

# Red Hat System Administration II

Complete Quick Reference Guide

---

## Ahmed Abdelwahed

Technical Trainer | Cloud & Infrastructure Specialist

 **Email**

[ahmed@abdelwahed.me](mailto:ahmed@abdelwahed.me)

 **Blog**

[www.abdelwahed.me](http://www.abdelwahed.me)

 **LinkedIn**

[linkedin.com/in/ahmedmct](https://linkedin.com/in/ahmedmct)

# 01. Linux Boot Sequence & GRUB2

## The 8-Stage Boot Process

#	Stage	What Happens
1	<b>Power On / POST</b>	BIOS/UEFI loads from non-volatile memory. Performs POST — checks CPU, RAM, disk controllers, hardware.
2	<b>Hardware Detection</b>	BIOS/UEFI probes and initializes storage controllers, disks, USB, network adapters, and graphics.
3	<b>Select Boot Device</b>	Reads boot order, scans bootable media (HDD, USB, PXE, DVD), hands off to boot sector or EFI loader.
4	<b>GRUB2 Boot Loader</b>	Loads from MBR or EFI System Partition. Reads /boot/grub2/grub.cfg. Presents boot menu.
5	<b>Kernel Initialization</b>	GRUB loads vmlinuz + initramfs. Kernel unpacks, loads drivers, detects hardware, mounts real root filesystem.
6	<b>systemd (PID 1)</b>	systemd starts as the first user-space process. Activates units in dependency order, switches to default target.
7	<b>Startup Services</b>	Networking, sshd, firewalld, cron/timers start. Init files loaded: /etc/profile, ~/.bash_profile, PAM modules.
8	<b>Login Available</b>	getty or GDM/SDDM login managers start. Users can authenticate via console or SSH.

## Systemd Targets (Runlevels)

Command / Option	Description
<code>systemctl get-default</code>	Show current default boot target
<code>systemctl --type=target</code>	List all loaded targets
<code>systemctl --type=target --all</code>	List all targets including inactive
<code>runlevel</code>	Show current runlevel (legacy compatibility)
<code>systemctl set-default multi-user.target</code>	Boot to CLI mode persistently
<code>systemctl set-default graphical.target</code>	Boot to GUI mode persistently
<code>systemctl isolate rescue.target</code>	Switch to rescue mode now (non-persistent)
<code>systemctl isolate multi-user.target</code>	Switch to multi-user CLI mode now
<code>systemctl poweroff</code>	Shut down the system
<code>systemctl reboot</code>	Reboot the system

Target	SysV Equiv	Purpose
<code>emergency.target</code>	—	Minimal shell; root filesystem read-only. Use when filesystems are corrupt.
<code>rescue.target</code>	Runlevel 1	Single-user mode; root shell + basic services. Use for config errors or password resets.
<code>multi-user.target</code>	Runlevel 3	Full multi-user CLI (no GUI). Standard for servers.

Target	SysV Equiv	Purpose
<code>graphical.target</code>	Runlevel 5	Full multi-user with GUI (GDM). Standard for desktops.
<code>reboot.target</code>	Runlevel 6	System reboot.

## GRUB2 Configuration

Never edit `grub.cfg` directly — it gets overwritten on updates. Use the method below instead.

### Method A: Permanent Changes

- `cp /etc/default/grub /etc/default/grub.bak` - backup first
- `nano /etc/default/grub` - edit `GRUB_CMDLINE_LINUX` (remove `'rhgb quiet'` for verbose boot)
- `grub2-mkconfig -o /boot/grub2/grub.cfg` - regenerate config

### Method B: Temporary (One-time) Boot Parameters

- Reboot → highlight kernel in GRUB menu
- Press `'e'` to edit the boot entry
- Find line starting with `'linux'` and append your target at the end
- Press `Ctrl+X` to boot with modified parameters

Command / Option	Description
<code>systemd.unit=rescue.target</code>	Boot to rescue mode (append to linux line)
<code>systemd.unit=emergency.target</code>	Boot to emergency mode (append to linux line)
<code>rd.break</code>	Stop before mounting real root — used for password resets
<code>init=/bin/bash</code>	Boot to a raw bash shell (no systemd)

## GRUB2 Password Protection

- `grub2-setpassword` - set GRUB superuser password
- `cat /boot/grub2/user.cfg` - view hashed password
- `grub2-mkconfig -o /boot/grub2/grub.cfg` - apply changes
- Reboot → press `'e'` to test password protection
- `rm /boot/grub2/user.cfg` - remove password if needed

## Reset Lost Root Password

- At GRUB menu, press `'e'` to edit boot entry
- Append `init=/bin/bash` to the linux line, then `Ctrl+X`
- `mount -o remount,rw /` - remount root as read-write
- `passwd` - set new root password

- 5 `touch /.autorelabel` - trigger SELinux relabeling on next boot
- 6 `exec /usr/lib/systemd/systemd - or reboot -f`

## Login Shell vs Non-Login Shell

Property	Login Shell	Non-Login Shell
<b>When</b>	SSH login, console login, su - user	SSH login, console login, su - user
<b>Files Read</b>	/etc/profile → ~/.bash_profile → ~/.bashrc	/etc/profile → ~/.bash_profile → ~/.bashrc
<b>Purpose</b>	Full environment initialization	Full environment initialization

## 02. Scheduling Jobs — at, cron, anacron, systemd Timers

### at — One-Time Scheduled Tasks

- 1 `dnf install at` - install the at package
- 2 `systemctl enable --now atd` - enable and start the atd service
- 3 `at 4:00 PM` - schedule a task (type commands, then Ctrl+D to save)

Command / Option	Description
<code>at 4:00 PM</code>	Run at 4:00 PM today
<code>at now + 2 hours</code>	Run 2 hours from now
<code>at 3:00 PM tomorrow</code>	Run at 3 PM tomorrow
<code>at 10:00 AM next Monday</code>	Run at 10 AM next Monday
<code>at 8:00 AM Jul 15</code>	Run at specific date
<code>atq</code>	View scheduled at jobs
<code>atrm &lt;job_number&gt;</code>	Remove a scheduled at job

### cron — Recurring Scheduled Tasks

#### Crontab Format

```
* * * * * command_to_run
| | | | |
| | | | | Day of week (0-7, 0/7=Sun)
| | | | | Month (1-12)
| | | | | Day of month (1-31)
| | | | | Hour (0-23)
| | | | | Minute (0-59)
```

#### crontab Management

Command / Option	Description
<code>crontab -e</code>	Edit current user's crontab
<code>crontab -l</code>	List current user's crontab
<code>crontab -r</code>	Remove current user's crontab
<code>crontab -e -u username</code>	Edit crontab for specific user (root)
<code>crontab -l -u username</code>	List crontab for specific user
<code>vim /etc/crontab</code>	Edit system-wide crontab (includes user field)

#### Cron Job Examples

Cron Expression	Meaning
<code>* * * * * /cmd</code>	Every minute

Cron Expression	Meaning
<code>0 5,17 * * * /cmd</code>	At 5 AM and 5 PM every day
<code>*/10 * * * * /cmd</code>	Every 10 minutes
<code>0 17 * * sun,fri /cmd</code>	5 PM on Sundays and Fridays
<code>3 12 * * sun /cmd</code>	12:03 PM every Sunday
<code>0 0 15 Jun,Jul,Aug * /cmd</code>	Midnight on 15th of Jun, Jul, Aug
<code>@reboot /cmd</code>	Once at system startup
<code>@daily /cmd</code>	Once per day at midnight
<code>@weekly /cmd</code>	Once per week on Sunday
<code>@monthly /cmd</code>	Once per month on the 1st

**NOTE** Cron security: Create `/etc/cron.allow` to whitelist users. Create `/etc/cron.deny` to blacklist. If `cron.allow` exists, only listed users can schedule.

## anacron — Missed Job Recovery

anacron runs missed scheduled jobs after the system powers on. Ideal for non-24/7 systems.

Command / Option	Description
<code>/etc/anacrontab</code>	Anacron configuration file
<code>7 10 my-backup /bin/bash /scripts/backup.sh</code>	Format: period(days) delay(mins) ID command

**NOTE** Anacron format: period delay job-id command. Period=days, Delay=minutes to wait after boot before running.

## systemd Timers — The Modern Cron

systemd timers replace cron with dependency-aware, logged, reliable job scheduling. Each timer requires a matching `.service` file.

### Timer Types

Command / Option	Description
<code>OnBootSec=5min</code>	Monotonic: run 5 minutes after boot
<code>OnUnitActiveSec=1h</code>	Monotonic: run 1 hour after last activation
<code>OnCalendar=daily</code>	Real-time: run daily at midnight
<code>OnCalendar=Mon *--* 08:00:00</code>	Real-time: every Monday at 8 AM
<code>OnCalendar=*--01 04:00:00</code>	Real-time: 1st of every month at 4 AM
<code>OnCalendar=*--* *:00/15:00</code>	Real-time: every 15 minutes

### Creating a Timer — Example (Daily Backup at 2 AM)

```
# /etc/systemd/system/backup.service
[Unit]
Description=Daily Website Backup
```

```
[Service]
Type=oneshot
ExecStart=/usr/local/bin/backup.sh

# /etc/systemd/system/backup.timer
[Unit]
Description=Run backup daily at 2AM

[Timer]
OnCalendar=*--* 02:00:00
Persistent=true
RandomizedDelaySec=300

[Install]
WantedBy=timers.target
```

Command / Option	Description
<code>systemctl daemon-reload</code>	Reload after creating/editing unit files
<code>systemctl enable --now backup.timer</code>	Enable and start the timer
<code>systemctl list-timers</code>	Show all active timers and next run time
<code>systemd-analyze calendar 'Mon *--* 08:00:00'</code>	Validate and test OnCalendar syntax
<code>journalctl -u backup.timer</code>	View timer execution logs
<code>journalctl -u backup.service</code>	View service execution logs

## 03. ACLs, Special Permissions & Shell Scripting

### Access Control Lists (ACLs)

ACLs extend standard Linux permissions to grant access to specific users or groups beyond the file owner, group, and others.

Command / Option	Description
<code>setfacl -m u:ahmed:rw file.txt</code>	Grant user ahmed read+write on file.txt
<code>setfacl -m g:sales:r data/</code>	Grant group sales read on directory
<code>setfacl -R -m g:sales:r data/</code>	Apply ACL recursively to all files
<code>setfacl -m d:u:ahmed:rw /mydir</code>	Set DEFAULT ACL (inherited by new files in dir)
<code>setfacl -d -m g:developers:rw projects/</code>	Default ACL for group on directory
<code>setfacl -x u:ahmed file.txt</code>	Remove ACL entry for user ahmed
<code>setfacl -b file.txt</code>	Remove ALL ACLs from file
<code>getfacl file.txt</code>	View ACLs on a file

**NOTE** When a file has ACLs, `ls -l` shows a '+' after the permissions string: `-rw-rw-r--+`

### Special Permission Bits

Bit	Octal	Symbol	Set Command	Effect
<b>SUID</b>	4	s (owner)	<code>chmod 4755 file</code> or <code>chmod u+s file</code>	File runs with owner's privileges. Find: <code>find / -perm /4000</code>
<b>SGID</b>	2	s (group)	<code>chmod 2755 dir</code> or <code>chmod g+s dir</code>	Files in dir inherit group ID. File runs as group. Find: <code>find / -perm /2000</code>
<b>Sticky Bit</b>	1	t (others)	<code>chmod 1777 dir</code> or <code>chmod +t dir</code>	Only owner can delete their files in dir (like /tmp). Find: <code>find / -perm /1000</code>

### Shell Scripting Fundamentals

#### Script Structure

```
#!/bin/bash           # Shebang - defines interpreter
set -e                # Exit immediately on error
set -u                # Error on unset variables
set -o pipefail       # Fail on pipeline errors

VARIABLE="value"      # Variable assignment (no spaces around =)
echo "Hello $VARIABLE" # Double quotes: variable expansion
echo 'Literal $VARIABLE' # Single quotes: no expansion
```

Command / Option	Description
<code>chmod +x script.sh</code>	Make script executable
<code>./script.sh</code>	Run script

Command / Option	Description
<code>bash script.sh</code>	Run without execute permission
<code>set -x</code>	Debug mode: print each command before running

## Special Variables

Command / Option	Description
<code>\$0</code>	Script name
<code>\$1, \$2, ...</code>	Positional arguments passed to script
<code>\$#</code>	Total number of arguments
<code>\$@</code>	All arguments as a list
<code>\$?</code>	Exit status of last command (0=success, non-zero=fail)
<code>\$\$</code>	PID of the current script

## Conditionals

```
if [ -f "/etc/passwd" ]; then
    echo "File exists"
elif [ -d "/tmp" ]; then
    echo "Directory exists"
else
    echo "Nothing found"
fi

# Short-circuit operators
cp data.txt /backup/ && echo "Success" || echo "Failed"
```

Command / Option	Description
<code>-f file</code>	File exists and is a regular file
<code>-d dir</code>	Directory exists
<code>-z str</code>	String is empty
<code>-n str</code>	String is not empty
<code>-eq / -ne</code>	Equal / Not equal (numeric)
<code>-lt / -gt</code>	Less than / Greater than (numeric)

## Loops

```
# for loop
for server in appl app2 db1; do
    echo "Pinging $server..."
    ping -c 1 $server
done

# while loop
count=1
while [ $count -le 5 ]; do
```

```
    echo "Count is $count"
    ((count++))
done
```

### trap — Self-Cleaning Scripts

```
#!/bin/bash
temp_file=$(mktemp)
trap "rm -f $temp_file; echo 'Cleanup done.'" EXIT
echo "Working with $temp_file..."
# Even Ctrl+C triggers cleanup
```

### getopts — Command-Line Flags

```
# Usage: ./adduser.sh -u john -g developers
while getopts "u:g:" opt; do
    case $opt in
        u) USERNAME="$OPTARG" ;;
        g) GROUP="$OPTARG" ;;
        *) echo "Usage: $0 -u <user> -g <group>"; exit 1 ;;
    esac
done
```

## 04. Storage: Partitions, LVM, RAID, Stratis & Encryption

### Storage Inspection Commands

Command / Option	Description
<code>lsblk -la</code>	List all block devices in tree format
<code>blkid</code>	Show UUIDs and filesystem types for all devices
<code>df -hT</code>	Show disk usage with filesystem types
<code>du -sh /path</code>	Show total size of a directory
<code>mount   column -t</code>	Show all currently mounted filesystems
<code>findmnt</code>	Show mounts in tree format
<code>findmnt -x</code>	Check for mount errors
<code>lsof +D /mnt/data</code>	List open files on a mount point

### Partitioning Tools Comparison

Tool	Table Type	Max Size	Best Use
<code>fdisk</code>	MBR	2 TB	Legacy systems, simple partitioning
<code>gdisk</code>	GPT	Exabytes	Disks > 2 TB, modern UEFI systems
<code>parted</code>	MBR/GPT	Both	Flexible, scriptable, supports both table types

#### fdisk — MBR Partitioning

Command / Option	Description
<code>fdisk -l /dev/sdb</code>	List partitions on sdb
<code>fdisk /dev/sdb</code>	Open interactive partitioning (n=new, d=delete, p=print, w=write)
<code>partprobe /dev/sdb</code>	Inform kernel of partition table changes

#### gdisk — GPT Partitioning


Command / Option	Description
<code>gdisk -l /dev/sdb</code>	Show partition table (shows if GPT is present)
<code>gdisk /dev/sdc</code>	Interactive GPT partitioning

#### parted — MBR & GPT

```
parted /dev/sdb
select /dev/sdb
mklabel gpt          # or msdos for MBR
mkpart primary 1 1024 # 1 MB to 1024 MB
p                   # print table
quit
```

## Filesystems

Command / Option	Description
<code>mkfs.xfs -L Data /dev/sdb2</code>	Create XFS with label 'Data'
<code>mkfs.ext4 /dev/sdb3</code>	Create ext4 filesystem
<code>xfs_info /dev/sdb2</code>	Show XFS filesystem info
<code>xfs_admin -U generate /dev/sdb2</code>	Generate new UUID for XFS filesystem
<code>blkid /dev/sdb2</code>	Show filesystem UUID, label, and type
<code>wipefs -a /dev/sdb</code>	Wipe all filesystem signatures from device

 **NOTE** XFS: max file size 8 EB, cannot shrink. ext4: max file size 2 TB, can shrink. XFS is the RHEL default.

## Mounting Filesystems


### Temporary Mount

Command / Option	Description
<code>mount /dev/sdb1 /data/sdbdata1</code>	Mount device to directory
<code>mount -t xfs /dev/sdb1 /data/</code>	Explicit filesystem type
<code>umount /data/sdbdata1</code>	Unmount (close all files first)

### Persistent Mount (/etc/fstab)

```
# Format: device mountpoint fstype options dump pass
UUID=abc123 /data xfs defaults 0 2
LABEL=Data /data xfs defaults 0 2
/dev/sdb1 /data xfs defaults 0 2
```

Command / Option	Description
<code>blkid</code>	Get UUIDs for fstab entries
<code>ls -l /dev/disk/by-uuid/</code>	View UUIDs as symlinks
<code>ls -l /dev/disk/by-label/</code>	View labels as symlinks
<code>mount -a</code>	Activate all fstab entries
<code>findmnt --verify</code>	Verify fstab for errors

 **NOTE** fstab pass field: 0=no fsck, 1=root filesystem, 2=other filesystems (checked after root)

## LVM — Logical Volume Management

LVM abstracts physical storage into flexible logical volumes. Order: Physical Volumes → Volume Groups → Logical Volumes.

### LVM Setup

- `dnf install lvm2 - install LVM tools`
- `pvcreate /dev/sdb /dev/sdc - create physical volumes`

- 3 `vgcreate vg1 /dev/sdb /dev/sdc - create volume group`
- 4 `lvcreate -n lv1 -L 2G vg1 - create 2GB logical volume`
- 5 `mkfs.xfs -L lvdata /dev/vg1/lv1 - format the LV`
- 6 `mkdir /lvml && mount /dev/vg1/lv1 /lvml - mount it`
- 7 Add to `/etc/fstab` for persistence

## LVM Management Commands

Command / Option	Description
<code>pvdisk / pvs</code>	Show physical volumes
<code>vgdisplay / vgs</code>	Show volume groups
<code>lvdisplay / lvs</code>	Show logical volumes
<code>vgextend vg1 /dev/sdd</code>	Add a new disk to volume group
<code>lvextend -L +50G -r /dev/vg1/lv1</code>	Extend LV by 50GB (auto-resizes filesystem)
<code>lvextend -l 100%FREE -r /dev/vg1/lv1</code>	Extend to use all free space
<code>xfs_growfs /lvml/</code>	Grow XFS filesystem manually (if -r not used)
<code>lvremove /dev/vg1/lv1</code>	Remove logical volume
<code>vgremove vg1</code>	Remove volume group
<code>pvremove /dev/sdb</code>	Remove physical volume

**NOTE** PE-based sizing: `lvcreate -l 50 -n lv1 vg1` (creates LV of 50 Physical Extents). Use `vgdisplay` to find PE size.

## Software RAID (mdadm)

Command / Option	Description
<code>cat /proc/mdstat</code>	View RAID status
<code>mdadm -C /dev/md0 -l 5 -n 3 /dev/sdb /dev/sdc /dev/sdd</code>	Create RAID-5 with 3 disks
<code>mdadm --create /dev/md0 -l 5 -n 3 /dev/sd[c-e]1 -x 1 /dev/sdf1</code>	RAID-5 with 1 hot spare
<code>mkfs.xfs -L Rdata /dev/md0</code>	Format RAID device
<code>mdadm --detail --scan &gt;&gt; /etc/mdadm.conf</code>	Make RAID config persistent
<code>mdadm --add /dev/md0 /dev/sde</code>	Add spare disk to array
<code>mdadm --grow /dev/md0 --raid-devices=4</code>	Increase active disks
<code>mdadm /dev/md0 --fail /dev/sdb</code>	Simulate disk failure (testing)
<code>watch cat /proc/mdstat</code>	Monitor RAID rebuild in real-time
<code>mdadm --stop /dev/md0</code>	Stop RAID device

## Stratis — Advanced Storage Management

Stratis simplifies thin-provisioning, snapshots, and monitoring. Pools contain filesystems that auto-grow.

- 1 `dnf install stratisd stratis-cli && systemctl enable --now stratisd`
- 2 `wipefs -a /dev/sdh /dev/sdi - clean disks first`
- 3 `stratis pool create pool_1 /dev/sdh /dev/sdi - create pool from disks`
- 4 `stratis fs create pool_1 fs_1 - create filesystem in pool`
- 5 `blkid -p /dev/stratis/pool_1/fs_1 - get UUID for fstab`
- 6 Add to `/etc/fstab`: `UUID=... /mount/point xfs defaults,x-systemd.requires=stratisd.service 0 0`

Command / Option	Description
<code>stratis pool list</code>	List pools
<code>stratis fs list</code>	List filesystems
<code>stratis blockdev list</code>	List block devices in pools
<code>stratis pool add-data pool_1 /dev/sdf</code>	Add disk to existing pool
<code>stratis fs snapshot pool_1 fs_1 snap_1</code>	Create snapshot
<code>stratis fs destroy pool_1 fs_1</code>	Delete filesystem
<code>stratis pool destroy pool_1</code>	Delete pool

## LUKS Disk Encryption

- 1 `dnf install -y cryptsetup - install cryptsetup`
- 2 `cryptsetup luksFormat /dev/sda1 - initialize LUKS (writes header)`
- 3 `cryptsetup luksOpen /dev/sda1 securedata - unlock as /dev/mapper/securedata`
- 4 `mkfs.ext4 /dev/mapper/securedata - create filesystem`
- 5 Add to `/etc/fstab`: `/dev/mapper/securedata /mnt/encrypted ext4 defaults 0 2`
- 6 Add to `/etc/crypttab`: `securedata /dev/sda1 none luks - auto-unlock at boot`

Command / Option	Description
<code>cryptsetup luksDump /dev/sda1</code>	Show LUKS header info and key slots
<code>cryptsetup status securedata</code>	Show mapped device status
<code>cryptsetup luksAddKey /dev/sda1</code>	Add a second passphrase/key
<code>cryptsetup luksRemoveKey /dev/sda1</code>	Remove a passphrase
<code>cryptsetup luksHeaderBackup /dev/sda1 -- header-backup-file /root/backup.img</code>	Backup LUKS header
<code>cryptsetup luksHeaderRestore /dev/sda1 -- header-backup-file /root/backup.img</code>	Restore LUKS header
<code>umount /mnt/encrypted &amp;&amp; cryptsetup luksClose securedata</code>	Unmount and re-lock

## Swap Space Management

Command / Option	Description
<code>free -h</code>	Show RAM and swap usage
<code>swapon -s</code>	List active swap spaces
<code>cat /proc/swaps</code>	Show swap usage details
<code>mkswap /dev/sdb5</code>	Initialize partition as swap
<code>swapon /dev/sdb5</code>	Enable swap partition
<code>swapoff /dev/sdb5</code>	Disable swap partition
<code>swapon -p 10 /dev/sdb1</code>	Enable with priority 10
<code>swapoff -a</code>	Disable all swap spaces
<code>dd if=/dev/zero of=swapfile bs=1G count=1</code>	Create a 1GB swap file
<code>mkswap swapfile &amp;&amp; chmod 0600 swapfile</code>	Prepare swap file
<code>swapon swapfile</code>	Enable swap file

## 05. System Logging — journald & Remote rsyslog

### systemd-journald

journald is systemd's logging subsystem. Logs are stored in binary format and queried with journalctl.

#### Basic journalctl Commands

Command / Option	Description
<code>journalctl -f</code>	Follow live system logs (like tail -f)
<code>journalctl -u sshd -f</code>	Follow logs for a specific service
<code>journalctl -b</code>	Logs from current boot
<code>journalctl --list-boots</code>	List all recorded boot sessions
<code>journalctl -b -1</code>	Logs from previous boot
<code>journalctl -r</code>	Show logs in reverse (newest first)
<code>journalctl -n 100</code>	Show last 100 log entries
<code>journalctl -k</code>	Show kernel messages only

#### Filtering by Time

Command / Option	Description
<code>journalctl --since today</code>	All logs since today midnight
<code>journalctl --since yesterday</code>	All logs since yesterday
<code>journalctl --since '1 hour ago'</code>	Logs from last hour
<code>journalctl --since '2025-08-21 08:00' --until '2025-08-22'</code>	Logs in a specific time window

#### Filtering by Priority

Command / Option	Description
<code>journalctl -p err</code>	Show errors and above (crit, alert, emerg)
<code>journalctl -p warn</code>	Show warnings and above
<code>journalctl -u sshd -p err --since yesterday</code>	SSH errors since yesterday

#### Filtering by Process / User

Command / Option	Description
<code>journalctl _PID=1234</code>	Logs from a specific PID
<code>journalctl _UID=1000</code>	Logs from a specific user
<code>journalctl _COMM=su --since '24 hours ago'</code>	Logs from 'su' command in last 24 hours
<code>journalctl --disk-usage</code>	Show disk space used by journal

## Persistent Journal Storage

- 1 `mkdir -p /var/log/journal - create persistent log directory`
- 2 Edit `/etc/systemd/journald.conf` → set: `Storage=persistent`
- 3 `systemctl restart systemd-journald`

💡 **NOTE** By default, journald uses volatile storage (lost on reboot). Creating `/var/log/journal` enables persistence.

## Redirecting Logs to a Different Volume

- 1 `mount -t xfs /dev/sda1 /x1 - mount new volume temporarily`
- 2 `cp -r /var/log/* /x1 - backup existing logs`
- 3 `rm -rf /var/log/* - clear /var/log directory`
- 4 `mount -t xfs /dev/sda1 /var/log - mount new volume to /var/log`
- 5 `echo '/dev/sda1 /var/log xfs defaults 0 2' >> /etc/fstab - make persistent`
- 6 `cp -r /x1/* /var/log/ - restore backed-up logs`

## Remote Log Forwarding (node1 → node2)

### On node1 (Source Server)

```
# /etc/rsyslog.d/remote.conf
*. * @node2:514 # TCP forwarding
*. * @node2:514 # UDP forwarding
*. * ~ # Stop local logging (optional, saves space)

systemctl restart rsyslog
```

### On node2 (Destination Server)

```
# /etc/rsyslog.d/receive.conf
module(load="imtcp")
input(type="imtcp" port="514")

template(name="RemoteLogs" type="string"
  string="/var/log/remote/%HOSTNAME%/%PROGRAMNAME%.log")

if ($fromhost-ip != '127.0.0.1') then {
  action(type="omfile" dynaFile="RemoteLogs")
  stop
}
```

- 1 `mkdir -p /var/log/remote && chmod 755 /var/log/remote`
- 2 `systemctl restart rsyslog`
- 3 `firewall-cmd --permanent --add-port=514/tcp && firewall-cmd --reload`
- 4 On node1: `logger 'Test message' - send test`

**5** On node2: `tail -f /var/log/remote/node1/ - verify receipt`

## 06. Systemd Advanced — Units, Sockets, Mounts & Timers

### systemd Unit Types

Extension	Type	Purpose	Example
<code>.service</code>	Service	Traditional daemon/process	sshd.service, nginx.service
<code>.socket</code>	Socket	Network listeners — on-demand activation	sshd.socket
<code>.timer</code>	Timer	Scheduled tasks — cron replacement	logrotate.timer
<code>.mount</code>	Mount	Filesystem mounts with dependencies	var-lib-docker.mount
<code>.target</code>	Target	Group of units — like runlevels	multi-user.target
<code>.path</code>	Path	Watch filesystem path for changes	cups.path

Command / Option	Description
<code>systemctl list-units --type=service</code>	List all loaded services
<code>systemctl list-units --type=socket</code>	List all loaded sockets
<code>systemctl list-units --type=timer</code>	List all loaded timers
<code>systemctl list-unit-files --type=service --state=enabled</code>	List enabled services
<code>systemctl cat nginx.service</code>	Show full unit file
<code>systemctl show nginx.service</code>	Show all settable properties
<code>systemctl edit nginx.service</code>	Create override (drop-in) file
<code>systemctl daemon-reload</code>	Reload after editing unit files

### Socket Activation

Socket activation starts services only when a connection arrives — reducing RAM usage for rarely-used services.

Command / Option	Description
<code>systemctl enable --now sshd.socket</code>	Enable SSH socket activation (starts sshd only on connection)
<code>systemctl disable sshd.socket</code>	Disable socket activation
<code>systemctl enable --now sshd.service</code>	Re-enable persistent sshd service

**NOTE** Socket activation benefit: sshd.service uses ~10MB RAM constantly. Socket activation = 0 MB until first connection.

### Systemd Mount Units

Mount units provide dependency-aware mounting — ensuring required services start before the mount.

**NOTE CRITICAL:** The filename MUST match the mount point path with slashes replaced by dashes. /data → data.mount | /mnt/backup → mnt-backup.mount | /var/www/html → var-www-html.mount

```
# /etc/systemd/system/mnt-data.mount
[Unit]
Description=Mount Data Disk
After=blockdev@dev-sdb1.target

[Mount]
What=/dev/sdb1
Where=/mnt/data
Type=xf
Options=defaults,noatime

[Install]
WantedBy=multi-user.target
```

Command / Option	Description
<code>systemctl daemon-reload</code>	Reload after creating unit file
<code>systemctl enable --now mnt-data.mount</code>	Enable and mount
<code>systemctl list-units --type=mount</code>	List all mount units
<code>journalctl -u mnt-data.mount</code>	View mount unit logs

## Service Resilience (Failure Recovery)

Add these directives to your [Service] section for automatic recovery:

```
[Service]
Restart=on-failure      # Restart if it crashes
RestartSec=5s          # Wait 5 seconds before restart
WatchdogSec=30s        # Kill and restart if hung for 30s
StartLimitBurst=5      # Max 5 restart attempts
StartLimitInterval=300 # ...within 300 seconds
```

## 07. File Transfer — rsync, scp & NFS/AutoFS

### rsync — Efficient Synchronization

rsync transfers only changed blocks, making it ideal for backups and large-scale synchronization.

Command / Option	Description
<code>rsync -zvh backup.tar /tmp/backups/</code>	Local file copy with compression and verbose output
<code>rsync -avzh /root/rpmpkgs /tmp/backups/</code>	Sync directory locally (archive mode)
<code>rsync -avz rpmpkgs/ root@192.168.0.101:/home/</code>	Push directory to remote server
<code>rsync -avzh root@192.168.0.100:/home/source/ /tmp/myrpms</code>	Pull from remote to local
<code>rsync -avzhe ssh --progress /home/ root@192.168.0.100:/root/</code>	With progress display
<code>rsync --max-size='200k' -avzhe ssh /var/lib/rpm/ root@host:/tmp/</code>	Limit file size
<code>rsync --remove-source-files -zvh backup.tar /tmp/</code>	Delete source after transfer
<code>rsync --bwlimit=100 -avzhe ssh /var/lib/rpm/ root@host:/tmp/</code>	Limit bandwidth to 100KB/s

### scp — Secure Copy Protocol

Command / Option	Description
<code>scp user@10.0.0.4:/etc/passwd .</code>	Download file from remote to current dir
<code>scp -P 1515 root@10.0.0.4:/etc/passwd .</code>	Download using non-standard SSH port
<code>scp file01 file02 file03 /repo</code>	Copy multiple files locally
<code>scp -r /local/dir user@host:/remote/dir</code>	Recursively copy directory to remote
<code>scp -C file.tar.gz user@host:/path/</code>	Copy with compression enabled
<code>scp user@host1:/file.txt user@host2:/path/</code>	Copy between two remote hosts

### NFS — Network File System

#### NFS Server Setup

- `1 yum install -y nfs-utils && systemctl enable --now nfs-server`
- `2 mkdir -p /srv/nfs/share && chmod 777 /srv/nfs/share`
- `3 nano /etc/exports → /srv/nfs/share 192.168.1.0/24(rw, sync, no_root_squash)`
- `4 exportfs -r - reload exports | exportfs -v - verify`
- `5 firewall-cmd --permanent --add-service=nfs && firewall-cmd --reload`

## NFS Client Setup

```
1 yum install -y nfs-utils
2 mkdir -p /mnt/nfs-share
3 mount -t nfs 192.168.1.100:/srv/nfs/share /mnt/nfs-share
4 df -h | grep nfs - verify mount
5 nano /etc/fstab → 192.168.1.100:/srv/nfs/share /mnt/nfs-share nfs defaults 0 0
6 mount -a - activate fstab
```

## AutoFS — Automatic NFS Mounting

AutoFS mounts filesystems on-demand and unmounts them after a period of inactivity.

```
1 yum install -y autofs
2 nano /etc/auto.master → /mnt/nfs-auto /etc/auto.nfs
3 nano /etc/auto.nfs → share -fstype=nfs,rw 192.168.1.100:/srv/nfs/share
4 systemctl enable --now autofs
5 ls /mnt/nfs-auto/share - access triggers auto-mount
```

Command / Option	Description
<code>/etc/auto.master</code>	Master map: defines mount points and map files
<code>/etc/auto.nfs</code>	Map file: defines share names and NFS sources
<code>/etc/autofs.conf</code>	AutoFS config: set timeout = 60 for 60-second idle unmount

## 08. Managing SELinux


### SELinux Overview

SELinux (Security-Enhanced Linux) provides Mandatory Access Control (MAC) on top of standard Linux permissions. Even root cannot bypass SELinux policy.

Mode	Behavior
<b>Enforcing</b>	SELinux enforces all policies. Violations are denied AND logged. Default in RHEL.
<b>Permissive</b>	Violations are LOGGED but NOT denied. Use for troubleshooting SELinux issues.
<b>Disabled</b>	SELinux is completely off. Requires relabeling the entire filesystem on re-enable. Avoid in production.

### Core SELinux Commands

Command / Option	Description
<code>getenforce</code>	Show current SELinux mode
<code>setenforce 0</code>	Temporarily set to Permissive (until reboot)
<code>setenforce 1</code>	Temporarily set to Enforcing
<code>sestatus</code>	Show detailed SELinux status and policy info
<code>vim /etc/sysconfig/selinux</code>	Edit persistent mode (SELINUX=enforcing/permissive/disabled)


 **NOTE** Changing SELINUX=disabled requires a full reboot AND filesystem relabeling. Set SELINUX=permissive for troubleshooting, not disabled.

### Managing SELinux Port Contexts

Command / Option	Description
<code>dnf install -y policycoreutils-python-utils</code>	Install semanage tools
<code>semanage port -l</code>	List all SELinux port contexts
<code>semanage port -l   grep http</code>	Check if port is allowed for HTTP
<code>semanage port -a -t http_port_t -p tcp 5555</code>	Allow port 5555 for HTTP
<code>semanage port -a -t ssh_port_t -p tcp 1414</code>	Allow port 1414 for SSH
<code>semanage port -a -t unreserved_port_t -p tcp 6789</code>	Allow port 6789 for any service
<code>semanage port -d -t http_port_t -p tcp 5555</code>	Remove custom port definition
<code>semanage port -l   grep sshd</code>	Verify SSH port configuration

### SELinux File Contexts

Command / Option	Description
<code>ls -Z /var/www/html/</code>	Show file SELinux contexts
<code>ps -eZ   grep httpd</code>	Show process contexts
<code>chcon -t httpd_sys_content_t file.html</code>	Change file context temporarily
<code>restorecon -v file.html</code>	Restore default context from policy
<code>restorecon -Rv /var/www/html/</code>	Recursively restore all contexts
<code>semanage fcontext -a -t httpd_sys_content_t '/mydir(/.*)?'</code>	Add persistent context rule

 **NOTE** TIP: Use `chcon` for temporary fixes, `semanage fcontext` + `restorecon` for permanent solutions.

## Troubleshooting SELinux Denials

Command / Option	Description
<code>ausearch -m avc -ts recent</code>	Search recent SELinux denial messages
<code>journalctl -t setroubleshoot</code>	View SELinux troubleshoot messages
<code>sealert -l &lt;alert-id&gt;</code>	Get detailed explanation and suggested fix
<code>audit2allow -a</code>	Generate allow rules from denials (review before applying)

## 09. Managing Firewall with Firewalld

### Firewalld Architecture

Firewalld uses zones to define trust levels for network connections. Every interface/source belongs to a zone, and each zone has its own set of rules.

Command / Option	Description
<code>systemctl status firewalld</code>	Check firewall service status
<code>systemctl enable --now firewalld</code>	Enable and start firewalld
<code>firewall-cmd --state</code>	Show if firewalld is running
<code>firewall-cmd --list-all</code>	Show all rules for default zone

### Managing Services

Command / Option	Description
<code>firewall-cmd --list-services</code>	List allowed services
<code>firewall-cmd --add-service=http</code>	Temporarily allow HTTP (lost on reload)
<code>firewall-cmd --add-service=http --permanent</code>	Permanently allow HTTP
<code>firewall-cmd --remove-service=http --permanent</code>	Remove HTTP service
<code>firewall-cmd --reload</code>	Apply permanent rules to runtime
<code>firewall-cmd --runtime-to-permanent</code>	Save current runtime rules as permanent

### Managing Ports

Command / Option	Description
<code>firewall-cmd --list-ports</code>	List all open ports
<code>firewall-cmd --add-port=5050/tcp --permanent</code>	Permanently open TCP port 5050
<code>firewall-cmd --add-port=514/tcp --permanent</code>	Open port for remote syslog
<code>firewall-cmd --remove-port=5050/tcp --permanent</code>	Remove port rule
<code>firewall-cmd --reload</code>	Apply changes


### Zones

Command / Option	Description
<code>firewall-cmd --get-zones</code>	List all available zones
<code>firewall-cmd --get-default-zone</code>	Show current default zone
<code>firewall-cmd --set-default-zone=internal</code>	Change default zone

Command / Option	Description
<code>firewall-cmd --list-all-zones</code>	Show rules for ALL zones
<code>firewall-cmd --zone=public --add-service=http</code>	Add service to public zone
<code>firewall-cmd --zone=public --list-services</code>	List services in public zone
<code>firewall-cmd --zone=public --add-interface=eth1</code>	Assign interface to zone

Zone	Trust Level
<code>drop</code>	All incoming dropped, no reply. Most restrictive.
<code>block</code>	All incoming rejected with icmp-admin-prohibited.
<code>public</code>	Default zone. Don't trust other hosts. Basic services allowed.
<code>work</code>	Mostly trusted — work machines on same network.
<code>home</code>	Mostly trusted — home machines.
<code>internal</code>	Trusted internal network. Most services allowed.
<code>trusted</code>	All connections accepted. Most permissive.

 **NOTE** Service definition XML files are in `/usr/lib/firewalld/services/`. Custom services go in `/etc/firewalld/services/`.

## 10. Performance Tuning (tuned) & KVM Virtualization

### tuned — Adaptive System Tuning

The tuned daemon automatically adjusts kernel and system parameters based on current workload using predefined profiles.

Command / Option	Description
<code>systemctl status tuned</code>	Check tuned service status
<code>dnf install tuned &amp;&amp; systemctl enable --now tuned</code>	Install and start tuned
<code>tuned-adm active</code>	Show currently active profile
<code>tuned-adm list</code>	List all available profiles
<code>tuned-adm recommend</code>	Show recommended profile for this system
<code>tuned-adm profile throughput-performance</code>	Apply throughput-performance profile
<code>tuned-adm profile balanced</code>	Apply balanced profile
<code>tuned-adm profile powersave</code>	Apply power-saving profile
<code>tuned-adm off</code>	Disable all dynamic tuning

Profile	Best For
<code>throughput-performance</code>	Servers requiring maximum I/O and network throughput
<code>latency-performance</code>	Real-time applications requiring low latency
<code>virtual-host</code>	KVM/QEMU hypervisor hosts
<code>virtual-guest</code>	VMs running inside hypervisors
<code>balanced</code>	General-purpose desktops and mixed workloads
<code>powersave</code>	Laptops and low-power environments
<code>desktop</code>	Desktop systems — improved responsiveness
<code>network-latency</code>	Low-latency network applications

### KVM — Kernel-based Virtual Machine

#### Installation

```
dnf install -y qemu-kvm libvirt libvirt-daemon libvirt-daemon-driver-qemu virt-install virt-manager
systemctl enable --now libvirtd
```

#### Create a VM (virt-install)

```
virt-install \
  --name test-vm \
  --memory 2048 \
  --vcpus 2 \
  --disk size=20 \
  --os-variant rhel9.0 \
```

```
--cdrom /path/to/rhel9.iso \  
--boot uefi \  
--graphics vnc  
  
# Without OS (import):  
virt-install --name test-vm --memory 512 --vcpus 1 --disk size=8 --os-variant  
generic --boot uefi --noautoconsole --import
```

## virsh — VM Management CLI

Command / Option	Description
<code>virsh list --all</code>	List all VMs (running and stopped)
<code>virsh start vm1</code>	Start VM named vm1
<code>virsh shutdown vm1</code>	Graceful shutdown of vm1
<code>virsh destroy vm1</code>	Force stop vm1 (like pulling the power cord)
<code>virsh reboot vm1</code>	Reboot vm1
<code>virsh autostart vm1</code>	Start vm1 automatically at host boot
<code>virsh autostart vm1 --disable</code>	Disable autostart for vm1
<code>virsh net-list</code>	List virtual networks
<code>virsh dominfo vm1</code>	Show VM configuration details
<code>virt-manager</code>	Open graphical VM management interface

# 11. Containers with Podman on RHEL

## Container Concepts


Concept	Description
<b>Image</b>	Read-only blueprint containing code, runtime, libraries, and config. The template for containers.
<b>Container</b>	A running instance of an image — isolated user-space process with its own filesystem, network, and PID space.
<b>Registry</b>	Storage for container images. Examples: registry.redhat.io, quay.io, Docker Hub.
<b>Pod</b>	Group of one or more containers sharing network/storage (Kubernetes concept). Podman supports pods natively.
<b>Volume</b>	Persistent storage that survives container removal. Mounted into the container at runtime.
<b>Orchestration</b>	Automates deployment, scaling, and management of containers (Kubernetes, OpenShift).

## Podman Installation & Images

Command / Option	Description
<code>dnf -y install podman</code>	Install Podman
<code>podman --version</code>	Verify installation
<code>podman search fedora</code>	Search registries for fedora images
<code>podman pull nginx</code>	Download nginx image
<code>podman images</code>	List downloaded images
<code>podman rmi ubuntu</code>	Remove image
<code>podman rmi -f \$(podman images -aq)</code>	Remove all images
<code>podman inspect nginx:latest</code>	Show image details and layers

## Running Containers

Command / Option	Description
<code>podman run hello-world</code>	Run and exit (test)
<code>podman run -d --name my-nginx -p 8080:80 nginx</code>	Run nginx in background, map port 8080→80
<code>podman run -it ubuntu bash</code>	Interactive Ubuntu terminal
<code>podman run -it -d --name myfirst ubuntu</code>	Detached interactive container
<code>podman run -d --restart=always --name my-nginx -p 8080:80 nginx</code>	Auto-restart on failure
<code>podman run -d --name my-mariadb -v my-data:/var/lib/mysql mariadb</code>	With persistent volume

 **NOTE** Port mapping: `-p HOST_PORT:CONTAINER_PORT` | `-d` = detached (background) | `-it` = interactive terminal

## Managing Containers

Command / Option	Description
<code>podman ps</code>	List running containers
<code>podman ps -a</code>	List all containers (including stopped)
<code>podman stop myfirst</code>	Gracefully stop container
<code>podman start myfirst</code>	Start a stopped container
<code>podman restart myfirst</code>	Stop and start container
<code>podman rm myapp</code>	Remove stopped container
<code>podman rm -f myapp</code>	Force-remove running container
<code>podman rm -f \$(podman ps -aq)</code>	Remove all containers

## Interacting with Running Containers


Command / Option	Description
<code>podman exec myfirst /bin/ps -aux</code>	Run command in running container
<code>podman exec -it myfirst bash</code>	Open interactive shell in container
<code>podman attach myfirst</code>	Re-attach to running container
<code>Ctrl-p then Ctrl-q</code>	Detach from container without stopping it
<code>podman logs my-nginx</code>	View container logs
<code>podman logs -f my-nginx</code>	Follow container logs live
<code>podman top myfirst</code>	Show processes in container
<code>podman stats myfirst</code>	Live resource usage (CPU, RAM, I/O)
<code>podman events --since '1h'</code>	Show container events from last hour

## Building Images with Buildah

```
# Build a custom Apache image on RHEL
containerID=$(buildah from fedora)
buildah run $containerID -- dnf install httpd -y
buildah config --entrypoint '["httpd", "-D", "FOREGROUND"]' $containerID
buildah commit $containerID my-apache:latest
```

```
# Then run it:
```

```
podman run -d -p 8080:80 my-apache:latest
```

 **NOTE** Podman is daemonless and rootless — runs without a background daemon. Buildah creates OCI-compliant images compatible with Docker.