

Version : **2022.01**

Dernière mise-à-jour : 2023/11/20 12:52

LDF512 - Gestion du Noyau et des Quotas

Contenu du Module

- **LDF512 - Gestion du Noyau et des Quotas**
 - Contenu du Module
 - Rôle du noyau
 - LAB #1 - Compilation et installation du noyau
 - 1.1 - Déplacer /home
 - 1.2 - Télécharger le Code Source du Noyau
 - 1.3 - Configurer le Noyau
 - 1.4 - Compiler le Noyau
 - 1.5 - Installer le Nouveau Noyau
 - 1.6 - Désinstaller un Noyau
 - LAB #2 - Mise à jour du Noyau avec le Gestionnaire des Paquets
 - LAB #3 - Gestion des Quotas
 - 3.1 - La Commande quotacheck
 - 3.2 - La Commande edquota
 - 3.3 - La Commande quotaon
 - 3.4 - La Commande repquota
 - 3.5 - La Commande quota
 - 3.6 - La Commande warnquota

Rôle du noyau

Le noyau ou *kernel* est la partie du système d'exploitation qui gère les entrées/sorties avec des périphériques. Dans certains cas il est préférable de

recompiler le noyau de Linux. La motivation de cette recompilation peut être :

- la diminution de la taille du noyau,
- la prise en charge de nouveau matériel,
- l'ajout de fonctionnalités,
- l'optimisation du code,
- la correction de bogues,
- le besoin d'une fonctionnalité expérimentale.

Commencez par identifier le noyau utilisé par votre machine :

```
root@debian11:~# uname -r
5.10.0-13-amd64
```

Dans le cas d'une utilisation courante de Linux, il est cependant préférable de faire appel aux **modules**. Les modules se trouvent dans le répertoire **/lib/modules/<version-du-noyau>** :

```
root@debian11:~# ls /lib/modules/`uname -r`/
kernel          modules.alias.bin  modules.builtin.alias.bin  modules.builtin.modinfo  modules.dep.bin
modules.order    modules.symbols
modules.alias    modules.builtin    modules.builtin.bin        modules.dep              modules.devname
modules.softdep  modules.symbols.bin
```

Les commandes pour manipuler les modules sont :

- insmod
- rmmod
- lsmod
- modprobe

Par exemple :

```
root@debian11:~# lsmod
Module          Size  Used by
```

rfkill	28672	1	
raid456	180224	1	
async_raid6_recov	24576	1	raid456
async_memcpy	20480	2	raid456,async_raid6_recov
async_pq	20480	2	raid456,async_raid6_recov
async_xor	20480	3	async_pq,raid456,async_raid6_recov
async_tx	20480	5	async_pq,async_memcpy,async_xor,raid456,async_raid6_recov
md_mod	180224	1	raid456
sg	36864	0	
virtio_balloon	24576	0	
joydev	28672	0	
qemu_fw_cfg	20480	0	
pcspkr	16384	0	
serio_raw	20480	0	
evdev	28672	2	
ecryptfs	122880	0	
parport_pc	40960	0	
ppdev	24576	0	
lp	20480	0	
parport	69632	3	parport_pc,lp,ppdev
fuse	167936	1	
configfs	57344	1	
ip_tables	32768	0	
x_tables	53248	1	ip_tables
autofs4	53248	2	
ext4	921600	1	
crc16	16384	1	ext4
mbcache	16384	1	ext4
jbd2	151552	1	ext4
hid_generic	16384	0	
btrfs	1568768	0	
usbhid	65536	0	
hid	147456	2	usbhid,hid_generic
blake2b_generic	20480	0	

xor	24576	2	async_xor,btrfs
raid6_pq	122880	4	async_pq,btrfs,raid456,async_raid6_recov
libcrc32c	16384	2	btrfs,raid456
crc32c_generic	16384	3	
usb_storage	81920	0	
dm_mod	163840	4	
sd_mod	61440	10	
t10_pi	16384	1	sd_mod
crc_t10dif	20480	1	t10_pi
crct10dif_generic	16384	1	
crct10dif_common	16384	2	crct10dif_generic,crc_t10dif
sr_mod	28672	0	
cdrom	73728	1	sr_mod
virtio_net	61440	0	
net_failover	24576	1	virtio_net
failover	16384	1	net_failover
virtio_scsi	24576	8	
bochs_drm	16384	0	
ata_generic	16384	0	
drm_vram_helper	20480	1	bochs_drm
uhci_hcd	53248	0	
drm_ttm_helper	16384	1	drm_vram_helper
ttm	114688	2	drm_vram_helper,drm_ttm_helper
ehci_hcd	98304	0	
drm_kms_helper	278528	4	drm_vram_helper,bochs_drm
cec	61440	1	drm_kms_helper
psmouse	184320	0	
drm	618496	6	drm_kms_helper,drm_vram_helper,bochs_drm,drm_ttm_helper,ttm
usbcore	323584	4	usbhid,usb_storage,ehci_hcd,uhci_hcd
ata_piix	36864	0	
virtio_pci	28672	0	
libata	290816	2	ata_piix,ata_generic
virtio_ring	36864	4	virtio_balloon,virtio_scsi,virtio_pci,virtio_net
virtio	16384	4	virtio_balloon,virtio_scsi,virtio_pci,virtio_net

i2c_piix4	28672	0
scsi_mod	262144	6 virtio_scsi,sd_mod,usb_storage,libata,sg,sr_mod
usb_common	16384	3 usbcore,ehci_hcd,uhci_hcd
floppy	90112	0
button	24576	0

Pour ajouter un module, on peut utiliser la commande **insmod** ou **modprobe**. Cette dernière ajoute non seulement le module passé en argument mais également ses dépendances :

```
root@debian11:~# modprobe bonding
```

```
root@debian11:~# lsmod | head
```

Module	Size	Used by
bonding	192512	0
rfskill	28672	1
raid456	180224	1
async_raid6_recov	24576	1 raid456
async_memcpy	20480	2 raid456,async_raid6_recov
async_pq	20480	2 raid456,async_raid6_recov
async_xor	20480	3 async_pq,raid456,async_raid6_recov
async_tx	20480	5 async_pq,async_memcpy,async_xor,raid456,async_raid6_recov
md_mod	180224	1 raid456

Pour supprimer un module, on peut utiliser la commande **rmmod** ou **modprobe -r**. Cette dernière essaie de supprimer les dépendances non-utilisées :

```
root@debian11:~# modprobe -r bonding
```

```
root@debian11:~# lsmod | head
```

Module	Size	Used by
rfskill	28672	1
raid456	180224	1
async_raid6_recov	24576	1 raid456
async_memcpy	20480	2 raid456,async_raid6_recov
async_pq	20480	2 raid456,async_raid6_recov
async_xor	20480	3 async_pq,raid456,async_raid6_recov

async_tx	20480	5	async_pq, async_memcpy, async_xor, raid456, async_raid6_recov
md_mod	180224	1	raid456
sg	36864	0	

Les dépendances des modules sont résolues par la commande **modprobe** grâce au fichier **/lib/modules/<version-du-noyau>/modules.dep**. Ce dernier peut être créé manuellement grâce à la commande **depmod** :

```
root@debian11:~# more /lib/modules/`uname -r`/modules.dep
kernel/arch/x86/events/amd/power.ko:
kernel/arch/x86/events/intel/intel-uncore.ko:
kernel/arch/x86/events/intel/intel-cstate.ko:
kernel/arch/x86/events/rapl.ko:
kernel/arch/x86/kernel/cpu/mce/mce-inject.ko:
kernel/arch/x86/kernel/msr.ko:
kernel/arch/x86/kernel/cpuid.ko:
kernel/arch/x86/crypto/glue_helper.ko:
kernel/arch/x86/crypto/twofish-x86_64.ko: kernel/crypto/twofish_common.ko
kernel/arch/x86/crypto/twofish-x86_64-3way.ko: kernel/arch/x86/crypto/twofish-x86_64.ko
kernel/crypto/twofish_common.ko kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/twofish-avx-x86_64.ko: kernel/crypto/crypto_simd.ko kernel/crypto/cryptd.ko
kernel/arch/x86/crypto/twofish-x86_64-3way.ko kernel/arch/x86/crypto/twofish-x86_64.ko kernel/crypto/twofish_co
mmon.ko kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/serpent-sse2-x86_64.ko: kernel/crypto/serpent_generic.ko kernel/crypto/crypto_simd.ko
kernel/crypto/cryptd.ko kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/serpent-avx-x86_64.ko: kernel/crypto/serpent_generic.ko kernel/crypto/crypto_simd.ko
kernel/crypto/cryptd.ko kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/serpent-avx2.ko: kernel/arch/x86/crypto/serpent-avx-x86_64.ko
kernel/crypto/serpent_generic.ko kernel/crypto/crypto_simd.ko kernel/crypto/cryptd.ko
kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/des3_ede-x86_64.ko: kernel/lib/crypto/libdes.ko
kernel/arch/x86/crypto/camellia-x86_64.ko: kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/camellia-aesni-avx-x86_64.ko: kernel/arch/x86/crypto/camellia-x86_64.ko
kernel/crypto/crypto_simd.ko kernel/crypto/cryptd.ko kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/camellia-aesni-avx2.ko: kernel/arch/x86/crypto/camellia-aesni-avx-x86_64.ko
```

kernel/arch/x86/crypto/camellia-x86_64.ko kernel/crypto/crypto_simd.ko kernel/crypto/cryptd.ko kernel/arch/x86/
crypto/glue_helper.ko
kernel/arch/x86/crypto/blowfish-x86_64.ko: kernel/crypto/blowfish_common.ko
kernel/arch/x86/crypto/cast5-avx-x86_64.ko: kernel/crypto/cast5_generic.ko kernel/crypto/cast_common.ko
kernel/crypto/crypto_simd.ko kernel/crypto/cryptd.ko
kernel/arch/x86/crypto/cast6-avx-x86_64.ko: kernel/crypto/cast6_generic.ko kernel/crypto/cast_common.ko
kernel/crypto/crypto_simd.ko kernel/crypto/cryptd.ko kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/aegis128-aesni.ko: kernel/crypto/crypto_simd.ko kernel/crypto/cryptd.ko
kernel/arch/x86/crypto/chacha-x86_64.ko: kernel/lib/crypto/libchacha.ko
kernel/arch/x86/crypto/aesni-intel.ko: kernel/lib/crypto/libaes.ko kernel/crypto/crypto_simd.ko
kernel/crypto/cryptd.ko kernel/arch/x86/crypto/glue_helper.ko
kernel/arch/x86/crypto/sha1-ssse3.ko:
kernel/arch/x86/crypto/sha256-ssse3.ko:
kernel/arch/x86/crypto/sha512-ssse3.ko: kernel/crypto/sha512_generic.ko
kernel/arch/x86/crypto/blake2s-x86_64.ko: kernel/lib/crypto/libblake2s-generic.ko
kernel/arch/x86/crypto/ghash-clmulni-intel.ko: kernel/crypto/cryptd.ko
kernel/arch/x86/crypto/crc32c-intel.ko:
kernel/arch/x86/crypto/crc32-pclmul.ko:
kernel/arch/x86/crypto/crct10dif-pclmul.ko: kernel/crypto/crct10dif_common.ko
kernel/arch/x86/crypto/poly1305-x86_64.ko:
kernel/arch/x86/crypto/curve25519-x86_64.ko: kernel/lib/crypto/libcurve25519-generic.ko
kernel/arch/x86/kvm/kvm.ko: kernel/virt/lib/irqbypass.ko
kernel/arch/x86/kvm/kvm-intel.ko: kernel/arch/x86/kvm/kvm.ko kernel/virt/lib/irqbypass.ko
kernel/arch/x86/kvm/kvm-amd.ko: kernel/drivers/crypto/ccp/ccp.ko kernel/drivers/char/hw_random/rng-core.ko
kernel/arch/x86/kvm/kvm.ko kernel/virt/lib/irqbypass.ko
kernel/mm/hwpoison-inject.ko:
kernel/mm/zsmalloc.ko:
kernel/mm/z3fold.ko:
kernel/fs/nfs_common/nfs_acl.ko: kernel/net/sunrpc/sunrpc.ko
kernel/fs/nfs_common/grace.ko:
kernel/fs/nfs_common/nfs_ssc.ko:
kernel/fs/quota/quota_v1.ko:
kernel/fs/quota/quota_v2.ko: kernel/fs/quota/quota_tree.ko
kernel/fs/quota/quota_tree.ko:

```
kernel/fs/nls/nls_cp437.ko:
kernel/fs/nls/nls_cp737.ko:
kernel/fs/nls/nls_cp775.ko:
kernel/fs/nls/nls_cp850.ko:
kernel/fs/nls/nls_cp852.ko:
kernel/fs/nls/nls_cp855.ko:
kernel/fs/nls/nls_cp857.ko:
kernel/fs/nls/nls_cp860.ko:
kernel/fs/nls/nls_cp861.ko:
kernel/fs/nls/nls_cp862.ko:
--More-- (0%)
[q]
```

Il est possible d'obtenir des informations sur un module grâce à la commande **modinfo** :

```
root@debian11:~# modinfo bonding
filename:      /lib/modules/5.10.0-13-amd64/kernel/drivers/net/bonding/bonding.ko
author:       Thomas Davis, tadavis@lbl.gov and many others
description:  Ethernet Channel Bonding Driver
license:      GPL
alias:        rtnl-link-bond
depends:
retpoline:    Y
intree:       Y
name:         bonding
vermagic:     5.10.0-13-amd64 SMP mod_unload modversions
sig_id:       PKCS#7
signer:       Debian Secure Boot CA
sig_key:      4B:6E:F5:AB:CA:66:98:25:17:8E:05:2C:84:66:7C:CB:C0:53:1F:8C
sig_hashalgo: sha256
signature:    52:9F:34:1A:E8:04:22:2E:4C:92:17:82:D2:22:83:59:38:E2:EB:D6:
              EC:81:C3:CD:F8:C1:B9:5F:FC:C2:6D:27:BC:7B:91:13:87:5B:2E:92:
              77:36:A1:3D:F8:41:5C:9B:1D:62:E1:90:F7:48:F4:4E:7E:85:F5:54:
              CD:51:99:A6:C4:E0:FB:2C:4F:D6:5F:11:15:93:7E:30:62:A9:FA:46:
```

```

29:90:B2:58:A4:B4:34:8B:EA:EF:14:AD:D6:5D:6C:4D:03:C9:AC:0F:
F1:28:A0:65:3C:AA:C8:CA:83:59:25:F3:09:20:F2:74:7F:5A:7D:BB:
84:17:57:F5:E6:16:E5:CB:A9:A3:0C:F0:A9:04:4A:B2:63:98:80:30:
73:EE:39:3A:86:3E:67:3C:2F:5C:38:1E:35:F9:79:6C:F6:60:7B:B5:
9D:3C:EA:0D:0F:23:E1:EC:EA:42:4C:EF:EC:AB:3E:07:D3:35:D7:E4:
4A:E4:D1:7A:50:F8:C1:50:5E:93:0C:A2:7D:D1:77:E2:3A:97:74:BA:
D7:38:7E:C0:4C:36:66:00:78:03:E0:E3:20:46:70:13:7D:15:10:3A:
85:B7:CB:10:3B:9A:DF:1A:4E:64:3D:BF:EC:21:7B:B1:4C:37:CE:30:
0D:90:FE:73:87:20:40:83:60:02:0F:4F:4D:85:78:E7

```

```

parm:      max_bonds:Max number of bonded devices (int)
parm:      tx_queues:Max number of transmit queues (default = 16) (int)
parm:      num_grat_arp:Number of peer notifications to send on failover event (alias of num_unsol_na) (int)
parm:      num_unsol_na:Number of peer notifications to send on failover event (alias of num_grat_arp) (int)
parm:      miimon:Link check interval in milliseconds (int)
parm:      updelay:Delay before considering link up, in milliseconds (int)
parm:      downdelay:Delay before considering link down, in milliseconds (int)
parm:      use_carrier:Use netif_carrier_ok (vs MII ioctls) in miimon; 0 for off, 1 for on (default) (int)
parm:      mode:Mode of operation; 0 for balance-rr, 1 for active-backup, 2 for balance-xor, 3 for
broadcast, 4 for 802.3ad, 5 for balance-tlb, 6 for balance-alb (charp)
parm:      primary:Primary network device to use (charp)
parm:      primary_reselect:Reselect primary slave once it comes up; 0 for always (default), 1 for only if
speed of primary is better, 2 for only on active slave failure (charp)
parm:      lacp_rate:LACPDU tx rate to request from 802.3ad partner; 0 for slow, 1 for fast (charp)
parm:      ad_select:802.3ad aggregation selection logic; 0 for stable (default), 1 for bandwidth, 2 for
count (charp)
parm:      min_links:Minimum number of available links before turning on carrier (int)
parm:      xmit_hash_policy:balance-alb, balance-tlb, balance-xor, 802.3ad hashing method; 0 for layer 2
(default), 1 for layer 3+4, 2 for layer 2+3, 3 for encap layer 2+3, 4 for encap layer 3+4 (charp)
parm:      arp_interval:arp interval in milliseconds (int)
parm:      arp_ip_target:arp targets in n.n.n.n form (array of charp)
parm:      arp_validate:validate src/dst of ARP probes; 0 for none (default), 1 for active, 2 for backup, 3
for all (charp)
parm:      arp_all_targets:fail on any/all arp targets timeout; 0 for any (default), 1 for all (charp)
parm:      fail_over_mac:For active-backup, do not set all slaves to the same MAC; 0 for none (default), 1

```

```
for active, 2 for follow (charp)
parm:      all_slaves_active:Keep all frames received on an interface by setting active flag for all slaves;
0 for never (default), 1 for always. (int)
parm:      resend_igmp:Number of IGMP membership reports to send on link failure (int)
parm:      packets_per_slave:Packets to send per slave in balance-rr mode; 0 for a random slave, 1 packet
per slave (default), >1 packets per slave. (int)
parm:      lp_interval:The number of seconds between instances where the bonding driver sends learning
packets to each slaves peer switch. The default is 1. (uint)
```

Dernièrement, les fichiers dans le repertoire **/etc/modprobe.d** sont utilisés pour spécifier les options éventuelles à passer aux modules lors de leur chargement ainsi que les alias utilisés pour leur faire référence :

```
root@debian11:~# ls /etc/modprobe.d
mdadm.conf

root@debian11:~# cat /etc/modprobe.d/mdadm.conf
# mdadm module configuration file
# set start_ro=1 to make newly assembled arrays read-only initially,
# to prevent metadata writes. This is needed in order to allow
# resume-from-disk to work - new boot should not perform writes
# because it will be done behind the back of the system being
# resumed. See http://bugs.debian.org/415441 for details.

options md_mod start_ro=1
```

LAB #1 - Compilation et installation du noyau

Commencez par installer les paquets nécessaires :

```
root@debian11:~# apt-get update
Hit:1 http://deb.debian.org/debian bullseye InRelease
Get:2 http://deb.debian.org/debian bullseye-updates InRelease [39.4 kB]
```

```
Get:3 http://security.debian.org/debian-security bullseye-security InRelease [44.1 kB]
Fetched 83.5 kB in 1s (158 kB/s)
Reading package lists... Done
```

```
root@debian11:~# apt-get -y install build-essential linux-source bc kmod cpio flex libncurses5-dev libelf-dev
libssl-dev dwarves debconf-utils dpkg-dev debhelper ncurses-dev
```

1.1 - Déplacer /home

Créez une seule partition sur **/dev/sdb** :

```
root@debian11:~# fdisk /dev/sdb

Welcome to fdisk (util-linux 2.36.1).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table.
Created a new DOS disklabel with disk identifier 0xa10c368d.

Command (m for help): n
Partition type
   p   primary (0 primary, 0 extended, 4 free)
   e   extended (container for logical partitions)
Select (default p):

Using default response p.
Partition number (1-4, default 1):
First sector (2048-134217727, default 2048):
Last sector, +/-sectors or +/-size{K,M,G,T,P} (2048-134217727, default 134217727):

Created a new partition 1 of type 'Linux' and of size 64 GiB.
```

```
Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
```

Créez maintenant un système de fichiers ext4 sur **/dev/sdb1** :

```
root@debian11:~# mkfs.ext4 /dev/sdb1
mke2fs 1.46.2 (28-Feb-2021)
Discarding device blocks: done
Creating filesystem with 16776960 4k blocks and 4194304 inodes
Filesystem UUID: 24f1821e-1d5b-4256-8ee3-c9ee6b382ddc
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424

Allocating group tables: done
Writing inode tables: done
Creating journal (65536 blocks): done
Writing superblocks and filesystem accounting information: done
```

Editez le fichier **/etc/ssh/sshd_config** et modifiez la directive **PermitRootLogin** :

```
root@debian11:~# vi /etc/ssh/sshd_config
root@debian11:~# cat /etc/ssh/sshd_config
#      $OpenBSD: sshd_config,v 1.103 2018/04/09 20:41:22 tj Exp $

# This is the sshd server system-wide configuration file.  See
# sshd_config(5) for more information.

# This sshd was compiled with PATH=/usr/bin:/bin:/usr/sbin:/sbin

# The strategy used for options in the default sshd_config shipped with
# OpenSSH is to specify options with their default value where
```

```
# possible, but leave them commented. Uncommented options override the
# default value.
```

```
Include /etc/ssh/sshd_config.d/*.conf
```

```
#Port 22
```

```
#AddressFamily any
```

```
#ListenAddress 0.0.0.0
```

```
#ListenAddress ::
```

```
#HostKey /etc/ssh/ssh_host_rsa_key
```

```
#HostKey /etc/ssh/ssh_host_ecdsa_key
```

```
#HostKey /etc/ssh/ssh_host_ed25519_key
```

```
# Ciphers and keying
```

```
#RekeyLimit default none
```

```
# Logging
```

```
#SyslogFacility AUTH
```

```
#LogLevel INFO
```

```
# Authentication:
```

```
#LoginGraceTime 2m
```

```
PermitRootLogin yes
```

```
#StrictModes yes
```

```
#MaxAuthTries 6
```

```
#MaxSessions 10
```

```
...
```

Re-démarrez le serveur ssh :

```
root@debian11:~# systemctl restart sshd
```

```
root@debian11:~# systemctl status sshd
● ssh.service - OpenBSD Secure Shell server
   Loaded: loaded (/lib/systemd/system/ssh.service; enabled; vendor preset: enabled)
   Active: active (running) since Sun 2022-05-01 15:35:50 CEST; 6s ago
     Docs: man:sshd(8)
           man:sshd_config(5)
  Process: 2793 ExecStartPre=/usr/sbin/sshd -t (code=exited, status=0/SUCCESS)
 Main PID: 2794 (sshd)
    Tasks: 1 (limit: 4656)
   Memory: 1.1M
      CPU: 25ms
   CGroup: /system.slice/ssh.service
           └─2794 sshd: /usr/sbin/sshd -D [listener] 0 of 10-100 startups

May 01 15:35:50 debian11 systemd[1]: Starting OpenBSD Secure Shell server...
May 01 15:35:50 debian11 sshd[2794]: Server listening on 0.0.0.0 port 22.
May 01 15:35:50 debian11 sshd[2794]: Server listening on :: port 22.
May 01 15:35:50 debian11 systemd[1]: Started OpenBSD Secure Shell server.
```



A Faire - Déconnectez-vous en ssh. Connectez-vous directement en tant que root en ssh.

Montez **/dev/sdb1** sur /mnt :

```
root@debian11:~# mount /dev/sdb1 /mnt
```

Copiez le contenu de /home vers /mnt :

```
root@debian11:~# cp -a /home/* /mnt
```

Démontez /dev/sdb1 :

```
root@debian11:~# umount /mnt
```

Identifiez l'UUID de /dev/sdb1 :

```
root@debian11:~# ls -l /dev/disk/by-uuid/ | grep sdb1
lrwxrwxrwx 1 root root 10 May  1 15:31 24f1821e-1d5b-4256-8ee3-c9ee6b382ddc -> ../../sdb1
```

Editez le fichier **/etc/fstab** en ajoutant la ligne pour le montage de /home :

```
root@debian11:~# vi /etc/fstab
root@debian11:~# cat /etc/fstab
# /etc/fstab: static file system information.
#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# systemd generates mount units based on this file, see systemd.mount(5).
# Please run 'systemctl daemon-reload' after making changes here.
#
# <file system> <mount point>  <type>  <options>          <dump>  <pass>
# / was on /dev/sda1 during installation
UUID=9887a74f-a680-4bde-8f04-db5ae9ea186e /          ext4      errors=remount-ro 0          1
UUID=24f1821e-1d5b-4256-8ee3-c9ee6b382ddc /home      ext4      defaults          0 0
# swap was on /dev/sda5 during installation
UUID=1f9439f5-4b19-49b1-b292-60c2c674cee9 none        swap      sw              0          0
/dev/sr0      /media/cdrom0  udf,iso9660 user,noauto     0          0
```

Créez le point de montage /home :

```
root@debian11:~# rm -rf /home
root@debian11:~# mkdir /home
```

Montez /dev/sdb1 :

```
root@debian11:~# mount -a

root@debian11:~# mount
sysfs on /sys type sysfs (rw,nosuid,nodev,noexec,relatime)
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
udev on /dev type devtmpfs (rw,nosuid,relatime,size=1986968k,nr_inodes=496742,mode=755)
devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000)
tmpfs on /run type tmpfs (rw,nosuid,nodev,noexec,relatime,size=402560k,mode=755)
/dev/sda1 on / type ext4 (rw,relatime,errors=remount-ro)
securityfs on /sys/kernel/security type securityfs (rw,nosuid,nodev,noexec,relatime)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
tmpfs on /run/lock type tmpfs (rw,nosuid,nodev,noexec,relatime,size=5120k)
cgroup2 on /sys/fs/cgroup type cgroup2 (rw,nosuid,nodev,noexec,relatime,nsdelegate,memory_recursiveprot)
pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
none on /sys/fs/bpf type bpf (rw,nosuid,nodev,noexec,relatime,mode=700)
systemd-1 on /proc/sys/fs/binfmt_misc type autofs
(rw,relatime,fd=29,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=11732)
mqueue on /dev/mqueue type mqueue (rw,nosuid,nodev,noexec,relatime)
debugfs on /sys/kernel/debug type debugfs (rw,nosuid,nodev,noexec,relatime)
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime,pagesize=2M)
tracefs on /sys/kernel/tracing type tracefs (rw,nosuid,nodev,noexec,relatime)
configfs on /sys/kernel/config type configfs (rw,nosuid,nodev,noexec,relatime)
fusectl on /sys/fs/fuse/connections type fusectl (rw,nosuid,nodev,noexec,relatime)
tmpfs on /run/user/0 type tmpfs (rw,nosuid,nodev,relatime,size=402556k,nr_inodes=100639,mode=700)
/dev/sdb1 on /home type ext4 (rw,relatime)
```

Notez la taille de /home :

```
root@debian11:~# df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            1.9G     0   1.9G   0% /dev
tmpfs           394M  892K   393M   1% /run
```

/dev/sda1	31G	4.1G	25G	14%	/
tmpfs	2.0G	0	2.0G	0%	/dev/shm
tmpfs	5.0M	0	5.0M	0%	/run/lock
tmpfs	394M	36K	394M	1%	/run/user/0
/dev/sdb1	63G	1.4M	60G	1%	/home



A Faire - Fermez la session de root et connectez-vous en tant que trainee en ssh.

1.2 - Télécharger le Code Source du Noyau

Le code source est disponible sur le site www.kernel.org :

```
trainee@debian11:~$ wget https://mirrors.edge.kernel.org/pub/linux/kernel/v5.x/linux-5.11.1.tar.gz
```

Désarchivez le tar.xz que vous avez téléchargé :

```
trainee@debian8:~$ tar xf linux-5.11.1.tar.gz
```

Notez que l'occupation disque du code source du noyau linux-5.11.1 est de 1.2 Go :

```
trainee@debian11:~$ du -hs ./linux-5.11.1
1.2G    ./linux-5.11.1
```

1.3 - Configurer le Noyau

Le fichier **Makefile** contient le nom du noyau et spécifie les informations suivantes :

- VERSION,

- PATCHLEVEL,
- SUBLEVEL,
- EXTRAVERSION.

Les trois premières informations sont gérées par **kernel.org** et Linus Torvalds en personne tandis que l'EXTRAVERSION est gérée par Debian :

```
trainee@debian11:~$ cat ./linux-5.11.1/Makefile | head
# SPDX-License-Identifier: GPL-2.0
VERSION = 5
PATCHLEVEL = 11
SUBLEVEL = 1
EXTRAVERSION =
NAME = □ Valentine's Day Edition □

# *DOCUMENTATION*
# To see a list of typical targets execute "make help"
# More info can be located in ./README
```



Important - La version 2.6 du noyau a vu le jour en **2003**. Les **SUBLEVEL** se suivaient régulièrement. Avec la version 2.6 du noyau, la valeur paire du **PATCHLEVEL** indiquait que le noyau était stable. Quand vous recompilez le noyau à partir des sources, vous devez modifier la valeur de l'EXTRAVERSION. Le passage à la version 3.0 fut décidé par Linus Torvalds à l'occasion des 20 ans du noyau Linux.

Utilisez maintenant la commande **make defconfig** pour créer le fichier de configuration .config :

```
trainee@debian11:~$ su -
Password: fenestros
root@debian11:~# cd /home/trainee/linux-5.11.1/
root@debian11:/home/trainee/linux-5.11.1# make defconfig
HOSTCC  scripts/basic/fixdep
HOSTCC  scripts/kconfig/conf.o
HOSTCC  scripts/kconfig/confdata.o
```

```
HOSTCC  scripts/kconfig/expr.o
LEX      scripts/kconfig/lexer.lex.c
YACC     scripts/kconfig/parser.tab.[ch]
HOSTCC   scripts/kconfig/lexer.lex.o
HOSTCC   scripts/kconfig/parser.tab.o
HOSTCC   scripts/kconfig/preprocess.o
HOSTCC   scripts/kconfig/symbol.o
HOSTCC   scripts/kconfig/util.o
HOSTLD   scripts/kconfig/conf
*** Default configuration is based on 'x86_64_defconfig'
#
# configuration written to .config
#
```

Ce fichier est configuré par une des trois commandes suivantes :

- make config
- make menuconfig
- make xconfig

Dans ce fichier, vous pouvez constater la présence de lignes correspondantes à des fonctionnalités suivies par une lettre ou une valeur. Dans le cas d'une lettre, la signification est la suivante :

- **y**
 - la fonctionnalité est incluse dans le noyau monolithique ou dans le cas d'une dépendance d'un module, dans le module concerné,
- **m**
 - la fonctionnalité est incluse en tant que module,
- **n**
 - la fonctionnalité n'est pas incluse. Cette option est rarement visible car dans bien les cas, la fonctionnalité est simplement commentée dans le fichier lui-même.

```
root@debian11:/home/trainee/linux-5.11.1# more .config
#
# Automatically generated file; DO NOT EDIT.
# Linux/x86 5.11.1 Kernel Configuration
```

```
#
CONFIG_CC_VERSION_TEXT="gcc (Debian 10.2.1-6) 10.2.1 20210110"
CONFIG_CC_IS_GCC=y
CONFIG_GCC_VERSION=100201
CONFIG_LD_VERSION=235020000
CONFIG_CLANG_VERSION=0
CONFIG_LLD_VERSION=0
CONFIG_CC_CAN_LINK=y
CONFIG_CC_CAN_LINK_STATIC=y
CONFIG_CC_HAS_ASM_GOTO=y
CONFIG_CC_HAS_ASM_INLINE=y
CONFIG_IRQ_WORK=y
CONFIG_BUILDTIME_TABLE_SORT=y
CONFIG_THREAD_INFO_IN_TASK=y
```

```
#
# General setup
#
CONFIG_INIT_ENV_ARG_LIMIT=32
# CONFIG_COMPILE_TEST is not set
CONFIG_LOCALVERSION=""
# CONFIG_LOCALVERSION_AUTO is not set
CONFIG_BUILD_SALT=""
CONFIG_HAVE_KERNEL_GZIP=y
CONFIG_HAVE_KERNEL_BZIP2=y
CONFIG_HAVE_KERNEL_LZMA=y
CONFIG_HAVE_KERNEL_XZ=y
CONFIG_HAVE_KERNEL_LZO=y
CONFIG_HAVE_KERNEL_LZ4=y
CONFIG_HAVE_KERNEL_ZSTD=y
CONFIG_KERNEL_GZIP=y
# CONFIG_KERNEL_BZIP2 is not set
# CONFIG_KERNEL_LZMA is not set
# CONFIG_KERNEL_XZ is not set
```

```
# CONFIG_KERNEL_LZ0 is not set
# CONFIG_KERNEL_LZ4 is not set
# CONFIG_KERNEL_ZSTD is not set
CONFIG_DEFAULT_INIT=""
CONFIG_DEFAULT_HOSTNAME="(none)"
CONFIG_SWAP=y
CONFIG_SYSVIPC=y
CONFIG_SYSVIPC_SYSCTL=y
CONFIG_POSIX_MQUEUE=y
CONFIG_POSIX_MQUEUE_SYSCTL=y
# CONFIG_WATCH_QUEUE is not set
CONFIG_CROSS_MEMORY_ATTACH=y
CONFIG_USELIB=y
CONFIG_AUDIT=y
CONFIG_HAVE_ARCH_AUDITSYSCALL=y
CONFIG_AUDITSYSCALL=y

#
# IRQ subsystem
#
CONFIG_GENERIC_IRQ_PROBE=y
--More-- (1%)
[q]
```

1.4 - Compiler le Noyau

La compilation du noyau peut prendre beaucoup de temps. La commande utilisée est la suivante :

```
root@debian11:/home/trainee/linux-5.11.1# make deb-pkg clean
...
dpkg-buildpackage: info: full upload (original source is included)
  CLEAN    arch/x86/entry/vdso
  CLEAN    arch/x86/kernel/cpu
```

```

CLEAN arch/x86/kernel
CLEAN arch/x86/realmode/rm
CLEAN arch/x86/lib
CLEAN certs
CLEAN drivers/firmware/efi/libstub
CLEAN drivers/scsi
CLEAN drivers/tty/vt
CLEAN drivers/video/logo
CLEAN lib
CLEAN net/wireless
CLEAN security/selinux
CLEAN usr/include
CLEAN usr
CLEAN arch/x86/boot/compressed
CLEAN arch/x86/boot
CLEAN arch/x86/tools
CLEAN vmlinux.symvers modules.builtin modules.builtin.modinfo

```

A l'issu du processus, les paquets du nouveau noyau se trouvent dans le répertoire **/home/trainee** :

```

root@debian11:/home/trainee/linux-5.11.1# cd ..
root@debian11:/home/trainee# ls
Desktop      linux-5.11.1                                linux-5.11.1_5.11.1-1.diff.gz    linux-5.11.1.tar.gz
linux-libc-dev_5.11.1-1_amd64.deb  Public
Documents    linux-5.11.1_5.11.1-1_amd64.buildinfo  linux-5.11.1_5.11.1-1.dsc      linux-
headers-5.11.1_5.11.1-1_amd64.deb  Music                                     Templates
Downloads    linux-5.11.1_5.11.1-1_amd64.changes    linux-5.11.1_5.11.1.orig.tar.gz  linux-
image-5.11.1_5.11.1-1_amd64.deb     Pictures                                 Videos

```

1.5 - Installer le Noyau

Installez maintenant les paquets **deb** :

```
root@debian11:/home/trainee# dpkg -i /home/trainee/linux*.deb
```

Constatez la création d'un nouveau grub.cfg :

```
root@debian11:/home/trainee# grep 5.11.1 /boot/grub/grub.cfg
echo      'Loading Linux 5.11.1 ...'
linux     /boot/vmlinuz-5.11.1 root=UUID=9887a74f-a680-4bde-8f04-db5ae9ea186e ro  quiet
initrd    /boot/initrd.img-5.11.1
menuentry 'Debian GNU/Linux, with Linux 5.11.1' --class debian --class gnu-linux --class gnu --class os
$menuentry_id_option 'gnulinux-5.11.1-advanced-9887a74f-a680-4bde-8f04-db5ae9ea186e' {
    echo      'Loading Linux 5.11.1 ...'
    linux     /boot/vmlinuz-5.11.1 root=UUID=9887a74f-a680-4bde-8f04-db5ae9ea186e ro  quiet
    initrd    /boot/initrd.img-5.11.1
    menuentry 'Debian GNU/Linux, with Linux 5.11.1 (recovery mode)' --class debian --class gnu-linux --class
gnu --class os $menuentry_id_option 'gnulinux-5.11.1-recovery-9887a74f-a680-4bde-8f04-db5ae9ea186e' {
    echo      'Loading Linux 5.11.1 ...'
    linux     /boot/vmlinuz-5.11.1 root=UUID=9887a74f-a680-4bde-8f04-db5ae9ea186e ro single
    initrd    /boot/initrd.img-5.11.1
```

1.6 - Désinstaller un Noyau

Lister les noyaux installés :

```
root@debian11:/home/trainee# dpkg -l | grep -i "linux-image*" | awk '{print $2}'
linux-image-5.10.0-13-amd64
linux-image-5.11.1
linux-image-amd64
```

Le noyau se désinstalle comme tout autre paquet :

```
root@debian11:/home/trainee# apt-get -y purge "linux-image-5.11.1"
Reading package lists... Done
```

```
Building dependency tree... Done
Reading state information... Done
The following packages will be REMOVED:
  linux-image-5.11.1*
0 upgraded, 0 newly installed, 1 to remove and 0 not upgraded.
After this operation, 15.1 MB disk space will be freed.
(Reading database ... 118012 files and directories currently installed.)
Removing linux-image-5.11.1 (5.11.1-1) ...
update-initramfs: Deleting /boot/initrd.img-5.11.1
Generating grub configuration file ...
Found background image: /usr/share/images/desktop-base/desktop-grub.png
Found linux image: /boot/vmlinuz-5.10.0-13-amd64
Found initrd image: /boot/initrd.img-5.10.0-13-amd64
done
(Reading database ... 117967 files and directories currently installed.)
Purging configuration files for linux-image-5.11.1 (5.11.1-1) ...
```

Vérifiez que le fichier grub.cfg a été modifié :

```
root@debian11:/home/trainee# grep 5.11.1 /boot/grub/grub.cfg
root@debian11:/home/trainee#
```

Dernièrement, listez les noyaux disponibles :

```
root@debian11:/home/trainee# dpkg -l | grep -i "linux-image*" | awk '{print $2}'
linux-image-5.10.0-13-amd64
linux-image-amd64
```

LAB #2 - Mise à Jour du Noyau avec le Gestionnaire des Paquets

Afin d'obtenir la **dernière version du noyau pré-compilée** par Debian, il convient d'ajouter les dépôts **unstable** :

```
root@debian11:/home/trainee# echo "deb http://deb.debian.org/debian unstable main contrib non-free" | sudo tee -a /etc/apt/sources.list
deb http://deb.debian.org/debian unstable main contrib non-free
root@debian11:/home/trainee# echo "deb-src http://deb.debian.org/debian unstable main contrib non-free" | sudo tee -a /etc/apt/sources.list
deb-src http://deb.debian.org/debian unstable main contrib non-free
```

Afin de ne faire que des mises à jour du **noyau** à partir de ce dépôt, créez le fichier **/etc/apt/preferences** :

```
root@debian11:/home/trainee# vi /etc/apt/preferences
root@debian11:/home/trainee# cat /etc/apt/preferences
Package: *
Pin: release a=bullseye
Pin-Priority: 500

Package: linux-image-amd64
Pin: release a=unstable
Pin-Priority: 1000

Package: *
Pin: release a=unstable
Pin-Priority: 100
```

Mettez à jour apt :

```
root@debian11:/home/trainee# apt update
Hit:1 http://security.debian.org/debian-security bullseye-security InRelease
Hit:2 http://deb.debian.org/debian bullseye InRelease
Get:3 http://deb.debian.org/debian bullseye-updates InRelease [39.4 kB]
Get:4 http://deb.debian.org/debian unstable InRelease [165 kB]
Get:5 http://deb.debian.org/debian unstable/contrib Sources [57.6 kB]
Get:6 http://deb.debian.org/debian unstable/main Sources [9,700 kB]
Get:7 http://deb.debian.org/debian unstable/non-free Sources [89.6 kB]
Get:8 http://deb.debian.org/debian unstable/main amd64 Packages [9,154 kB]
```

```
Get:9 http://deb.debian.org/debian unstable/main Translation-en [6,780 kB]
Get:10 http://deb.debian.org/debian unstable/contrib amd64 Packages [65.4 kB]
Get:11 http://deb.debian.org/debian unstable/contrib Translation-en [55.9 kB]
Get:12 http://deb.debian.org/debian unstable/non-free amd64 Packages [112 kB]
Get:13 http://deb.debian.org/debian unstable/non-free Translation-en [103 kB]
Fetched 26.3 MB in 5s (4,955 kB/s)
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
2 packages can be upgraded. Run 'apt list --upgradable' to see them.
```

Installez ensuite les mises-à-jour du noyau :

```
root@debian11:/home/trainee# apt -y upgrade
```

Re-démarrez votre VM :

```
root@debian11:/home/trainee# reboot
```

Connectez-vous de nouveau à la VM et contrôlez la version du noyau :

```
trainee@debian11:~$ uname -a
Linux debian11 5.17.0-1-amd64 #1 SMP PREEMPT Debian 5.17.3-1 (2022-04-18) x86_64 GNU/Linux
```

Déconnectez-vous et re-connectez-vous directement en tant que root.

LAB #3 - Gestion des Quotas

Sous Linux il est possible de mettre en place des quotas par utilisateur et par groupe. Ceci étant, Linux ne sait pas gérer des quotas par répertoire, uniquement des quotas par partition. L'administrateur met souvent des quotas en place sur l'arborescence de /home pour limiter l'espace de stockage occupé par les utilisateurs.

Commencez par vérifiez que le paquet **quota** est bien installé :

```
root@debian11:~# dpkg --get-selections | grep quota
root@debian11:~# apt -y install quota
```

Editez le fichier **/etc/fstab** en ajoutant les options **usrquota** et **grpquota** à la ligne **/home** :

```
root@debian11:~# vi /etc/fstab
root@debian11:~# cat /etc/fstab
# /etc/fstab: static file system information.
#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# systemd generates mount units based on this file, see systemd.mount(5).
# Please run 'systemctl daemon-reload' after making changes here.
#
# <file system> <mount point>    <type>  <options>          <dump>  <pass>
# / was on /dev/sda1 during installation
UUID=9887a74f-a680-4bde-8f04-db5ae9ea186e /          ext4      errors=remount-ro 0      1
UUID=24f1821e-1d5b-4256-8ee3-c9ee6b382ddc /home      ext4      defaults,usrquota,grpquota 0 0
# swap was on /dev/sda5 during installation
UUID=1f9439f5-4b19-49b1-b292-60c2c674cee9 none       swap      sw           0      0
/dev/sr0      /media/cdrom0    udf,iso9660 user,noauto   0      0
```

Démontez puis remontez /home :

```
root@debian11:~# umount /home
root@debian11:~# mount -a
```

Vérifiez ensuite que les options soient prises en compte :

```
root@debian11:~# cat /etc/mtab
```

```
sysfs /sys sysfs rw,nosuid,nodev,noexec,relatime 0 0
proc /proc proc rw,nosuid,nodev,noexec,relatime 0 0
udev /dev devtmpfs rw,nosuid,relatime,size=1976600k,nr_inodes=494150,mode=755,inode64 0 0
devpts /dev/pts devpts rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000 0 0
tmpfs /run tmpfs rw,nosuid,nodev,noexec,relatime,size=402384k,mode=755,inode64 0 0
/dev/sda1 / ext4 rw,relatime,errors=remount-ro 0 0
securityfs /sys/kernel/security securityfs rw,nosuid,nodev,noexec,relatime 0 0
tmpfs /dev/shm tmpfs rw,nosuid,nodev,inode64 0 0
tmpfs /run/lock tmpfs rw,nosuid,nodev,noexec,relatime,size=5120k,inode64 0 0
cgroup2 /sys/fs/cgroup cgroup2 rw,nosuid,nodev,noexec,relatime,nsdelegate,memory_recursiveprot 0 0
pstore /sys/fs/pstore pstore rw,nosuid,nodev,noexec,relatime 0 0
bpf /sys/fs/bpf bpf rw,nosuid,nodev,noexec,relatime,mode=700 0 0
systemd-1 /proc/sys/fs/binfmt_misc autofs
rw,relatime,fd=29,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=13172 0 0
mqueue /dev/mqueue mqueue rw,nosuid,nodev,noexec,relatime 0 0
hugetlbfs /dev/hugepages hugetlbfs rw,relatime,pagesize=2M 0 0
tracefs /sys/kernel/tracing tracefs rw,nosuid,nodev,noexec,relatime 0 0
debugfs /sys/kernel/debug debugfs rw,nosuid,nodev,noexec,relatime 0 0
configfs /sys/kernel/config configfs rw,nosuid,nodev,noexec,relatime 0 0
fusectl /sys/fs/fuse/connections fusectl rw,nosuid,nodev,noexec,relatime 0 0
tmpfs /run/user/0 tmpfs rw,nosuid,nodev,relatime,size=402380k,nr_inodes=100595,mode=700,inode64 0 0
/dev/sdb1 /home ext4 rw,relatime,quota,usrquota,grpquota 0 0
```

3.1 - La Commande quotacheck

Pour activer les quotas sur /home, il convient d'utiliser la commande **quotacheck** :

```
root@debian11:~# quotacheck -cugvm -f /dev/sdb1
quotacheck: Your kernel probably supports ext4 quota feature but you are using external quota files. Please
switch your filesystem to use ext4 quota feature as external quota files on ext4 are deprecated.
quotacheck: Scanning /dev/sdb1 [/home] done
quotacheck: Cannot stat old user quota file /home/aquota.user: No such file or directory. Usage will not be
subtracted.
```

```
quotacheck: Cannot stat old group quota file /home/aquota.group: No such file or directory. Usage will not be
subtracted.
quotacheck: Cannot stat old user quota file /home/aquota.user: No such file or directory. Usage will not be
subtracted.
quotacheck: Cannot stat old group quota file /home/aquota.group: No such file or directory. Usage will not be
subtracted.
quotacheck: Checked 7199 directories and 81583 files
quotacheck: Old file not found.
quotacheck: Old file not found.
```

Les options de la commande quotacheck sont :

```
root@debian11:~# quotacheck --help
Utility for checking and repairing quota files.
quotacheck [-gucbfinvdmMR] [-F <quota-format>] filesystem|-a

-u, --user          check user files
-g, --group         check group files
-c, --create-files  create new quota files
-b, --backup        create backups of old quota files
-f, --force         force check even if quotas are enabled
-i, --interactive   interactive mode
-n, --use-first-dquot use the first copy of duplicated structure
-v, --verbose       print more information
-d, --debug         print even more messages
-m, --no-remount    do not remount filesystem read-only
-M, --try-remount   try remounting filesystem read-only,
                   continue even if it fails
-R, --exclude-root  exclude root when checking all filesystems
-F, --format=formatname check quota files of specific format
-a, --all           check all filesystems
-h, --help          display this message and exit
-V, --version       display version information and exit
```

```
Bugs to jack@suse.cz
```

Les quotas ont été activés et les fichiers **aquota.user** et **aquota.group** ont été créés dans le répertoire `/home` :

```
root@debian11:~# ls -la /home
total 44
drwxr-xr-x  4 root    root    4096 May  1 18:05 .
drwxr-xr-x 20 root    root    4096 May  1 17:41 ..
-rw-----  1 root    root    7168 May  1 18:05 aquota.group
-rw-----  1 root    root    7168 May  1 18:05 aquota.user
drwx-----  2 root    root   16384 May  1 15:31 lost+found
drwxr-xr-x 15 trainee trainee 4096 May  1 16:28 trainee
```

Créez maintenant un utilisateur **fenestros** avec le mot de passe **fenestros** :

```
root@debian11:~# groupadd fenestros && useradd -m fenestros -c Fenestr0s -d /home/fenestros -g fenestros -s /bin/bash
root@debian11:~# passwd fenestros
New password: fenestros
Retype new password: fenestros
passwd: password updated successfully
```

3.2 - La Commande edquota

Mettez en place maintenant un quota de 10Mo pour l'utilisateur **fenestros** :

```
root@debian8:~# EDITOR=/usr/bin/vi
root@debian8:~# export EDITOR
root@debian8:~# edquota -u fenestros -f /home
```

L'éditeur **vi** se lance et vous obtiendrez un résultat similaire à celui-ci :

Disk quotas for user fenestros (uid 1001):

Filesystem	blocks	soft	hard	inodes	soft	hard
/dev/sdb1	0	0	0	0	0	0

Modifiez ce fichier ainsi :

Disk quotas for user fenestros (uid 1001):

Filesystem	blocks	soft	hard	inodes	soft	hard
/dev/sdb1	0	8000	10000	0	0	0

Les options de la commande **edquota** sont :

```
root@debian11:~# edquota --help
```

edquota: Usage:

```
edquota [-rm] [-u] [-F formatname] [-p username] [-f filesystem] username ...
edquota [-rm] -g [-F formatname] [-p groupname] [-f filesystem] groupname ...
edquota [-rm] -P [-F formatname] [-p projectname] [-f filesystem] projectname ...
edquota [-u|g|-P] [-F formatname] [-f filesystem] -t
edquota [-u|g|-P] [-F formatname] [-f filesystem] -T username|groupname|projectname ...
```

-u, --user	edit user data
-g, --group	edit group data
-P, --project	edit project data
-r, --remote	edit remote quota (via RPC)
-m, --no-mixed-pathnames	trim leading slashes from NFSv4 mountpoints
-F, --format=formatname	edit quotas of a specific format
-p, --prototype=name	copy data from a prototype user/group
--always-resolve	always try to resolve name, even if it is composed only of digits
-f, --filesystem=filesystem	edit data only on a specific filesystem
-t, --edit-period	edit grace period
-T, --edit-times	edit grace time of a user/group
-h, --help	display this help text and exit
-V, --version	display version information and exit

Bugs to: jack@suse.cz



Important - Pour mettre en place un quota par group, la procédure est similaire. Il suffit d'utiliser l'option -g de la commande edquota.

3.3 - La Commande quotaon

Appliquez maintenant les quotas :

```
root@debian11:~# quotaon -a
```

Les options de la commande **quotaon** sont :

```
root@debian11:~# quotaon --help
quotaon: Usage:
    quotaon [-guPvp] [-F quotaformat] [-x state] -a
    quotaon [-guPvp] [-F quotaformat] [-x state] filesystems ...

-a, --all                turn quotas on for all filesystems
-f, --off                turn quotas off
-u, --user               operate on user quotas
-g, --group              operate on group quotas
-P, --project            operate on project quotas
-p, --print-state        print whether quotas are on or off
-x, --xfs-command=cmd    perform XFS quota command
-F, --format=formatname  operate on specific quota format
-v, --verbose            print more messages
-h, --help               display this help text and exit
-V, --version            display version information and exit
```

De cette manière vous avez mis en place un quota **souple** pour fenestros de 8 000 Ko et un quota **stricte** de 10 000 Ko.

Quand l'utilisateur fenestros aura dépassé le quota **souple**, il recevra un message d'avertissement. Quand il dépasse le quota **stricte**, il ne pourra plus enregistrer dans /home, sauf dans le cas où il supprime des fichiers pour retomber en dessous de la limite **stricte**.

Il est à noter que vous pouvez soit mettre en place un quota en taille, soit mettre en place un quota basé sur le nombre d'inodes utilisés par l'utilisateur.



Important - La commande pour désactiver les quotas est **quotaoff**.

3.4 - La Commande repquota

Pour visualiser les quotas utilisez la commande **repquota** :

```
root@debian11:~# repquota /home
*** Report for user quotas on device /dev/sdb1
Block grace time: 7days; Inode grace time: 7days
```

User	Block limits				File limits			
	used	soft	hard	grace	used	soft	hard	grace
root	-- 241004	0	0		2528	0	0	
trainee	-- 1409108	0	0		86253	0	0	



Important - Notez que l'utilisateur fenestros ne figure pas dans la liste. Le quota n'est pas visible tant que l'utilisateur ne s'est pas connecté pour la première fois. Notez aussi les période de grâce de **7** jours.

Les options de la commande **repquota** sont :

```
root@debian11:~# repquota --help
```

repquota: Utility for reporting quotas.

Usage:

repquota [-vugsi] [-c|C] [-t|n] [-F quotaformat] [-O (default | xml | csv)] (-a | mntpoint)

-v, --verbose	display also users/groups without any usage
-u, --user	display information about users
-g, --group	display information about groups
-P, --project	display information about projects
-s, --human-readable[=units]	display numbers in human friendly units (MB, GB, ...). Units can be also specified explicitly by an optional argument in format [kgt],[kgt] where the first character specifies space units and the second character specifies inode units
-t, --truncate-names	truncate names to 9 characters
-p, --raw-grace	print grace time in seconds since epoch
-n, --no-names	do not translate uid/gid to name
-i, --no-autofs	avoid autofs mountpoints
-c, --cache	translate big number of ids at once
-C, --no-cache	translate ids one by one
-F, --format=formatname	report information for specific format
-O, --output=format	format output as xml or csv
-a, --all	report information for all mount points with quotas
-h, --help	display this help message and exit
-V, --version	display version information and exit

Bugs to jack@suse.cz

3.5 - La Commande quota

Pour visualiser les quotas d'un utilisateur spécifique, il convient d'utiliser la commande **quota** :

```
root@debian11:~# quota fenestros
```

```
Disk quotas for user fenestros (uid 1001): no limited resources used
```

```
root@debian11:~# su - fenestros
```

```
fenestros@debian11:~$ touch test
```

```
fenestros@debian11:~$ exit
```

```
logout
```

```
root@debian11:~# quota fenestros
```

```
Disk quotas for user fenestros (uid 1001):
```

Filesystem	blocks	quota	limit	grace	files	quota	limit	grace
/dev/sdb1	4	8000	10000		2	0	0	

Les options de la commande **quota** sont :

```
root@debian11:~# quota --help
```

```
quota: Usage: quota [-guPqvswim] [-l | [-Q | -A]] [-F quotaformat]
```

```
    quota [-qvswim] [-l | [-Q | -A]] [-F quotaformat] -u username ...
```

```
    quota [-qvswim] [-l | [-Q | -A]] [-F quotaformat] -g groupname ...
```

```
    quota [-qvswim] [-l | [-Q | -A]] [-F quotaformat] -P projectname ...
```

```
    quota [-qvswugPQm] [-F quotaformat] -f filesystem ...
```

```
-u, --user                display quota for user
```

```
-g, --group               display quota for group
```

```
-P, --project             display quota for project
```

```
-q, --quiet               print more terse message
```

```
-v, --verbose             print more verbose message
```

```
-s, --human-readable[=units] display numbers in human friendly units (MB, GB, ...). Units can be also specified explicitly by an optional argument in format [kgt],[kgt] where the first character specifies space units and the second character specifies inode units
```

```
--always-resolve          always try to translate name to id, even if it is
```

	composed of only digits
-w, --no-wrap	do not wrap long lines
-p, --raw-grace	print grace time in seconds since epoch
-l, --local-only	do not query NFS filesystems
-Q, --quiet-refuse	do not print error message when NFS server does not respond
-i, --no-autofs	do not query autofs mountpoints
-F, --format=formatname	display quota of a specific format
-f, --filesystem-list	display quota information only for given filesystems
--filesystem=path	display quota information only for given filesystem, remaining command line arguments are still treated as user/group/project names
-A, --all-nfs	display quota for all NFS mountpoints
-m, --no-mixed-pathnames	trim leading slashes from NFSv4 mountpoints
--show-mntpoint	show mount point of the file system in output
--hide-device	do not show file system device in output
-h, --help	display this help message and exit
-V, --version	display version information and exit

Bugs to: jack@suse.cz

3.6 - La Commande warnquota

La commande **warnquota** vérifie le ou les disques et envoie un message par mail à tout utilisateur qui a dépassé la limite soft. Elle est enrègle générale appelée par un job cron. Cependant elle peut aussi est appelée d'une manière interactive.

Les options de la commande **warnquota** sont :

```
root@debian11:~# warnquota --help
warnquota: Usage:
  warnquota [-ugsid] [-F quotaformat] [-c configfile] [-q quotatabfile] [-a adminsfile] [filesystem...]
```

```
-u, --user          warn users
-g, --group         warn groups
-s, --human-readable[=units] display numbers in human friendly units (MB,
                        GB, ...). Units can be also specified
                        explicitly by an optional argument in format
                        [kgt],[kgt] where the first character specifies
                        space units and the second character specifies
                        inode units
-i, --no-autofs     avoid autofs mountpoints
-d, --no-details    do not send quota information itself
-F, --format=formatname use quotafiles of specific format
-c, --config=config-file non-default config file
-q, --quota-tab=quotatab-file non-default quotatab
-a, --admins-file=admins-file non-default admins file
-I, --ignore-config-errors ignore unknown statements in config file
-h, --help          display this help message and exit
-V, --version        display version information and exit
```

warnquota: Bugs to jack@suse.cz