

reduxio

Reduxio Best Practices for
BackDating™ of LVM

For more information, refer to Reduxio website at <http://www.reduxio.com>.
If you have comments about this documentation, submit your feedback to docs@reduxio.com.

© 2016 Reduxio Systems Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of Reduxio.

Reduxio™, the Reduxio logo, NoDup™, BackDating™ and Tier-X™ are trademarks or registered trademarks of Reduxio in the United States and/or other countries.

Linux is a registered trademark of Linus Torvalds.

Windows is a registered trademark of Microsoft Corporation.

UNIX is a registered trademark of The Open Group.

ESX and VMWare are registered trademarks of VMWare, Inc.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

The Reduxio system hardware, software, user interface and/or information contained herein is Reduxio Systems Inc. proprietary and confidential. Any and all rights including all intellectual property rights associated therewith are reserved and shall remain with Reduxio Systems Inc. Rights to use, if any, shall be subject to the acceptance of the End User License Agreement provided with the system.

Information in this document is subject to change without notice.

Reduxio Systems, Inc.
111 Pine Avenue
South San Francisco, CA, 94080
United States
www.reduxio.com

Contents

Overview	4
Introduction.....	4
Software Configuration.....	4
LVM.....	4
LVM Structure	4
Setting up LVM on Reduxio.....	4
Recovering LVs using Reduxio BackDating	5
Conclusion	7
References.....	7

Overview

Introduction

The Linux Logical Volume Manager – LVM, is a device mapper target that provides logical volume management for the Linux kernel. Most modern Linux distributions provide built-in support for LVM, typically creating their root file system on an LVM volume.

This tech note describes best practices of recovering and cloning LVM devices using Reduxio BackDating™.

Software Configuration

LVM and Reduxio BackDating were tested using Red Hat Enterprise 7.2 and Ubuntu Server 14.04.01.

LVM

LVM Structure

An LVM configuration consists of the following layers:

PV (Physical Volume)	A physical device (LUN), i.e. Reduxio volume.
VG (Volume Group)	A logical storage pool. A VG is configured on one or more PVs.
LV (Logical Volume)	A logical volume. An LV is configured on a VG. File systems are created on LVs.

Each object in LVM contains a unique identifier (UUID). Each PV contains an LVM header. Stored in the first megabyte of the PV, it contains the complete copy of the entire volume group's layout, including the UUIDs of all other PVs, the UUIDs of all LVs and the mapping between the logical extents (LE) of LVs to the physical extents (PE) stored on PVs.

Setting up LVM on Reduxio

To set up a logical volume on a Reduxio volume, first create a volume, assign it to the Linux host and rescan devices. This will result in a multipath device – in this case `/dev/mpatha`.

1. Validate that a multipath device is ready.

```
root@moon:~# multipath -ll
mpatha (36f4032f0003a000000000000000000093) dm-0 REDUXIO,TCAS
size=100G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|  |- 6:0:0:1 sdc 8:32 active ready running
|  `-- 5:0:0:1 sdd 8:48 active ready running
`--+- policy='round-robin 0' prio=10 status=enabled
    |- 3:0:0:1 sdb 8:16 active ready running
    `-- 4:0:0:1 sde 8:64 active ready running
```

2. Create a PV on the device.

```
root@moon:~# pvcreate /dev/mapper/mpatha
Physical volume "/dev/dm-0" successfully created
```

```

root@moon:~# pvdisplay
"/dev/mapper/mpatha" is a new physical volume of "100.00 GiB"
--- NEW Physical volume ---
PV Name           /dev/mapper/mpatha
VG Name
PV Size           100.00 GiB
Allocatable       NO
PE Size           0
Total PE          0
Free PE           0
Allocated PE      0
PV UUID           qCNoc0-hdfV-xmGP-0NX5-qcCv-B1h2-egrdfE

```

3. Create a VG on the PV (the reference to the PV is still /dev/mapper/mpatha)

```

root@moon:~# vgcreate vg1 /dev/mapper/mpatha
volume group "vg1" successfully created

```

4. Create an LV on the volume group and create a file system on it.

```

root@moon:~# lvcreate --name lv1 --size 40G vg1
root@moon:~# mkfs -t ext4 /dev/vg1/lv1

```

5. /dev/vg1/lv1 can now be mounted and used.

Recovering LVs using Reduxio BackDating

When cloning a Reduxio volume and attaching it to the same host, the LVM will typically identify the PV clone as a duplicate device, since it contains the same UUID as the original PV. This behavior is demonstrated in the following set of commands.

1. Before attaching the clone, validate the entire chain of LV-VG-PV.

```

root@moon:~# lvdisplay
--- Logical volume ---
LV Path           /dev/vg1/lv1
LV Name           lv1
VG Name           vg1
LV UUID           ycemmm-F2om-OpTb-uoTz-tGbp-P0YN-GAc7Q3
LV Write Access   read/write
LV Creation host, time lvm, 2016-05-10 12:25:09 +0300
LV Status         available
# open           1
LV Size           40.00 GiB
Current LE        10240
Segments         1
Allocation        inherit
Read ahead sectors    auto
- currently set to 256
Block device      252:2

```

```

root@moon:~# vgdisplay
--- Volume group ---
VG Name           vg1
System ID
Format            lvm2
Metadata Areas    1
Metadata Sequence No 2
VG Access         read/write
VG Status         resizable
MAX LV            0
Cur LV           1
Open LV           1

```

```

Max PV          0
Cur PV         1
Act PV         1
VG Size        100.00 GiB
PE Size        4.00 MiB
Total PE       25599
Alloc PE / Size 10240 / 40.00 GiB
Free PE / Size 15359 / 60.00 GiB
VG UUID        Io9ECT-gKm2-xaFC-PIqH-v7H0-re9C-ELALgd

```

```

root@moon:~# pvdisk
--- Physical volume ---
PV Name        /dev/mapper/mpatha
VG Name        vg1
PV Size        100.00 GiB / not usable 4.00 MiB
Allocatable    yes
PE Size        4.00 MiB
Total PE       25599
Free PE        15359
Allocated PE   10240
PV UUID        qCNoc0-hdfV-xmGP-0NX5-qcCv-B1h2-egrdfE

```

2. Clone the volume storing /dev/mapper/mpatha and assign it to the host.

```

# ssh rdxadmin@reduxio1 volumes clone pv1 -name pv1clone
# ssh rdxadmin@reduxio1 volumes assign pv1clone -host moon

```

3. Rescan the SCSI bus.

```

# sudo apt-get install scsitools
# sudo rescan-scsi-bus

```

4. LVM identifies and reports a duplicate PV and incorrectly switches to the clone PV.

```

root@moon:~# multipath -ll
mpathb (36f4032f0003a00000000000000000095) dm-3 REDUXIO,TCAS
size=100G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
|  |- 5:0:0:3 sdl 8:176 active ready running
|  `-- 6:0:0:3 sdm 8:192 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
   |- 3:0:0:3 sdj 8:144 active ready running
   `-- 4:0:0:3 sdk 8:160 active ready running
mpatha (36f4032f0003a00000000000000000093) dm-0 REDUXIO,TCAS
size=100G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
|  |- 5:0:0:1 sdd 8:48 active ready running
|  `-- 6:0:0:1 sdc 8:32 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
   |- 3:0:0:1 sdb 8:16 active ready running
   `-- 4:0:0:1 sde 8:64 active ready running
root@moon:~# pvdisk
Found duplicate PV qCNoc0hdfVxmGP0NX5qcCvB1h2egrdfE: using /dev/mapper/mpathb not
/dev/mapper/mpatha
--- Physical volume ---
PV Name        /dev/mapper/mpathb
VG Name        vg1
PV Size        100.00 GiB / not usable 4.00 MiB
Allocatable    yes
PE Size        4.00 MiB
Total PE       25599
Free PE        15359
Allocated PE   10240
PV UUID        qCNoc0-hdfV-xmGP-0NX5-qcCv-B1h2-egrdfE

```

Note that in some distributions the duplicate PV error may not be displayed. However, it is still important to perform the following steps to prevent an incorrect import of a clone volume. Immediately after assigning the clone, use the `vgimportclone` command to configure the clone PV as a separate volume group. Do not run any command such as `pvdisplay` shown above. `vgimportclone` must be executed first to import the PV as a clone type.

```
root@moon:~# vgimportclone --basevgname vg1_backdate1 /dev/mapper/mpathb
WARNING: Activation disabled. No device-mapper interaction will be attempted.
Physical volume "/tmp/snap.0qicqndm/vgimport0" changed
1 physical volume changed / 0 physical volumes not changed
WARNING: Activation disabled. No device-mapper interaction will be attempted.
Volume group "vg1" successfully changed
Volume group "vg1" successfully renamed to "vg1_backdate1"
Reading all physical volumes. This may take a while...
Found volume group "vg1_backdate1" using metadata type lvm2
Found volume group "vg1" using metadata type lvm2
root@moon:~# pvscan
PV /dev/mapper/mpathb   VG vg1_backdate1   lvm2 [100.00 GiB / 60.00 GiB free]
PV /dev/mapper/mpatha  VG vg1              lvm2 [100.00 GiB / 60.00 GiB free]
Total: 2 [199.99 GiB] / in use: 2 [199.99 GiB] / in no VG: 0 [0  ]
root@moon:~# vgscan
Reading all physical volumes. This may take a while...
Found volume group "vg1_backdate1" using metadata type lvm2
Found volume group "vg1" using metadata type lvm2
root@moon:~# lvscan
ACTIVE          '/dev/vg1_backdate1/lv1' [40.00 GiB] inherit
ACTIVE          '/dev/vg1/lv1' [40.00 GiB] inherit
```

5. The clone LV can now be mounted and used independently from the original LV.

Conclusion

LVM-based environments can greatly benefit from the Reduxio BackDating capability. Using best practice procedures, LVM logical volumes can be safely cloned and reverted to any second in the available history of the respective Reduxio volumes.

References

Reduxio Documentation

- *Reduxio Administration Guide*

Linux Documentation

- [Red Hat Enterprise Linux 7 - LVM Administrator Guide](#)