RHEL 6 w/ ASM Devices

Ben Breard, RHCA
Solutions Architect, Red Hat
bbreard@redhat.com

Agenda

- ASMlib Components and Usage with RHEL
- UDEV intro
- Configuring SCSI devices for ASM
- Configuring multipath devices for ASM
- Additional Resources

ASMLib consists of the following components:

- kmod-oracleasm

  An open source (GPL) kernel module package

- oracleasm-support

  An open source (GPL) utilities package

- oracleasmlib

  A closed source (proprietary) library package

ASM features and functionality are available without ASMLib.

The use of ASMLib does not affect database performance.

ASMLib Pros

- Well documented and recommended by Oracle.
- Many Oracle DBAs and SysAdmins are trained in how to use ASM with ASMLib and are comfortable with this environment.
- Optimized for database applications via direct and async I/O provided by the ASMLib kernel driver.

Note: Red Hat Enterprise Linux kernel 2.6.X supports optimized direct and async I/O for all supported filesystem EXT3/4, XFS, GFS and NFS.

ASMLib Cons

- Requires ASMLib kernel driver that is not included in the mainline Linux kernel.
- ASMLib delivers no known performance benefits.
- Red Hat Enterprise Linux with ASMLib does not have government security certification.
- ASMLib is not compatible with SELinux.
- Non POSIX system calls for device access lead to issues with trouble-shooting and performance monitoring.

ASMlib on RHEL 6.4

- kmod-oracleasm is available via Supplementary repository
- oracleasm-support & oracleasmlib are available from Oracle.
- For setup details visit:
  - Red Hat's Knowledge Article 315643

    https://access.redhat.com/knowledge/solutions/315643
  - Oracle's ASMlib on RHEL 6 downloads

    http://www.oracle.com/technetwork/server-storage/linux/asmlib/rhel6-1940776.html

UDEV - Userspace Device Management

UDEV - Introduction

- Native device manager for the Linux kernel.
- Persistent storage can easily be configured with the proper ownership and permissions for ASM.
- Rules look complicated, but the format is simple and well documented.

UDEV in RHEL 6 – What you need to know

- Rules are located in /etc/udev/rules.d/ and /lib/udev/rules.d/
- Changes are monitored via inotify.
- Rules are parsed in alphanumeric order.
  - 10-XX before 10-XY before 20-XX, etc.
- In most cases it's not necessary to reload the config after a rule change. This can be forced via:

```
udevadm control --reload-rules
```

UDEV – Rules

- Rules are built from comma separated key-value pairs.
- **Match keys** are conditions used to identify the device which the rule is acting upon.
- When *all* match keys in a rule correspond to the device being handled, then the rule is applied and the actions of the **assignment keys** are invoked.
- Every rule should consist of at least one match key and at least one assignment key.

Example: KERNEL=="sdb", NAME="my_spare_disk"

Match key

Assignment key

Source: /usr/share/doc/udev-147/writing_udev_rules/index.html

UDEV – ASM Example Rule:

KERNEL=="sd?1", BUS=="scsi", PROGRAM=="/sbin/scsi_id --whitelisted --replace-whitespace /dev/$parent", RESULT=="1ATA_OCZ-NOCTI_OCZ-G0L319451R4Z671G", NAME="asm-disk1", OWNER="oracle", GROUP="dba", MODE="0660"

- **Match first partition on any SCSI disk attached to the SCSI bus:** KERNEL=="sd?1", BUS=="scsi"
- **Execute this program:** PROGRAM=="/sbin/scsi_id --whitelisted --replace-whitespace /dev/$parent"
- **If the scsi_id returns:** RESULT=="1ATA_OCZ-NOCTI_OCZ-G0L319451R4Z671G"
- **Create /dev/asm-disk1, owned by oracle:dba 0660:**

  NAME="asm-disk1", OWNER="oracle", GROUP="dba", MODE="0660"

UDEV – Identifying Storage

- scsi_id can always be used to generate a match key, but it might not be the best option.
- /lib/udev/rules.d/60-persistent-storage.rules collects the WWIDs for any attached storage.
- For performance reasons, consider creating rules that run after 60-*, and reference ENV{ID_SERIAL} or ENV{DM_UUID} instead of querying the storage a second time.
- SCSI devices can be identified by ENV{ID_SERIAL}
- dm-multipath devices can be identified by ENV{DM_UUID}
- Query via:

```
# udevadm info --query=all --name=/dev/sdb | grep ID_SERIAL
or
# udevadm info --query=all –name=/dev/mapper/my_vgg-data | grep DM_UUID
```

UDEV: Simple SCSI Example

1) Query the storage

Identify the WWID via udevadm:

```
# udevadm info --query=all --name=/dev/sdb | grep ID_SERIAL
```

Returns:

```
E: ID_SERIAL=OCZ-NOCTI_OCZ-G0L319451R4Z671G
E: ID_SERIAL_SHORT=OCZ-G0L319451R4Z671G
```

Note: in this example we will ignore ID_SERIAL_SHORT

2) Create a Rule

- Create /etc/udev/rules.d/99-oracle-asmdevices.rules, and add the rule(s) for any current or added ASM storage.

Rules should look similar to:

ENV{ID_SERIAL}=="<wwid>", ENV{DEVTYPE}=="partition", SYMLINK+="asm-disk1", OWNER="oracle", GROUP="dba", MODE="0660"

Note: for each device, replace wwid w/ the output from udevadm

3) Test/Apply your Rule

udevadm test /sys/block/sdb/sdb1

And confirm the rule had the desired effect:

ls -lh /dev/asm-disk1

brw-rw----. 1 oracle dba 8, 17 Mar 11 20:56 /dev/asm-disk1 -> /dev/sdb1

UDEV – Repeatable Process for Adding Storage

1. Add the lun(s)
2. Identify the ID_SERIAL(s)
3. Create UDEV rules for additional luns(s)
4. Partition the lun(s).
5. If the ENV{ID_SERIAL}=="<wwid>" syntax was used, the UDEV rules will immediately take effect. Otherwise run:

```
# udevadm test /sys/block/sdb/sdb1
```

UDEV: Simple Device-Mapper-Multipath Example

1) Identify and Query the storage

Identify the dm- device number via:

```
# ls -l /dev/mapper/mpathb
lrwxrwxrwx. 1 root root 7 Mar 19 08:58 /dev/mapper/mpathb-> ../dm-2
```

Identify the UUID via:

```
udevadm info --query=all --name=/dev/mapper/mpathb | grep DM_UUID
```

2) Create a Rule

Create /etc/udev/rules.d/99-oracle-asmdevices.rules. Add the rule(s) for any current or added storage.

Rules will look similar to this example:

ENV{DM_UUID}=="<wwid>", ENV{DEVTYPE}=="disk", SYMLINK+="asm-disk1", OWNER="oracle", GROUP="dba", MODE="0660"

- Note: for each device, replace %wwid w/ the output from udevadm
- SYMLINK is not required if it's desired to change permissions directly on the dm-* device.

3) Test/Apply your Rule

udevadm test /sys/block/dm-2

And confirm the rule had the desired effect:

ls -lh /dev/asm-disk2

brw-rw----. 1 oracle dba 8, 17 Mar 11 20:56 /dev/asm-disk2 -> /dev/dm-2

UDEV Tips

- Avoid writing rules from scratch.
  - Consider prestaging commented out rule files that will reduce configuration errors and administrative time. Can use RHN Satellite or Puppet.
- The previous examples create symlinks to the actual device. This is preferred so the disk in the kernel messages can be easily identified.
- To rename devices simply replace “SYMLINK+=” with “NAME=” in the UDEV rule(s). (not recommended)

KERNEL=="sd?1", BUS=="scsi", PROGRAM=="/sbin/scsi_id --whitelisted --replace-whitespace /dev/$parent", RESULT=="1ATA_OCZ-NOCTI_OCZ-G0L319451R4Z671G", **NAME=**"asm-disk1", OWNER="oracle", GROUP="dba", MODE="0660"

Available Resources

RHEL 6 Reference Architecture:
https://access.redhat.com/knowledge/articles/216813

How to replace ASMlib with UDEV:

https://access.redhat.com/knowledge/articles/216353

ASMlib on RHEL 6:
https://access.redhat.com/knowledge/solutions/315643

Configuring ASM storage from Oracle:
http://docs.oracle.com/cd/E11882_01/install.112/e17212/storage.htm#CDEBFDEH

redhat®