

ATOL: Filesystems and Their Management

Marek Grác
xgrac@fi.muni.cz

Red Hat Czech s.r.o. / Faculty of Informatics, Masaryk University

Advanced Topics of Linux Administration

Partitions

- ▶ Split disk space into logically independent sections
- ▶ MBR (Master Boot Record)
 - ▶ max. 4 primary partition, extended partitions
 - ▶ max. addressable size is 2TiB (one HDD)
- ▶ EFI GPT (GUID Partition Table)
 - ▶ EFI is Intel's attempt to replace BIOS
 - ▶ max. addressable size is in zeta-bytes
 - ▶ unsupported by some OS, problems with some tools even in Linux

Creating partitions

- ▶ *fdisk*, *cdisk*, *parted* – view and manage partition tables
- ▶ List partition tables from command line
- ▶ *partprobe* – inform the OS of partition table changes
- ▶ `cat /proc/partitions`

Making Filesystems

- ▶ *mkfs*
- ▶ *mkfs.ext2*, *mkfs.ext3*, *mkfs.msdos*
- ▶ Specific filesystem utilities can be called directly
 - ▶ `mke2fs [options] device`

Filesystem Labels

- ▶ Alternate way to refer to devices
- ▶ Device independent
 - ▶ `e2label devfile [fslabel]`
 - ▶ `mount [options] LABEL=fslabel mountpoint`
- ▶ *blkid* – used to see labels and filesystems type of all devices

Mount Points and `/etc/fstab`

- ▶ Configuration of the filesystem hierarchy
- ▶ Used by *mount*, *fsck* and other programs
- ▶ Maintains the hierarchy between system reboots
- ▶ May use filesystem volume labels in the device field
- ▶ The `mount -a` command can be used to mount all filesystems listed in `/etc/fstab`
- ▶ Option `_netdev` should be used for network devices (NFS, iSCSI, ...)

Unmounting Filesystems

- ▶ `umount [options] device|mountpoint`
- ▶ You cannot unmount a filesystem that is in use
 - ▶ Use *fuser* to check and/or kill processes
- ▶ Use the *remount* option to change a mounted filesystem's options atomically
 - ▶ `mount -o remount,ro /data`

Handling Swap Files and Partitions

- ▶ Swap space is a supplement to system RAM
- ▶ Basic setup involves:
 - ▶ Create a swap partition or file
 - ▶ Write special signature using *mkswap*
 - ▶ Add appropriate entries to */etc/fstab*
 - ▶ Activate swap space with `swapon -a`

Software RAID Configuration

- ▶ Create and define RAID devices using *mdadm*
 - ▶ `mdadm -C /dev/md0 -a yes -l 1 -n 2 -x 1 elements`
- ▶ Format each RAID device with a filesystem
 - ▶ `mke2fs -k /dev/md0`
- ▶ Test the RAID devices
- ▶ *mdadm* allows to check the status of your RAID devices
 - ▶ `mdadm --detail /dev/md0`

Software RAID Testing and Recovery

- ▶ Simulating disk failures
 - ▶ `mdadm /dev/md0 -f /dev/sda1`
- ▶ Recovering from a software RAID disk failure
 - ▶ replace the failed hard drive and power on
 - ▶ reconstruct partitions on the replacement drive
 - ▶ `mdadm /dev/md0 -a /dev/sda1`
- ▶ `mdadm`, `/proc/mdstat` and `syslog` messages

What is Logical Volume Manager?

- ▶ A layer of abstraction that allows easy manipulation of volumes. Including resizing of filesystems.
- ▶ Allow reorganization of filesystems across multiple physical devices
 - ▶ Devices are designated as Physical Volumes (PV)
 - ▶ One or more PV are used to create a Volume Group (VG)
 - ▶ PV are defined with Physical Extents of a fixed size
 - ▶ Logical Volumes (LV) are created on PV and are composed of Physical Extents
 - ▶ Filesystems may be created on Logical Volumes

Creating Logical Volumes

- ▶ Create physical volumes
 - ▶ `pvcreate /dev/sda3`
- ▶ Assign physical volumes to volume groups
 - ▶ `vgcreate vg0 /dev/sda3`
- ▶ Create logical volumes from volume groups
 - ▶ `lvcreate -L 256M -n data vg0`
 - ▶ `mke2fs -j /dev/vg0/data`

Resizing Logical Volumes

- ▶ Growing Volumes
 - ▶ *lvextend* can grow logical volumes
 - ▶ *resize2fs* can grow EXT3 filesystems online
 - ▶ *vgextend* adds new physical volumes to an existing volume group
- ▶ Shrinking Volumes
 - ▶ Filesystem have to be reduced first
 - ▶ Requires a filesystem check and cannot be performed online
 - ▶ *lvreduce* can then reduce volume
- ▶ Volume Groups can be reduced with:
 - ▶ `pvmove /dev/sda3`
 - ▶ `vgreduce vg0 /dev/sda3`

Lab: Installation

- ▶ Goals:
 - ▶ Deploy LVM on the software RAID device
 - ▶ Create a group with two partitions such that new partition could be added, and the filesystem could be extended

Lab: Prepare a paper

- ▶ Themes:
 - ▶ Compare software and hardware RAID
 - ▶ Compare filesystems in Linux (ext2, ext3, reiser, gfs, ...)
- ▶ Format:
 - ▶ Short presentation (15–20 minutes; 5-7 slides)
 - ▶ Paper containing comparison (1000 words)