg01" does not exist in the "/etc/lvmtab" file.
Volume Group configuration has been restored to /dev/rdisk/disk348
vgcfgrestore: Volume group "/dev/vg01" does not exist in the "/etc/lvmtab" file.
Volume Group configuration has been restored to /dev/rdisk/disk349
The Volume Group already exists in lvmtab/lvmtab_p
Exporting the Volume group
vgexport: Volume group "/dev/vg01" has been successfully removed.
vgimport: Volume group "/dev/vg01" has been successfully created.
Warning: A backup of this volume group may not exist on this machine.
Please remember to take a backup using the vgcfgbackup command after activating
the volume group.
Completed restore of VG /dev/vg01

```

Examples

The following examples assume the volume group is already activated. This is not a requirement as `vgversion` activates the volume group, if needed. If `vgversion` activates the volume group, `vgversion` deactivates it before completion.

Migrating a Version 1.0 Volume Group to a Version 2.x Volume Group

Verify that all prerequisites are met

Run `vgversion` in review mode. The volume group can be active while running `vgversion` in review mode. If any of the migration prerequisites are not met, the review option lists them. You must correct any prerequisites that are not met. The following example shows the migration of a Version 1.0 volume group with a physical volume marked as spare and a physical volume marked as a bootable PV:

```
# vgversion -V 2.1 -v -r vg01
vgversion: Error: The Physical Volume "/dev/disk/disk348" is configured
as a spare.
vgversion: Error: The Physical Volume "/dev/disk/disk349" is a boot disk.
vgversion: Error: Volume Group version 2.1 does not support bootable
Physical Volumes. Use vgmodify(1M) to make the Physical Volume non-
bootable.
vgversion: Error: Volume Group version 2.1 does not support sparing.
Use pvchange(1M) to change to a regular Physical Volume.
```

```
Review complete. Volume group not modified
```

In this example, the physical volume `/dev/disk/disk348` is configured as spare. Version 2.x volume groups do not support the spare feature. To mark the physical volume as a regular physical volume, use the `pvchange` command when the volume group is active. For example:

```
# pvchange -z n /dev/disk/disk348
Physical volume "/dev/disk/disk348" has been successfully changed.
Volume Group configuration for /dev/vg01 has been saved in
/etc/lvmconf/vg01.conf
```

The `vgversion` review output shows that physical volume `/dev/disk/disk349` is a bootable physical volume. A Version 2.x volume group cannot be booted. To proceed with the migration, you must mark `/dev/disk/disk349` as a nonbootable physical volume. Use the `vgmodify` command to convert the physical volume as nonbootable. In the HP-UX 11i v3 March 2009 Update release and later, `vgmodify` supports the online volume group modification feature. To convert the bootable physical volume to nonbootable when the volume group is active, use the `-a` option while running `vgmodify -B n`. For example:

```
# vgmodify -a -B n vg1 /dev/rdisk/disk378
```

```
Current Volume Group settings:
```

Max LV	255
Max PV	16
Max PE per PV	1016
PE Size (Mbytes)	256
VGRA Size (Kbytes)	176

```
An update to the Volume Group IS required
```

```
New Volume Group settings:
```

```

                Max LV      255
                Max PV      16
                Max PE per PV 1016
                PE Size (Mbytes) 256
                VGRA Size (Kbytes) 176
New Volume Group configuration for "vg1" has been saved in
"/etc/lvmconf/vg1.conf"
Old Volume Group configuration for "vg1" has been saved in
"/etc/lvmconf/vg1.conf.old"
Starting the modification by writing to all Physical Volumes
Applying the configuration to all Physical Volumes from
"/etc/lvmconf/vg1.conf"
Completed the modification process.
New Volume Group configuration for "vg1" has been saved in
"/etc/lvmconf/vg1.conf.old"
Volume group "vg1" has been successfully changed.

```

Perform the migration

After all of the prerequisites are met, deactivate the volume group and run `vgversion`. For example:

```

# vgchange -a n vg01
Volume group "vg01" has been successfully changed.

# vgversion -V 2.1 vg01
Performing "vgchange -a y -l -p -s vg01" to collect data
Activated volume group
Volume group "vg01" has been successfully activated.

Old Volume Group configuration for "/dev/vg01" has been saved in
"/etc/lvmconf/vgversion_vg01/vg01_1.0.conf"
Deactivating Volume Group "vg01"
Volume group "vg01" has been successfully deactivated.

New Volume Group configuration for "/dev/vg01" has been saved in
"/etc/lvmconf/vgversion_vg01/vg01_2.1.conf"
Removing the Volume Group /dev/vg01 from /etc/lvmtab

Applying the configuration to all Physical Volumes from
"/etc/lvmconf/vgversion_vg01/vg01_2.1.conf"
Volume Group configuration has been restored to /dev/rdisk/disk349
Volume Group configuration has been restored to /dev/rdisk/disk348
Creating the Volume Group of version 2.1 with minor number 0x1000.
Adding the Volume Group /dev/vg01 to /etc/lvmtab_p
Original Volume Group Version was 1.0
New Volume Group Version is 2.1
Volume Group version has been successfully changed to 2.1

```

Activate the volume group

```

# vgchange -a y vg01
Activated volume group.
Volume group "vg01" has been successfully changed.

```

Verify the volume group

```

# vgs
--- Volume groups ---
VG Name                /dev/vg01
VG Write Access        read/write
VG Status              available
Max LV                 2047
Cur LV                0

```

Open LV	0
Max PV	2048
Cur PV	2
Act PV	2
Max PE per PV	131072
VGDA	4
PE Size (Mbytes)	128
Total PE	32766
Alloc PE	0
Free PE	32766
Total PVG	0
Total Spare PVs	0
Total Spare PVs in use	0
VG Version	2.1
VG Max Size	32766g
VG Max Extents	262128

This completes the migration of a volume group from Version 1.0 to Version 2.1.

Migrating a Version 2.0 Volume Group to Version 2.1

By default, a Version 2.0 volume group fulfills all the prerequisites to migrate it to Version 2.1.

Deactivate the volume group and run the `vgversion` command as follows:

```
# vgversion -V 2.0 vg01
Performing "vgchange -a y -l -p -s vg01" to collect data
Activated volume group.
Volume group "vg01" has been successfully changed.

Old Volume Group configuration for "/dev/vg01" has been saved in
"/etc/lvmconf/vgversion_vg01/vg01_2.1.conf"

New Volume Group configuration for "/dev/vg01" has been saved in
"/etc/lvmconf/vgversion_vg01/vg01_2.0.conf"
Deactivating Volume Group "vg01"
Volume group "vg01" has been successfully changed.
Removing the Volume Group /dev/vg01 from /etc/lvmtab_p

Applying the configuration to all Physical Volumes from
"/etc/lvmconf/vgversion_vg01/vg01_2.0.conf"
Volume Group configuration has been restored to /dev/rdisk/disk349
Volume Group configuration has been restored to /dev/rdisk/disk348
Adding the Volume Group /dev/vg01 to /etc/lvmtab_p
Original Volume Group Version was 2.1
New Volume Group Version is 2.0
Volume Group version has been successfully changed to 2.0
```

vgversion Status Messages

While using the `vgversion` command you may see status messages similar to the following:

- Performing `vgchange -a y -l -p -s vg01` to collect data
When `vgversion` runs, it internally activates the volume group to ensure that all physical volumes are present and all logical volume opens are blocked. Stale extents are not resynchronized during this activation.
- Old Volume Group configuration for `"/dev/vg01"` has been saved in `"/etc/lvmconf/vgversion_vg01/vg01_1.0.conf"`
The `vgversion` command stored the configuration backup file corresponding to the initial configuration of the volume group. You can use this file to undo the volume group migration operation, if the volume group does not under go any configuration changes after the migration.
- New Volume Group configuration for `"/dev/vg01"` has been saved in `"/etc/lvmconf/vgversion_vg01/vg01_2.1.conf"`
While migrating a volume group to the specified volume group version, `vgversion` created the configuration backup file of the volume group corresponding to the new volume group version.
- Applying the configuration to all Physical Volumes from `"/etc/lvmconf/vgversion_vg01/vg01_2.1.conf"`
The `vgversion` command called `vgcfgrestore` to apply the configuration corresponding to the final volume group version to the physical volumes.
- Removing the Volume Group `/dev/vg01` from `/etc/lvmtab`
While migrating a volume group from Version 1.0 to Version 2.x, `vgversion` removed the volume group entry from the `/etc/lvmtab` file and added the entry to `/etc/lvmtab_p` file.
- Adding the Volume Group `/dev/vg01` to `/etc/lvmtab_p`
After successfully applying the configuration corresponding to the new volume group to the physical volumes, `vgversion` added the volume group entry to `/etc/lvmtab_p` file.
- Error: The Physical Volume `"/dev/disk/disk348"` is a cluster lock disk.
The `vgversion` command cannot continue with the migration because one of the physical volumes belonging to the volume group is a cluster lock disk for a high availability HP Serviceguard cluster. To continue with the migration, modify the cluster configuration so the physical volume is not a cluster lock disk.
- Error: Cannot change the version of a cluster aware Volume Group.
The `vgversion` command failed the migration operation because the volume group is configured as cluster aware. When a volume group is configured as cluster aware and deactivated on the local node, it might be active on another node. In these cases, continuing with migration causes undesirable consequences.
- Error: The Physical Volume `"/dev/disk/disk348"` is a boot disk.
A Version 2.x volume group cannot be a bootable volume group. If any of the physical volumes of a Version 1.0 volume group are marked as bootable, `vgversion` fails the migration. To continue the migration, you must convert all bootable physical volumes to nonbootable by using the `vgmodify -B n` command.

Note: After the migration, you cannot use the volume group as bootable volume group.

- Error: The Physical Volume `"/dev/disk/disk348"` is configured as a spare. Version 2.x volume groups do not support spare physical volumes. If any physical volume of a Version 1.0 volume group is marked as a spare physical volume, `vgversion` fails the migration. To continue with the migration, you must convert all spare physical volumes to regular physical volumes.
- Error: The Physical Volume `"/dev/disk/disk348"` has bad blocks relocated on it. Version 2.x volume groups cannot handle bad blocks on physical volumes. If any physical volumes of a Version 1.0 volume group have relocated bad blocks, `vgversion` fails the migration. You cannot migrate the volume group to Version 2.x.
- The space required for Volume Group version 2.1 metadata on Physical Volume `/dev/disk/disk14` is 7408 KB, but available free space is 3072 KB. 2 free user extents from the end of the Physical Volume `/dev/disk/disk14` will be utilized to accommodate the Volume Group version 2.1 metadata.
The space required for the metadata of the final volume group version is more than the existing free space available. Some (in this case 2) of the unused free user extents at the end of the physical volume are to be converted into metadata space.
- `vgversion`: More space for metadata is required. Take one of the following actions:
 1. Free 1 extents more from the end of the Physical Volume
 2. Increase the disk size by 4336 KB to avoid using the free user extents
 3. Increase the disk size by 240 KB and use the free user extents
 The space required for the metadata of the final volume group version is more than the existing free space available, and `vgversion` cannot acquire the required space by converting the free user extents at the end of the physical volume into metadata space or by using the free space available beyond the end of the user data. In this case, you have three choices to create space for metadata.
- A version 2.x Volume Group with the minor number `0x1000` already exists
The minor number is already used by a Version 2.x volume group. The `vgversion` command selects the first available minor number. While converting a Version 1.0 volume group to a Version 2.x volume group, `vgversion` tries to retain the volume group minor number as is.
- Error: Failed to update the `/etc/lvmtab_p` with the Volume Group `"/dev/vg01"`
The addition of volume group entries to the `/etc/lvmtab_p` file failed. The physical volumes belonging to the volume group are vulnerable to data loss. HP recommends you run the restore script after fixing the reason for the error. When a Version 1.0 volume group is converted to Version 2.x, the `vgversion` command removes the volume group entry from `/etc/lvmtab` and adds it to `/etc/lvmtab_p`.
- Error: Could not find any free minor number for volume group version 2.1
The `vgversion` command did not find a free minor number to create the volume group `group` file because all the minor numbers are used by the existing Version 2.x volume groups. This can occur when the maximum supported number of volume groups were already created.
- Error: The number of Logical Volumes in the Volume Group `"/dev/vg01"` is 520. But Volume Group version 2.0 supports up to 511 Logical Volumes.
The minor number of the logical volume is greater than what is supported by the requested final volume group version; `vgversion` failed the migration.

- Error: The Physical Volume `"/dev/disk/disk348 "` has the `pvkey` 512, which is not supported by volume group version 2.0

The `pvkey` of a physical volume of a Version 2.x volume group is more than 510, even though the total number of physical volumes in the volumes group is less than 511; `vgversion` failed the migration.

- Error: The Logical Volume `"/dev/vg01/lvol512"` has the `lv` number 512, which is not supported by volume group version 2.0

The `lv` number of a logical volume is more than 511; `vgversion` failed the migration. To continue with the migration, delete those logical volumes and recreate them.

- Error: Failed to apply the new configuration on all Physical Volumes of the Volume Group.

The `vgversion` command calls `vgcfgrestore` to apply the new configuration on the physical volumes. While applying the configuration if any of the physical volumes becomes unavailable, `vgversion` tries to revert back to the initial volume group version metadata.

- Error: Kernel does not support version change operation.

The `vgversion` command does not support migration to the requested volume group.

For more information

To learn more about LVM and HP-UX system administration, see the following documents on the HP documentation website (<http://docs.hp.com>):

- *HP-UX System Administrator's Guide: Logical Volume Configuration*
<http://docs.hp.com/en/5992-4589/5992-4589.pdf>
- LVM White Papers:
 - *LVM Volume Group Dynamic LUN expansion (DLE) / vgmodify* (September 2008)
http://docs.hp.com/en/vgmodify/vgmodify_wp_modified_final.pdf
 - *LVM Supported Limits*
http://docs.hp.com/en/6054/Limits_wp.pdf
 - *LVM Migration from Legacy to Agile Naming Model*
http://docs.hp.com/en/LVMmigration1/LVM_Migration_to_Agile.pdf
 - *LVM New Features in HP-UX 11i v3*
http://docs.hp.com/en/LVM-11iv3features/LVM_New_Features_11iv3_final.pdf
 - *LVM Online Disk Replacement (LVM OLR)*
http://docs.hp.com/en/7161/LVM_OLR_whitepaper.pdf
 - *LVM Volume Group Quiesce/Resume*
http://docs.hp.com/en/lvm-quiesce/LVM_Quiesce-Resume.pdf
 - *SLVM Single-Node Online Reconfiguration (SLVM SNOR)*
http://docs.hp.com/en/7389/LVM_SNOR_whitepaper.pdf
 - *When Good Disks Go Bad: Dealing with Disk Failures under LVM*
http://docs.hp.com/en/5991-1236/When_Good_Disks_Go_Bad_WP.pdf

Call to action

HP welcomes your input. Please give us comments about this white paper, or suggestions for LVM or related documentation, through our technical documentation feedback website:

<http://docs.hp.com/en/feedback.html>

© 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

5992-6592 March 2009

