

Version : **2022.01**

Updated: 2022/06/05 17:55

LCE511 - Managing Integrated Peripherals

Contents

- **LCE511 - Managing Integrated Peripherals**
 - Contents
 - Special Files
 - LAB #1 - Commands
 - 1.1 - The lspci Command
 - 1.2 - The lsusb Command
 - 1.3 - The dmidecode Command
 - LAB #2 - The sysctl Command
 - 2.1 - The /proc Directory
 - Files
 - /proc/cpuinfo
 - /proc/interrupts
 - /proc/dma
 - /proc/ioports
 - /proc/devices
 - /proc/module
 - /proc/diskstats
 - /proc/partitions
 - /proc/swaps
 - /proc/loadavg
 - /proc/meminfo
 - /proc/version
 - Directories

- ide/scsi
- acpi
- bus
- net
- sys
- 2.2 - Using the sysctl Command
- LAB #3 - Interpreting Information in /proc
 - 3.1 - free
 - 3.2 - uptime ou w
 - 3.3 - iostat
 - 3.4 - hdparm
 - 3.5 - vmstat
 - 3.6 - mpstat
 - 3.7 - sar
- USB Modules
- udev
 - The udevadm Command
- The /sys Filesystem
- LAB #4 - Limiting Ressources
 - 4.1 - ulimit
 - 4.2 - CGroups
 - Limiting Memeory Usage
 - The cgcreate Command
 - The cgdelete Command
 - The /etc/cgconfig.conf File
 - The cgconfigparser Command

Special Files

In a PC, peripherals are connected to a **controler** which communicates with the processor via a **bus**. The controller and associated peripherals require specific drivers. Under Linux, these drivers are normally supplied as kernel **modules**. Each peripheral is represented by a special file in the **/dev** directory. Each special file contains the information required by the system in order for it to use the driver.



Peripherals that require the system to be halted prior to plugging or unplugging them are referred to as **Cold Plug Devices**. Peripherals that can be plugged/unplugged whilst the system is running are referred to as **Hot Plug Devices**.

The following output shows the typical content of the /dev directory:

```
[root@centos8 ~]# ls -l /dev | more
total 0
crw-r--r--. 1 root root    10, 235 Jun 28 02:04 autofs
drwxr-xr-x. 2 root root    180 Jun 28 02:04 block
drwxr-xr-x. 2 root root    100 Jun 28 02:04 bsg
drwxr-xr-x. 3 root root     60 Jun 28 02:04 bus
lrwxrwxrwx. 1 root root     3 Jun 28 02:04 cdrom -> sr0
drwxr-xr-x. 2 root root   2940 Jun 28 02:04 char
drwxr-xr-x. 2 root root     80 Jun 28 02:04 cl_centos8
crw----- 1 root root     5,  1 Jun 28 02:04 console
lrwxrwxrwx. 1 root root    11 Jun 28 02:04 core -> /proc/kcore
drwxr-xr-x. 10 root root    200 Jun 28 02:04 cpu
crw----- 1 root root    10,  62 Jun 28 02:04 cpu_dma_latency
drwxr-xr-x. 6 root root    120 Jun 28 02:04 disk
brw-rw---- 1 root disk   253,  0 Jun 28 02:04 dm-0
brw-rw---- 1 root disk   253,  1 Jun 28 02:04 dm-1
drwxr-xr-x. 3 root root     80 Jun 28 02:04 dri
crw-rw---- 1 root video  29,  0 Jun 28 02:04 fb0
lrwxrwxrwx. 1 root root    13 Jun 28 02:04 fd -> /proc/self/fd
crw-rw-rw-. 1 root root     1,  7 Jun 28 02:04 full
crw-rw-rw-. 1 root root    10, 229 Jun 28 02:04 fuse
crw----- 1 root root   245,  0 Jun 28 02:04 hidraw0
crw----- 1 root root    10, 228 Jun 28 02:04 hpet
drwxr-xr-x. 3 root root     0 Jun 28 02:04 hugepages
crw----- 1 root root    10, 183 Jun 28 02:04 hwrng
```

```
lrwxrwxrwx. 1 root root      12 Jun 28 02:04 initctl -> /run/initctl
drwxr-xr-x. 4 root root     280 Jun 28 02:04 input
crw-r--r--. 1 root root      1, 11 Jun 28 02:04 kmsg
lrwxrwxrwx. 1 root root      28 Jun 28 02:04 log -> /run/systemd/journal/dev-log
crw-rw----. 1 root disk    10, 237 Jun 28 02:04 loop-control
crw-rw----. 1 root lp       6,  0 Jun 28 02:04 lp0
crw-rw----. 1 root lp       6,  1 Jun 28 02:04 lp1
crw-rw----. 1 root lp       6,  2 Jun 28 02:04 lp2
crw-rw----. 1 root lp       6,  3 Jun 28 02:04 lp3
drwxr-xr-x. 2 root root     100 Jun 28 02:04 mapper
crw-----. 1 root root    10, 227 Jun 28 02:04 mcelog
crw-r-----. 1 root kmem    1,  1 Jun 28 02:04 mem
drwxrwxrwt. 2 root root      40 Jun 28 02:04 mqueue
drwxr-xr-x. 2 root root      60 Jun 28 02:04 net
crw-rw-rw-. 1 root root      1,  3 Jun 28 02:04 null
--More--
```

As you can see, certain files refer to **block** devices whilst others refer to **character** devices:

```
...
brw-rw----. 1 root disk      8,  1 Jun 28 02:04 sda1
...
crw-rw-rw-. 1 root tty       5,  0 Jun 28 02:04 tty
...
```

The major difference between these two types lies in the way that the communication between the system and the peripheral takes place. In the case of a block file, that communication uses a buffer whilst in the case of a character file the communication takes place directly byte by byte.

The figures that can be seen immediately before the date of the special file are called respectively the **major** and the **minor** :

- the **major** identifies the peripheral's driver,
- the **minor** identifies the peripheral. For instance 8,1 indicates the first partition of the **sda** disk.

LAB #1 - Commands

1.1 - The lspci Command

This command show a list of the peripherals connected to the PCI, AGP and PCI express buses:

```
[root@centos8 ~]# lspci
00:00.0 Host bridge: Intel Corporation 440FX - 82441FX PMC [Natoma] (rev 02)
00:01.0 ISA bridge: Intel Corporation 82371SB PIIX3 ISA [Natoma/Triton II]
00:01.1 IDE interface: Intel Corporation 82371SB PIIX3 IDE [Natoma/Triton II]
00:01.2 USB controller: Intel Corporation 82371SB PIIX3 USB [Natoma/Triton II] (rev 01)
00:01.3 Bridge: Intel Corporation 82371AB/EB/MB PIIX4 ACPI (rev 03)
00:02.0 VGA compatible controller: Device 1234:1111 (rev 02)
00:03.0 Unclassified device [00ff]: Red Hat, Inc. Virtio memory balloon
00:07.0 SATA controller: Intel Corporation 82801IR/IO/IH (ICH9R/D0/DH) 6 port SATA Controller [AHCI mode] (rev 02)
00:12.0 Ethernet controller: Red Hat, Inc. Virtio network device
00:1e.0 PCI bridge: Red Hat, Inc. QEMU PCI-PCI bridge
00:1f.0 PCI bridge: Red Hat, Inc. QEMU PCI-PCI bridge
```

To obtain peripheral specific information, use the **-v** or **-vv** switches whilst specifying the peripheral ID:

```
[root@centos8 ~]# lspci -v -s 00:03.0
00:03.0 Unclassified device [00ff]: Red Hat, Inc. Virtio memory balloon
    Subsystem: Red Hat, Inc. Device 0005
    Physical Slot: 3
    Flags: bus master, fast devsel, latency 0, IRQ 10
    I/O ports at e000 [size=64]
    Memory at fe400000 (64-bit, prefetchable) [size=16K]
    Capabilities: [84] Vendor Specific Information: VirtIO: <unknown>
    Capabilities: [70] Vendor Specific Information: VirtIO: Notify
    Capabilities: [60] Vendor Specific Information: VirtIO: DeviceCfg
```

```
Capabilities: [50] Vendor Specific Information: VirtIO: ISR
Capabilities: [40] Vendor Specific Information: VirtIO: CommonCfg
Kernel driver in use: virtio-pci
```

```
[root@centos8 ~]# lspci -vv -s 00:03.0
00:03.0 Unclassified device [00ff]: Red Hat, Inc. Virtio memory balloon
Subsystem: Red Hat, Inc. Device 0005
Physical Slot: 3
Control: I/O+ Mem+ BusMaster+ SpecCycle- MemWINV- VGASnoop- ParErr- Stepping- SERR+ FastB2B- DisINTx-
Status: Cap+ 66MHz- UDF- FastB2B- ParErr- DEVSEL=fast >TAbort- <TAbort- <MAbort- >SERR- <PERR- INTx-
Latency: 0
Interrupt: pin A routed to IRQ 10
Region 0: I/O ports at e000 [size=64]
Region 4: Memory at fe400000 (64-bit, prefetchable) [size=16K]
Capabilities: [84] Vendor Specific Information: VirtIO: <unknown>
        BAR=0 offset=00000000 size=00000000
Capabilities: [70] Vendor Specific Information: VirtIO: Notify
        BAR=4 offset=00003000 size=00001000 multiplier=00000004
Capabilities: [60] Vendor Specific Information: VirtIO: DeviceCfg
        BAR=4 offset=00002000 size=00001000
Capabilities: [50] Vendor Specific Information: VirtIO: ISR
        BAR=4 offset=00001000 size=00001000
Capabilities: [40] Vendor Specific Information: VirtIO: CommonCfg
        BAR=4 offset=00000000 size=00001000
Kernel driver in use: virtio-pci
```

Command Line Switches

The command line switches of this command are:

```
[root@centos8 ~]# lspci --help
lspci: invalid option -- '-'
Usage: lspci [<switches>]
```

Basic display modes:

-mm Produce machine-readable output (single -m for an obsolete format)
-t Show bus tree

Display options:

-v Be verbose (-vv or -vvv for higher verbosity)
-k Show kernel drivers handling each device
-x Show hex-dump of the standard part of the config space
-xxx Show hex-dump of the whole config space (dangerous; root only)
-xxxx Show hex-dump of the 4096-byte extended config space (root only)
-b Bus-centric view (addresses and IRQ's as seen by the bus)
-D Always show domain numbers
-P Display bridge path in addition to bus and device number
-PP Display bus path in addition to bus and device number

Resolving of device ID's to names:

-n Show numeric ID's
-nn Show both textual and numeric ID's (names & numbers)
-q Query the PCI ID database for unknown ID's via DNS
-qq As above, but re-query locally cached entries
-Q Query the PCI ID database for all ID's via DNS

Selection of devices:

-s [[:<domain>]:]<bus>[:<slot>][.<func>] Show only devices in selected slots
-d [<vendor>]:<device>[:<class>] Show only devices with specified ID's

Other options:

-i <file> Use specified ID database instead of /usr/share/hwdata/pci.ids
-p <file> Look up kernel modules in a given file instead of default modules.pcimap
-M Enable `bus

1.2 - The lsusb Command

This command show a list of the peripherals connected to the USB bus:

```
[root@centos8 ~]# lsusb
Bus 001 Device 002: ID 0627:0001 Adomax Technology Co., Ltd
Bus 001 Device 001: ID 1d6b:0001 Linux Foundation 1.1 root hub

[root@centos8 ~]# lsusb -vt
/: Bus 01.Port 1: Dev 1, Class=root_hub, Driver=uhci_hcd/2p, 12M
   |__ Port 1: Dev 2, If 0, Class=Human Interface Device, Driver=usbhid, 12M
```

Command Line Switches

The command line switches of this command are:

```
[root@centos8 ~]# lsusb --help
Usage: lsusb [options]...
List USB devices
-v, --verbose
    Increase verbosity (show descriptors)
-s [[bus:]][devnum]
    Show only devices with specified device and/or
    bus numbers (in decimal)
-d vendor:[product]
    Show only devices with the specified vendor and
    product ID numbers (in hexadecimal)
    .LAB#1
-D device
    Selects which device lsusb will examine
-t, --tree
    Dump the physical USB device hierarchy as a tree
```

```
-V, --version
    Show version of program
-h, --help
    Show usage and help
```

1.3 - The dmidecode Command

The **dmidecode** Command reads the **DMI** (*Desktop Management Interface*) table, also called the **SMBIOS** (*System Management BIOS*) and provides information on:

- the current status of each peripheral,
- possible extensions.

```
[root@centos8 ~]# dmidecode
# dmidecode 3.2
Getting SMBIOS data from sysfs.
SMBIOS 2.8 present.
11 structures occupying 511 bytes.
Table at 0x000F5870.

Handle 0x0000, DMI type 0, 24 bytes
BIOS Information
    Vendor: SeaBIOS
    Version: rel-1.14.0-0-g155821a1990b-prebuilt.qemu.org
    Release Date: 04/01/2014
    Address: 0xE8000
    Runtime Size: 96 kB
    ROM Size: 64 kB
    Characteristics:
        BIOS characteristics not supported
        Targeted content distribution is supported
    BIOS Revision: 0.0
```

Handle 0x0100, DMI type 1, 27 bytes

System Information

Manufacturer: QEMU
Product Name: Standard PC (i440FX + PIIX, 1996)
Version: pc-i440fx-5.2
Serial Number: Not Specified
UUID: 95bd69e3-4a74-44a7-b58c-b74fbfb86df2
Wake-up Type: Power Switch
SKU Number: Not Specified
Family: Not Specified

Handle 0x0300, DMI type 3, 22 bytes

Chassis Information

Manufacturer: QEMU
Type: Other
Lock: Not Present
Version: pc-i440fx-5.2
Serial Number: Not Specified
Asset Tag: Not Specified
Boot-up State: Safe
Power Supply State: Safe
Thermal State: Safe
Security Status: Unknown
OEM Information: 0x00000000
Height: Unspecified
Number Of Power Cords: Unspecified
Contained Elements: 0
SKU Number: Not Specified

Handle 0x0400, DMI type 4, 42 bytes

Processor Information

Socket Designation: CPU 0
Type: Central Processor
Family: Other

Manufacturer: QEMU
ID: 61 0F 00 00 FF FB 8B 07
Version: pc-i440fx-5.2
Voltage: Unknown
External Clock: Unknown
Max Speed: 2000 MHz
Current Speed: 2000 MHz
Status: Populated, Enabled
Upgrade: Other
L1 Cache Handle: Not Provided
L2 Cache Handle: Not Provided
L3 Cache Handle: Not Provided
Serial Number: Not Specified
Asset Tag: Not Specified
Part Number: Not Specified
Core Count: 4
Core Enabled: 4
Thread Count: 1
Characteristics: None

Handle 0x0401, DMI type 4, 42 bytes

Processor Information

Socket Designation: CPU 1
Type: Central Processor
Family: Other
Manufacturer: QEMU
ID: 61 0F 00 00 FF FB 8B 07
Version: pc-i440fx-5.2
Voltage: Unknown
External Clock: Unknown
Max Speed: 2000 MHz
Current Speed: 2000 MHz
Status: Populated, Enabled
Upgrade: Other

L1 Cache Handle: Not Provided
L2 Cache Handle: Not Provided
L3 Cache Handle: Not Provided
Serial Number: Not Specified
Asset Tag: Not Specified
Part Number: Not Specified
Core Count: 4
Core Enabled: 4
Thread Count: 1
Characteristics: None

Handle 0x1000, DMI type 16, 23 bytes

Physical Memory Array

Location: Other
Use: System Memory
Error Correction Type: Multi-bit ECC
Maximum Capacity: 4 GB
Error Information Handle: Not Provided
Number Of Devices: 1

Handle 0x1100, DMI type 17, 40 bytes

Memory Device

Array Handle: 0x1000
Error Information Handle: Not Provided
Total Width: Unknown
Data Width: Unknown
Size: 4 GB
Form Factor: DIMM
Set: None
Locator: DIMM 0
Bank Locator: Not Specified
Type: RAM
Type Detail: Other
Speed: Unknown

Manufacturer: QEMU
Serial Number: Not Specified
Asset Tag: Not Specified
Part Number: Not Specified
Rank: Unknown
Configured Memory Speed: Unknown
Minimum Voltage: Unknown
Maximum Voltage: Unknown
Configured Voltage: Unknown

Handle 0x1300, DMI type 19, 31 bytes
Memory Array Mapped Address
Starting Address: 0x000000000000
Ending Address: 0x000BFFFFFFF
Range Size: 3 GB
Physical Array Handle: 0x1000
Partition Width: 1

Handle 0x1301, DMI type 19, 31 bytes
Memory Array Mapped Address
Starting Address: 0x001000000000
Ending Address: 0x0013FFFFFFF
Range Size: 1 GB
Physical Array Handle: 0x1000
Partition Width: 1

Handle 0x2000, DMI type 32, 11 bytes
System Boot Information
Status: No errors detected

Handle 0x7F00, DMI type 127, 4 bytes
End Of Table

Command Line Switches

The command line switches of this command are:

```
[root@centos7 ~]# dmidecode --help
Usage: dmidecode [OPTIONS]
Options are:
  -d, --dev-mem FILE      Read memory from device FILE (default: /dev/mem)
  -h, --help              Display this help text and exit
  -q, --quiet             Less verbose output
  -s, --string KEYWORD    Only display the value of the given DMI string
  -t, --type TYPE         Only display the entries of given type
  -u, --dump              Do not decode the entries
  --dump-bin FILE        Dump the DMI data to a binary file
  --from-dump FILE       Read the DMI data from a binary file
  -V, --version           Display the version and exit
```

LAB #2 - The sysctl Command

2.1 - The /proc Directory

The /proc directory contains virtual files and directories which are created dynamically when consulted. Only root can consult all of the information in /proc.

```
[root@centos8 ~]# ls /proc
1      16391  19    2212  2427  2622  431   59    84    999    mdstat
10     16476  1931  2215  2428  2659  432   6     842   acpi   meminfo
1007   16534  1956  2222  2431  2667  433   60    8465  buddyinfo  misc
11     16576  1960  2226  2432  2686  434   61    866   bus    modules
11805  16593  2     2230  2435  27    435   63    867   cgroups  mounts
12     16598  20    2237  2439  28    436   64    868   cmdline  mtrr
```

1219	16600	2007	2238	244	29	437	65	869	consoles	net
1228	16613	2029	2241	2443	3	44	6568	87	cpuinfo	pagetypeinfo
1232	16646	203	2244	2445	31	446	66	870	crypto	partitions
1234	16673	2034	2247	2449	32	45	67	871	devices	sched_debug
1235	16677	2037	2260	2451	33	46	674	872	diskstats	schedstat
1247	16711	2054	2262	2465	34	47	68	874	dma	scsi
13	16712	2062	2267	2472	35	4790	69	875	driver	self
1307	16729	21	2268	2473	37	49	70	878	execdomains	slabinfo
1339	16742	210	2274	2474	38	50	701	879	fb	softirqs
1356	17	2118	2275	2475	39	5076	71	880	filesystems	stat
14	1764	2121	2280	2476	4	51	714	884	fs	swaps
1441	180	2124	2287	2478	40	52	72	887	interrupts	sys
1443	181	2126	2292	2481	402	53	73	9	iomem	sysrq-trigger
1444	1817	2156	23	2484	41	532	74	901	ioports	sysvipc
1446	182	2160	2302	25	419	539	75	903	irq	thread-self
14977	1828	2164	2307	2536	420	55	76	9144	kallsyms	timer_list
15	1829	2165	2310	2539	421	568	8	916	kcore	tty
15067	183	2167	2330	2571	422	569	808	918	keys	uptime
1536	1845	2169	2332	2578	423	570	809	919	key-users	version
1553	185	2177	2349	2579	425	571	81	921	kmsg	vmallocinfo
15594	186	2187	2358	259	426	572	833	969	kpagecgroup	vmstat
15735	187	2190	2373	2593	427	573	835	986	kpagecount	zoneinfo
16	1880	2194	2384	26	428	574	837	989	kpageflags	
16165	1883	22	239	2602	43	575	838	990	loadavg	
16167	1888	2204	241	2608	430	576	839	993	locks	

Files

/proc/cpuinfo

```
[root@centos8 ~]# cat /proc/cpuinfo
processor      : 0
```

```
vendor_id      : GenuineIntel
cpu family    : 15
model         : 6
model name    : Common KVM processor
stepping     : 1
microcode    : 0x1
cpu MHz      : 1999.987
cache size   : 16384 KB
physical id  : 0
siblings     : 4
core id      : 0
cpu cores    : 4
apicid       : 0
initial apicid : 0
fpu          : yes
fpu_exception : yes
cpuid level  : 13
wp           : yes
flags        : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse
sse2 ht syscall nx lm constant_tsc nopl xtopology cpuid tsc_known_freq pni cx16 x2apic hypervisor lahf_lm
cpuid_fault pti
bugs         : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit
bogomips     : 3999.97
clflush size : 64
cache_alignment : 128
address sizes : 40 bits physical, 48 bits virtual
power management:

processor     : 1
vendor_id    : GenuineIntel
cpu family   : 15
model       : 6
model name  : Common KVM processor
stepping    : 1
```

```
microcode      : 0x1
cpu MHz        : 1999.987
cache size     : 16384 KB
physical id    : 0
siblings       : 4
core id        : 1
cpu cores      : 4
apicid         : 1
initial apicid : 1
fpu            : yes
fpu_exception  : yes
cpuid level    : 13
wp             : yes
flags          : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse
sse2 ht syscall nx lm constant_tsc nopl xtopology cpuid tsc_known_freq pni cx16 x2apic hypervisor lahf_lm
cpuid_fault pti
bugs           : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit
bogomips       : 3999.97
clflush size   : 64
cache_alignment : 128
address sizes  : 40 bits physical, 48 bits virtual
power management:

processor      : 2
vendor_id      : GenuineIntel
cpu family     : 15
model          : 6
model name     : Common KVM processor
stepping       : 1
microcode      : 0x1
cpu MHz        : 1999.987
cache size     : 16384 KB
physical id    : 0
siblings       : 4
```

```
core id      : 2
cpu cores   : 4
apicid      : 2
initial apicid : 2
fpu         : yes
fpu_exception : yes
cpuid level : 13
wp          : yes
flags       : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse
sse2 ht syscall nx lm constant_tsc nopl xtopology cpuid tsc_known_freq pni cx16 x2apic hypervisor lahf_lm
cpuid_fault pti
bugs        : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit
bogomips    : 3999.97
clflush size : 64
cache_alignment : 128
address sizes : 40 bits physical, 48 bits virtual
power management:

processor    : 3
vendor_id   : GenuineIntel
cpu family  : 15
model       : 6
model name  : Common KVM processor
stepping    : 1
microcode   : 0x1
cpu MHz     : 1999.987
cache size  : 16384 KB
physical id : 0
siblings    : 4
core id     : 3
cpu cores   : 4
apicid      : 3
initial apicid : 3
fpu         : yes
```

```
fpu_exception : yes
cpuid level   : 13
wp           : yes
flags        : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse
sse2 ht syscall nx lm constant_tsc nopl xtopology cpuid tsc_known_freq pni cx16 x2apic hypervisor lahf_lm
cpuid_fault  pti
bugs         : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit
bogomips     : 3999.97
clflush size : 64
cache_alignment : 128
address sizes : 40 bits physical, 48 bits virtual
power management:

processor    : 4
vendor_id    : GenuineIntel
cpu family   : 15
model        : 6
model name   : Common KVM processor
stepping     : 1
microcode    : 0x1
cpu MHz      : 1999.987
cache size   : 16384 KB
physical id   : 1
siblings     : 4
core id      : 0
cpu cores    : 4
apicid       : 4
initial apicid : 4
fpu          : yes
fpu_exception : yes
cpuid level   : 13
wp           : yes
flags        : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse
sse2 ht syscall nx lm constant_tsc nopl xtopology cpuid tsc_known_freq pni cx16 x2apic hypervisor lahf_lm
```

```
cpuid_fault pti
bugs          : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit
bogomips      : 3999.97
clflush size  : 64
cache_alignment : 128
address sizes : 40 bits physical, 48 bits virtual
power management:

processor     : 5
vendor_id    : GenuineIntel
cpu family   : 15
model        : 6
model name   : Common KVM processor
stepping     : 1
microcode    : 0x1
cpu MHz      : 1999.987
cache size   : 16384 KB
physical id  : 1
siblings     : 4
core id      : 1
cpu cores    : 4
apicid       : 5
initial apicid : 5
fpu          : yes
fpu_exception : yes
cpuid level  : 13
wp           : yes
flags        : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse
sse2 ht syscall nx lm constant_tsc nopl xtopology cpuid tsc_known_freq pni cx16 x2apic hypervisor lahf_lm
cpuid_fault pti
bugs          : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit
bogomips      : 3999.97
clflush size  : 64
cache_alignment : 128
```

```
address sizes : 40 bits physical, 48 bits virtual
power management:
```

```
processor      : 6
vendor_id     : GenuineIntel
cpu family    : 15
model         : 6
model name    : Common KVM processor
stepping      : 1
microcode     : 0x1
cpu MHz       : 1999.987
cache size    : 16384 KB
physical id   : 1
siblings      : 4
core id       : 2
cpu cores     : 4
apicid        : 6
initial apicid : 6
fpu           : yes
fpu_exception : yes
cpuid level   : 13
wp            : yes
flags         : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse
sse2 ht syscall nx lm constant_tsc nopl xtopology cpuid tsc_known_freq pni cx16 x2apic hypervisor lahf_lm
cpuid_fault pti
bugs          : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit
bogomips      : 3999.97
clflush size  : 64
cache_alignment : 128
address sizes : 40 bits physical, 48 bits virtual
power management:

processor      : 7
vendor_id     : GenuineIntel
```


RTR:	0	0	0	0	0	0	0	0	0	APIC ICR read retries
RES:	178171	67749	43116	65994	71361	118585	54139	53452		Rescheduling interrupts
CAL:	11095	11007	10841	10343	14679	9998	12165	12443		Function call interrupts
TLB:	2295	2297	2000	1728	2330	2338	1991	1861		TLB shutdowns
TRM:	0	0	0	0	0	0	0	0		Thermal event interrupts
THR:	0	0	0	0	0	0	0	0		Threshold APIC interrupts
DFR:	0	0	0	0	0	0	0	0		Deferred Error APIC interrupts
MCE:	0	0	0	0	0	0	0	0		Machine check exceptions
MCP:	288	288	288	288	288	288	288	288		Machine check polls
HYP:	0	0	0	0	0	0	0	0		Hypervisor callback interrupts
HRE:	0	0	0	0	0	0	0	0		Hyper-V reenlightenment interrupts
HVS:	0	0	0	0	0	0	0	0		Hyper-V stimer0 interrupts
ERR:	0									
MIS:	0									
PIN:	0	0	0	0	0	0	0	0		Posted-interrupt notification event
NPI:	0	0	0	0	0	0	0	0		Nested posted- interrupt event
PIW:	0	0	0	0	0	0	0	0		Posted-interrupt wakeup event



Important : The use of an IRQ by a peripheral is exclusive.

/proc/dma

```
[root@centos8 ~]# cat /proc/dma
4: cascade
```

/proc/ioports

```
root@centos8 ~]# cat /proc/ioports | more
0000-0cf7 : PCI Bus 0000:00
  0000-001f : dma1
  0020-0021 : pic1
  0040-0043 : timer0
  0050-0053 : timer1
  0060-0060 : keyboard
  0064-0064 : keyboard
  0070-0077 : rtc0
  0080-008f : dma page reg
  00a0-00a1 : pic2
  00c0-00df : dma2
  00f0-00ff : fpu
  0170-0177 : 0000:00:01.1
    0170-0177 : ata_piix
  01f0-01f7 : 0000:00:01.1
    01f0-01f7 : ata_piix
  0376-0376 : 0000:00:01.1
    0376-0376 : ata_piix
  03c0-03df : vga+
  03f6-03f6 : 0000:00:01.1
    03f6-03f6 : ata_piix
--More--
```



Important - If two peripherals use the same IO Port, both become unusable.

/proc/devices

```
[root@centos8 ~]# cat /proc/devices
Character devices:
 1 mem
 4 /dev/vc/0
 4 tty
 4 ttyS
 5 /dev/tty
 5 /dev/console
 5 /dev/ptmx
 7 vcs
10 misc
13 input
21 sg
29 fb
128 ptm
136 pts
162 raw
180 usb
188 ttyUSB
189 usb_device
202 cpu/msr
203 cpu/cpuid
226 drm
244 aux
245 hidraw
246 usbmon
```

```
247 bsg
248 watchdog
249 ptp
250 pps
251 rtc
252 dax
253 tpm
254 gpiochip
```

Block devices:

```
 8 sd
 9 md
11 sr
65 sd
66 sd
67 sd
68 sd
69 sd
70 sd
71 sd
128 sd
129 sd
130 sd
131 sd
132 sd
133 sd
134 sd
135 sd
253 device-mapper
254 mdp
259 blkext
```

/proc/modules

```
[root@centos8 ~]# cat /proc/modules | more
xt_CHECKSUM 16384 1 - Live 0xffffffffc09a8000
ipt_MASQUERADE 16384 3 - Live 0xffffffffc09a3000
xt_conntrack 16384 1 - Live 0xffffffffc099e000
ipt_REJECT 16384 2 - Live 0xffffffffc0999000
nft_compat 20480 16 - Live 0xffffffffc0993000
nf_nat_tftp 16384 0 - Live 0xffffffffc098b000
nft_objref 16384 1 - Live 0xffffffffc0986000
nf_conntrack_tftp 16384 3 nf_nat_tftp, Live 0xffffffffc0981000
nft_counter 16384 33 - Live 0xffffffffc097c000
tun 53248 1 - Live 0xffffffffc096e000
bridge 192512 0 - Live 0xffffffffc093e000
stp 16384 1 bridge, Live 0xffffffffc0939000
llc 16384 2 bridge,stp, Live 0xffffffffc0930000
nft_fib_inet 16384 1 - Live 0xffffffffc08f5000
nft_fib_ipv4 16384 1 nft_fib_inet, Live 0xffffffffc08ed000
nft_fib_ipv6 16384 1 nft_fib_inet, Live 0xffffffffc08e8000
nft_fib 16384 3 nft_fib_inet,nft_fib_ipv4,nft_fib_ipv6, Live 0xffffffffc08e3000
nft_reject_inet 16384 5 - Live 0xffffffffc08de000
nf_reject_ipv4 16384 2 ipt_REJECT,nft_reject_inet, Live 0xffffffffc08d9000
nf_reject_ipv6 16384 1 nft_reject_inet, Live 0xffffffffc08d4000
nft_reject 16384 1 nft_reject_inet, Live 0xffffffffc08cf000
--More--
```

/proc/diskstats

```
[root@centos8 ~]# cat /proc/diskstats
 8      0 sda 15481 112 1445637 154103 10272 2377 277530 890611 0 237219 1044714 0 0 0 0
 8      1 sda1 402 3 66754 13349 31 18 392 4632 0 2824 17981 0 0 0 0
 8      2 sda2 14915 109 1375516 140528 8450 2359 277138 869788 0 225416 1010316 0 0 0 0
```

```
 8      16 sdb 230 0 5991 36 0 0 0 0 0 110 36 0 0 0 0
11      0 sr0 10 0 4 2 0 0 0 0 0 9 2 0 0 0 0
253     0 dm-0 11651 0 1364532 72138 12121 0 288727 1208138 0 227630 1280276 0 0 0 0
253     1 dm-1 104 0 4440 79 0 0 0 0 0 71 79 0 0 0 0
```

/proc/partitions

```
[root@centos8 ~]# cat /proc/partitions
major minor #blocks name

 8      0 33554432 sda
 8      1  1048576 sda1
 8      2 32504832 sda2
 8     16  4194304 sdb
11      0  1048575 sr0
253     0 29143040 dm-0
253     1  3358720 dm-1
```

/proc/swaps

```
[root@centos8 ~]# cat /proc/swaps
Filename                                Type      Size      Used      Priority
/dev/dm-1                              partition 3358716 0         -2
```

/proc/loadavg

```
[root@centos8 ~]# cat /proc/loadavg
0.00 0.00 0.00 1/697 16936
```

/proc/meminfo

```
[root@centos8 ~]# cat /proc/meminfo
MemTotal:      3825032 kB
MemFree:       1862116 kB
MemAvailable:  2420560 kB
Buffers:       3300 kB
Cached:        750496 kB
SwapCached:    0 kB
Active:        315572 kB
Inactive:      1400260 kB
Active(anon):  1856 kB
Inactive(anon): 974728 kB
Active(file):  313716 kB
Inactive(file): 425532 kB
Unevictable:   0 kB
Mlocked:       0 kB
SwapTotal:     3358716 kB
SwapFree:      3358716 kB
Dirty:         0 kB
Writeback:     0 kB
AnonPages:     962004 kB
Mapped:        261084 kB
Shmem:         14552 kB
KReclaimable:  46980 kB
Slab:          118396 kB
SReclaimable:  46980 kB
SUnreclaim:   71416 kB
KernelStack:  11280 kB
PageTables:    46532 kB
NFS_Unstable:  0 kB
Bounce:        0 kB
WritebackTmp:  0 kB
```

```
CommitLimit:      5271232 kB
Committed_AS:     5072744 kB
VmallocTotal:     34359738367 kB
VmallocUsed:      0 kB
VmallocChunk:     0 kB
Percpu:           5920 kB
HardwareCorrupted: 0 kB
AnonHugePages:    546816 kB
ShmemHugePages:   0 kB
ShmemPmdMapped:   0 kB
FileHugePages:    0 kB
FilePmdMapped:    0 kB
HugePages_Total:  0
HugePages_Free:   0
HugePages_Rsvd:   0
HugePages_Surp:   0
Hugepagesize:     2048 kB
Hugetlb:          0 kB
DirectMap4k:      173944 kB
DirectMap2M:      4020224 kB
```

/proc/version

```
[root@centos8 ~]# cat /proc/version
Linux version 4.18.0-305.3.1.el8.x86_64 (mockbuild@kbuilder.bsys.centos.org) (gcc version 8.4.1 20200928 (Red Hat 8.4.1-1) (GCC)) #1 SMP Tue Jun 1 16:14:33 UTC 2021
```

Répertoires

ide/scsi

This sub-directory contains disk capacity, disk type and disk geometry information.

acpi

This sub-directory contains information on energy management, temperatures, fan speeds and battery levels.

bus

This sub-directory contains a sub-directory for each bus.

net

This sub-directory contains information concerning the network.

sys

This sub-directory contains files used by root to configure the kernel. For instance, the following command configures the kernel to ignore incoming pings:

```
# echo 1 > /proc/sys/net/ipv4/icmp_echo_ignore_all [Enter]
```

2.2 - Using the sysctl Command

Files in the **/proc/sys** can be administered by using the **sysctl** command.

The **sysctl** command applies rules at system boot that are defined in the **/etc/sysctl.conf** file:

```
[root@centos8 ~]# cat /etc/sysctl.conf
# sysctl settings are defined through files in
# /usr/lib/sysctl.d/, /run/sysctl.d/, and /etc/sysctl.d/.
#
# Vendors settings live in /usr/lib/sysctl.d/.
# To override a whole file, create a new file with the same in
# /etc/sysctl.d/ and put new settings there. To override
# only specific settings, add a file with a lexically later
# name in /etc/sysctl.d/ and put new settings there.
#
# For more information, see sysctl.conf(5) and sysctl.d(5).

[root@centos8 ~]# ls -l /etc/sysctl.d/
total 0
lrwxrwxrwx. 1 root root 14 Mar 16 15:42 99-sysctl.conf -> ../sysctl.conf
[root@centos8 ~]# ls -l /usr/lib/sysctl.d/
total 24
-rw-r--r--. 1 root root 1810 Dec 22 2020 10-default-yama-scope.conf
-rw-r--r--. 1 root root 524 Mar 16 15:42 50-coredump.conf
-rw-r--r--. 1 root root 1270 Mar 16 15:42 50-default.conf
-rw-r--r--. 1 root root 246 Jun 15 2020 50-libkcapioptmem_max.conf
-rw-r--r--. 1 root root 636 Mar 16 15:42 50-pid-max.conf
-rw-r--r--. 1 root root 499 Nov 26 2019 60-libvirtd.conf

[root@centos8 ~]# cat /usr/lib/sysctl.d/50-default.conf
# This file is part of systemd.
#
# systemd is free software; you can redistribute it and/or modify it
# under the terms of the GNU Lesser General Public License as published by
# the Free Software Foundation; either version 2.1 of the License, or
# (at your option) any later version.

# See sysctl.d(5) and core(5) for documentation.
```

```
# To override settings in this file, create a local file in /etc
# (e.g. /etc/sysctl.d/90-override.conf), and put any assignments
# there.

# System Request functionality of the kernel (SYNC)
#
# Use kernel.sysrq = 1 to allow all keys.
# See https://www.kernel.org/doc/html/latest/admin-guide/sysrq.html for a list
# of values and keys.
kernel.sysrq = 16

# Append the PID to the core filename
kernel.core_uses_pid = 1

# https://bugzilla.redhat.com/show\_bug.cgi?id=1689346
kernel.kptr_restrict = 1

# Source route verification
net.ipv4.conf.all.rp_filter = 1

# Do not accept source routing
net.ipv4.conf.all.accept_source_route = 0

# Promote secondary addresses when the primary address is removed
net.ipv4.conf.all.promote_secondaries = 1

# Fair Queue CoDel packet scheduler to fight bufferbloat
net.core.default_qdisc = fq_codel

# Enable hard and soft link protection
fs.protected_hardlinks = 1
fs.protected_symlinks = 1
```

The command line switches of this command are:

Usage:

```
sysctl [options] [variable[=value] ...]
```

Options:

```
-a, --all          display all variables
-A                alias of -a
-X                alias of -a
  --deprecated    include deprecated parameters to listing
-b, --binary      print value without new line
-e, --ignore      ignore unknown variables errors
-N, --names       print variable names without values
-n, --values      print only values of the given variable(s)
-p, --load[=<file>] read values from file
-f                alias of -p
  --system        read values from all system directories
-r, --pattern <expression>
                  select setting that match expression
-q, --quiet       do not echo variable set
-w, --write       enable writing a value to variable
-o               does nothing
-x               does nothing
-d               alias of -h

-h, --help       display this help and exit
-V, --version    output version information and exit
```

For more details see `sysctl(8)`.

LAB#3 - Interpreting Information in /proc

The information found in files in the /proc filesystem can be interpreted using the following commands:

- free,
- uptime et w,
- iostat,
- hdparm
- vmstat,
- mpstat,
- sar.

3.1 - The free Command

The **free** command shows total, used, free, shared, buffered, cached and swapped memory:

```
[root@centos8 ~]# free -m
```

	total	used	free	shared	buff/cache	available
Mem:	3735	1135	1818	14	782	2363
Swap:	3279	0	3279			

In the above example, you can see:

- 3735 MB of total physical memory,
- 1135 MB of used physical memory and 1818 MB of free physical memory,
- 3279 MB of swap space with 0MB being used.

The command line switches of this command are:

```
[root@centos8 ~]# free --help
```

Usage:
free [options]

Options:
-b, --bytes show output in bytes
--kilo show output in kilobytes

```
--mega      show output in megabytes
--giga      show output in gigabytes
--tera      show output in terabytes
--peta      show output in petabytes
-k, --kibi  show output in kibibytes
-m, --mebi  show output in mebibytes
-g, --gibi  show output in gibibytes
--tebi      show output in tebibytes
--pebi      show output in pebibytes
-h, --human show human-readable output
--si        use powers of 1000 not 1024
-l, --lohi  show detailed low and high memory statistics
-t, --total show total for RAM + swap
-s N, --seconds N repeat printing every N seconds
-c N, --count N repeat printing N times, then exit
-w, --wide  wide output

--help      display this help and exit
-V, --version output version information and exit
```

For more details see `free(1)`.

3.2 - The uptime and w Commands

Each of these commands show the load average over the past 1, 5 and 15 minutes:

```
[root@centos8 ~]# uptime
04:39:03 up 1 day,  2:34,  1 user,  load average: 0.00, 0.00, 0.00

[root@centos8 ~]# w
04:39:04 up 1 day,  2:34,  1 user,  load average: 0.00, 0.00, 0.00
USER      TTY      FROM          LOGIN@   IDLE   JCPU   PCPU   WHAT
```

```
trainee pts/0 10.0.2.1 03:58 0.00s 0.11s 0.02s sshd: trainee [priv]
```

The **load average** indicates the number of processes being executed and waiting to be executed for the period concerned.

If the load average of a single-core system was **3.48 4.00 3.85** this would indicate a bottleneck since, on average:

- 2.48 processes would have been waiting to be executed over the last minute,
- 3.00 processes would have been waiting to be executed over the last 5 minutes,
- 2.85 processes would have been waiting to be executed over the last 15 minutes,

The command line switches of these commands are:

```
[root@centos8 ~]# uptime --help
```

Usage:

```
uptime [options]
```

Options:

```
-p, --pretty    show uptime in pretty format  
-h, --help      display this help and exit  
-s, --since     system up since  
-V, --version   output version information and exit
```

For more details see `uptime(1)`.

```
[root@centos8 ~]# w --help
```

Usage:

```
w [options]
```

Options:

```
-h, --no-header    do not print header  
-u, --no-current   ignore current process username  
-s, --short        short format
```

```
-f, --from      show remote hostname field
-o, --old-style old style output
-i, --ip-addr   display IP address instead of hostname (if possible)

--help        display this help and exit
-V, --version   output version information and exit
```

For more details see `w(1)`.

3.3 - The `iostat` Command

The **`iostat`** command show disk, terminal and streamer statistics:

```
[root@centos8 ~]# iostat
bash: iostat: command not found...
Install package 'sysstat' to provide command 'iostat'? [N/y] y

* Waiting in queue...
The following packages have to be installed:
  lm_sensors-libs-3.4.0-22.20180522git70f7e08.el8.x86_64 Lm_sensors core libraries
  sysstat-11.7.3-5.el8.x86_64      Collection of performance monitoring tools for Linux
Proceed with changes? [N/y] y

* Waiting in queue...
* Waiting for authentication...
* Waiting in queue...
* Downloading packages...
* Requesting data...
* Testing changes...
* Installing packages...
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_(8 CPU)
```

```
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.03    0.00    0.03    0.01    0.00   99.93
```

```
Device            tps    kB_read/s    kB_wrtn/s    kB_read    kB_wrtn
sda                0.28         7.67         1.49       735338     142510
sdb                0.00         0.03         0.00        2995         0
scd0               0.00         0.00         0.00         2         0
dm-0               0.26         7.25         1.55       694786     148837
dm-1               0.00         0.02         0.00        2220         0
```

```
[root@centos8 ~]# iostat -d -x
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)          30/06/21      _x86_64_      (8 CPU)

Device            r/s      w/s      rkB/s      wkB/s      rrqm/s      wrqm/s      %rrqm      %wrqm      r_await      w_await      aqu-sz      rareq-sz
wareq-sz  svctm  %util
sda                0.20     0.16     11.67       1.81       0.00       0.03       0.48     17.45       6.79       69.99       0.01     58.28
11.33   9.13   0.33
sdb                0.00     0.00       0.02       0.00       0.00       0.00       0.00       0.00       0.13       0.00       0.00       8.81
0.00   0.51   0.00
scd0               0.00     0.00       0.00       0.00       0.00       0.00       0.00       0.00       0.20       0.00       0.00       0.20
0.00   0.90   0.00
dm-0               0.12     0.19       3.98       1.90       0.00       0.00       0.00       0.00       4.18       79.07       0.02     32.88
10.07  10.34  0.32
dm-1               0.00     0.00       0.01       0.00       0.00       0.00       0.00       0.00       0.76       0.00       0.00     21.35
0.00   0.68   0.00
```

The command line switches of this command are:

```
[root@centos8 ~]# iostat --help
Usage: iostat [ options ] [ <interval> [ <count> ] ]
Options are:
[ -c ] [ -d ] [ -h ] [ -k | -m ] [ -N ] [ -s ] [ -t ] [ -V ] [ -x ] [ -y ] [ -z ]
[ -j { ID | LABEL | PATH | UUID | ... } ] [ --human ] [ -o JSON ]
[ [ -H ] -g <group_name> ] [ -p [ <device> [,...] | ALL ] ]
```

```
[ <device> [...] | ALL ]
```

3.4 - The hdparm Command

The hdparm command measures disk reads:

```
[root@centos8 ~]# hdparm -t /dev/sda

/dev/sda:
Timing buffered disk reads: 1410 MB in 3.00 seconds = 469.98 MB/sec
```

3.5 - La Commande vmstat

The **vmstat** commands shows memory, pagination and processor statistics:

```
[root@centos8 ~]# vmstat 1 10
procs -----memory----- ---swap-- -----io----- -system-- -----cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo    in   cs  us  sy  id  wa  st
 0  0     0 1765216  2256 866336    0    0     3     0    9   12  0  0 100  0  0
 0  0     0 1765136  2256 866336    0    0     0     0   57   80  0  0 100  0  0
 0  0     0 1765136  2256 866376    0    0     0     0   54   77  0  0 100  0  0
 0  0     0 1765136  2256 866376    0    0     0     0   66  100  0  0 100  0  0
 0  0     0 1765136  2256 866376    0    0     0     0  103  125  0  0 100  0  0
 0  0     0 1765108  2256 866376    0    0     0     0   64   86  0  0 100  0  0
 0  0     0 1765108  2256 866376    0    0     0     0   62   88  0  0 100  0  0
 0  0     0 1765108  2256 866376    0    0     0     0   68   97  0  0 100  0  0
 0  0     0 1765108  2256 866376    0    0     0     0   60   88  0  0 100  0  0
 0  0     0 1765108  2256 866376    0    0     0     0  177  251  0  0 100  0  0
```

The command line switches of this command are:

```
[root@centos8 ~]# vmstat --help
```

Usage:

```
vmstat [options] [delay [count]]
```

Options:

```
-a, --active          active/inactive memory
-f, --forks           number of forks since boot
-m, --slabs           slabinfo
-n, --one-header      do not redisplay header
-s, --stats           event counter statistics
-d, --disk            disk statistics
-D, --disk-sum        summarize disk statistics
-p, --partition <dev> partition specific statistics
-S, --unit <char>    define display unit
-w, --wide            wide output
-t, --timestamp       show timestamp

-h, --help           display this help and exit
-V, --version        output version information and exit
```

For more details see `vmstat(8)`.



Important : By default `vmstat` shows statistics from boot until current time.

3.6 - The `mpstat` Command

La commande **`mpstat`** affiche des statistiques détaillées sur le CPU :


```

04:55:19      6  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:19      7  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00

04:55:19      CPU  %usr  %nice  %sys %iowait  %irq  %soft  %steal  %guest  %gnice  %idle
04:55:21     all  0.06  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  99.94
04:55:21      0  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:21      1  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:21      2  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:21      3  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:21      4  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:21      5  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:21      6  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
04:55:21      7  0.50  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  99.50

Average:      CPU  %usr  %nice  %sys %iowait  %irq  %soft  %steal  %guest  %gnice  %idle
Average:     all  0.01  0.00  0.00  0.00  0.02  0.00  0.01  0.00  0.00  0.00  99.95
Average:      0  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
Average:      1  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
Average:      2  0.00  0.00  0.00  0.00  0.00  0.10  0.00  0.00  0.00  0.00  99.90
Average:      3  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
Average:      4  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.10  0.00  0.00  99.90
Average:      5  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
Average:      6  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00 100.00
Average:      7  0.10  0.00  0.00  0.00  0.00  0.10  0.00  0.00  0.00  0.00  99.80

```

The command line switches of this command are:

```

[root@centos8 ~]# mpstat --help
Usage: mpstat [ options ] [ <interval> [ <count> ] ]
Options are:
[ -A ] [ -n ] [ -u ] [ -V ] [ -I { SUM | CPU | SCPU | ALL } ]
[ -N { <node_list> | ALL } ] [ -o JSON ] [ -P { <cpu_list> | ALL } ]

```

3.7 - The sar Command

The **sar** command can survey all system resources dependant upon the switch that is used. Several important switches are:

Switch	Description
-u	CPU usage in %
-q	Number of processes in the process queue
-r	Memory usage
-w	Swap usage
-p	Pagination usage
-b	Buffer usage
-d	Disk usage

The **/usr/lib64/sa/sadc** command is used to collect data:

```
[root@centos8 ~]# ls /usr/lib64/sa
sa1 sa2 sadc
```

The **/usr/lib64/sa/sa1** script calls the **/usr/lib/sa/sadc** command and can use two switches:

Switch	Description
-t	Interval
-n	Count

The **/usr/lib64/sa/sa2** script also creates a log at **/var/log/sa/sar<dd>**, where <dd> is the day of the month.

```
[root@centos8 ~]# ls /var/log/sa/
sa29 s
```

```
ar29
```

Using CentOS / RHEL 8, the interval between collects is configured using systemd **timers** de systemd and not cron as was previously the case:

```
[root@centos8 ~]# cat /usr/lib/systemd/system/sysstat-collect.timer
# /usr/lib/systemd/system/sysstat-collect.timer
# (C) 2014 Tomasz Torcz <tomek@pipebreaker.pl>
#
# sysstat-11.7.3 systemd unit file:
#     Activates activity collector every 10 minutes

[Unit]
Description=Run system activity accounting tool every 10 minutes

[Timer]
OnCalendar=*:00/10

[Install]
WantedBy=sysstat.service
```

The **OnCalendar** value indicates a collect every 10 minutes.

To change this value, you need to create an **override** file in **/etc/systemd/system/** by using the **systemctl edit** command. You should never edit directly files in **/usr/lib/systemd/system** :

```
[root@centos8 ~]# systemctl edit sysstat-collect.timer
[root@centos8 ~]# cat /etc/systemd/system/sysstat-collect.timer.d/override.conf
[Unit]
Description=Run system activity accounting tool every 2 minutes

[Timer]
OnCalendar=
OnCalendar=*:00/2
AccuracySec=0
```



Important : Note the line **OnCalendar=** which is required to override the default value.

Now check if the configuration has been applied:

```
[root@centos8 ~]# systemctl status sysstat-collect.timer
● sysstat-collect.timer - Run system activity accounting tool every 2 minutes
  Loaded: loaded (/usr/lib/systemd/system/sysstat-collect.timer; enabled; vendor preset: disabled)
  Drop-In: /etc/systemd/system/sysstat-collect.timer.d
           └─override.conf
  Active: active (waiting) since Tue 2021-06-29 06:16:04 EDT; 3h 2min ago
  Trigger: Tue 2021-06-29 09:20:00 EDT; 1min 19s left

Jun 29 06:16:04 centos8.ittraining.loc systemd[1]: Started Run system activity accounting tool every 10 minutes.
```

```
[root@centos8 ~]# journalctl -g sysstat-collect.service
-- Logs begin at Mon 2021-06-28 02:04:10 EDT, end at Tue 2021-06-29 09:18:00 EDT. --
Jun 29 06:20:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 06:26:29 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 06:30:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 06:40:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 06:50:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:00:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:10:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:20:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:30:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:40:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:50:33 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:53:56 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:54:00 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:56:00 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
Jun 29 07:58:00 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
```

```
Jun 29 08:00:00 centos8.ittraining.loc systemd[1]: sysstat-collect.service: Succeeded.
...
```

Execute the sar command:

```
[root@centos8 ~]# sar
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_      (8 CPU)

06:16:04      LINUX RESTART      (8 CPU)

06:20:33      CPU      %user      %nice      %system      %iowait      %steal      %idle
06:26:29      all      0.03      0.00      0.03      0.00      0.00      99.94
06:30:33      all      0.02      0.00      0.03      0.00      0.00      99.95
06:40:33      all      0.02      0.00      0.03      0.00      0.00      99.94
06:50:33      all      0.02      0.00      0.02      0.00      0.00      99.95
07:00:33      all      0.02      0.00      0.02      0.00      0.00      99.95
07:10:33      all      0.02      0.00      0.02      0.00      0.00      99.95
07:20:33      all      0.02      0.00      0.03      0.00      0.00      99.95
07:30:33      all      0.02      0.01      0.03      0.00      0.00      99.94
07:40:33      all      0.03      0.00      0.04      0.00      0.00      99.93
07:50:33      all      0.03      0.00      0.03      0.00      0.00      99.94
07:53:56      all      0.08      0.00      0.06      0.00      0.00      99.86
07:54:00      all      0.09      0.00      0.06      0.00      0.00      99.85
Average:      all      0.03      0.00      0.03      0.00      0.00      99.94

07:55:44      LINUX RESTART      (8 CPU)

07:56:00      CPU      %user      %nice      %system      %iowait      %steal      %idle
07:58:00      all      0.03      0.00      0.03      0.00      0.00      99.94
08:00:00      all      0.02      0.00      0.04      0.00      0.00      99.94
08:02:00      all      0.02      0.00      0.03      0.00      0.00      99.94
08:04:00      all      0.02      0.00      0.03      0.00      0.00      99.95
08:06:00      all      0.02      0.00      0.03      0.00      0.00      99.95
08:08:00      all      0.02      0.00      0.04      0.00      0.00      99.94
```

08:10:00	all	0.02	0.00	0.03	0.00	0.00	99.95
08:12:00	all	0.03	0.00	0.03	0.00	0.00	99.95
08:14:00	all	0.02	0.00	0.03	0.01	0.00	99.94
08:16:00	all	0.02	0.00	0.03	0.00	0.00	99.95
08:18:00	all	0.02	0.00	0.03	0.00	0.00	99.95
08:20:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:22:00	all	0.02	0.00	0.03	0.00	0.00	99.95
08:24:00	all	0.02	0.00	0.02	0.00	0.00	99.95
08:26:00	all	0.02	0.00	0.03	0.00	0.00	99.95
08:28:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:30:00	all	0.02	0.05	0.05	0.00	0.00	99.87
08:32:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:34:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:36:00	all	0.03	0.00	0.04	0.00	0.00	99.94
08:38:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:40:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:42:00	all	0.02	0.00	0.03	0.00	0.00	99.94
08:44:00	all	0.03	0.00	0.03	0.00	0.00	99.94
08:46:00	all	0.02	0.00	0.03	0.00	0.00	99.94
08:48:00	all	0.03	0.00	0.03	0.00	0.00	99.95
08:50:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:52:00	all	0.02	0.00	0.06	0.00	0.00	99.92
08:54:00	all	0.02	0.00	0.03	0.00	0.00	99.95
08:56:00	all	0.02	0.00	0.04	0.00	0.00	99.94
08:58:00	all	0.02	0.00	0.02	0.00	0.00	99.96
09:00:00	all	0.07	0.00	0.05	0.00	0.00	99.88
09:02:00	all	0.02	0.00	0.04	0.00	0.00	99.94
09:04:00	all	0.02	0.00	0.03	0.00	0.00	99.95
09:06:00	all	0.02	0.00	0.04	0.00	0.00	99.94
09:08:00	all	0.02	0.00	0.04	0.00	0.00	99.94
09:10:00	all	0.02	0.00	0.03	0.00	0.00	99.95
09:12:00	all	0.02	0.00	0.03	0.00	0.00	99.94
09:12:00	CPU	%user	%nice	%system	%iowait	%steal	%idle

09:14:00	all	0.02	0.00	0.03	0.00	0.00	99.95
09:16:00	all	0.02	0.00	0.06	0.00	0.00	99.92
09:18:00	all	0.03	0.00	0.03	0.00	0.00	99.95
09:20:00	all	0.02	0.00	0.03	0.00	0.00	99.94
Average:	all	0.02	0.00	0.03	0.00	0.00	99.94

CPU Stats

Use the -u switch:

```
[root@centos8 ~]# sar -u 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_      (8 CPU)

09:22:52      CPU      %user      %nice      %system      %iowait      %steal      %idle
09:22:57      all      0.03      0.00      0.03      0.00      0.00      99.95
09:23:02      all      0.03      0.00      0.03      0.00      0.00      99.95
09:23:07      all      0.00      0.00      0.03      0.00      0.00      99.97
Average:      all      0.02      0.00      0.03      0.00      0.00      99.96
```

More information can be shown by using the **ALL** argument:

```
[root@centos8 ~]# sar -u ALL 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_(8 CPU)

01:49:14      CPU      %usr      %nice      %sys      %iowait      %steal      %irq      %soft      %guest      %gnice
%idle
01:49:19      all      0.03      0.00      0.00      0.00      0.00      0.03      0.00      0.00      0.00
99.95
01:49:24      all      0.03      0.00      0.03      0.00      0.00      0.10      0.05      0.00      0.00
99.80
01:49:29      all      0.00      0.00      0.00      0.25      0.00      0.10      0.05      0.00      0.00
99.60
```

Average:	all	0.02	0.00	0.01	0.08	0.00	0.08	0.03	0.00	0.00
99.78										

To see the statistics from a specific core, use the **-P** switch:

```
[root@centos8 ~]# sar -u -P 1 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_      (8 CPU)

01:51:52      CPU      %user      %nice      %system      %iowait      %steal      %idle
01:51:57        1      0.00      0.00      0.00      0.00      0.00     100.00
01:52:02        1      0.20      0.00      0.00      0.00      0.00     99.80
01:52:07        1      0.00      0.00      0.00      0.00      0.00     100.00
Average:        1      0.07      0.00      0.00      0.00      0.00     99.93

[root@centos8 ~]# sar -u -P 5 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_      (8 CPU)

01:52:16      CPU      %user      %nice      %system      %iowait      %steal      %idle
01:52:21        5      0.00      0.00      0.00      0.00      0.00     100.00
01:52:26        5      0.00      0.00      0.00      0.00      0.00     100.00
01:52:31        5      0.00      0.00      0.00      0.00      0.00     100.00
Average:        5      0.00      0.00      0.00      0.00      0.00     100.00
```

Memory and Swap Statistics

Use the **-r** switch to see memory statistics:

```
[root@centos8 ~]# sar -r 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      30/06/21      _x86_64_ (8 CPU)

07:33:32      kbmemfree  kbavail  kbmemused  %memused  kbbuffers  kbcached  kbcommit  %commit  kbactive  kbinact
kbdirty
07:33:37      1647240   2297232  2177792    56.94     3356      827396   5096432   70.94    359072   1486368
```

```

0
07:33:42      1647232    2297224    2177800    56.94    3356    827396    5095788    70.93    359072    1486300
0
07:33:47      1647232    2297224    2177800    56.94    3356    827396    5095788    70.93    359072    1486376
0
Average:      1647235    2297227    2177797    56.94    3356    827396    5096003    70.94    359072    1486348
0

```

Use the **-S** switch to see swap statistics:

```

[root@centos8 ~]# sar -S 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      30/06/21      _x86_64_ (8 CPU)

07:31:58      kbswpfree kbswpused  %swpused  kbswpcad  %swpcad
07:32:03      3358716      0      0.00      0      0.00
07:32:08      3358716      0      0.00      0      0.00
07:32:13      3358716      0      0.00      0      0.00
Average:      3358716      0      0.00      0      0.00

```

I/O Statistics

Use the **-b** switch:

```

[root@centos8 ~]# sar -b 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_      (8 CPU)

09:24:49      tps      rtps      wtps      bread/s      bwrtn/s
09:24:54      0.00      0.00      0.00      0.00      0.00
09:24:59      1.20      0.00      1.20      0.00      20.20
09:25:04      0.00      0.00      0.00      0.00      0.00
Average:      0.40      0.00      0.40      0.00      6.73

```

Disk I/O Statistics

Use the **-d** switch:

```
[root@centos8 ~]# sar -d 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      29/06/21      _x86_64_      (8 CPU)

09:25:45      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
09:25:50      dev8-0    0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:50      dev8-16   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:50      dev11-0   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:50      dev253-0  0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:50      dev253-1  0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00

09:25:50      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
09:25:55      dev8-0    0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:55      dev8-16   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:55      dev11-0   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:55      dev253-0  0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:25:55      dev253-1  0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00

09:25:55      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
09:26:00      dev8-0    0.60      0.00      0.30      0.50      0.01      13.00      13.00      0.78
09:26:00      dev8-16   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:26:00      dev11-0   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
09:26:00      dev253-0  0.60      0.00      0.50      0.83      0.01      12.67      13.00      0.78
09:26:00      dev253-1  0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00

Average:      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
Average:      dev8-0    0.20      0.00      0.10      0.50      0.00      13.00      13.00      0.26
Average:      dev8-16   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
Average:      dev11-0   0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
Average:      dev253-0  0.20      0.00      0.17      0.83      0.00      12.67      13.00      0.26
```

Average:	dev253-1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
----------	----------	------	------	------	------	------	------	------	------

The **DEV** column identifies the disks by their major/minor numbers. To see the names of the disks add the, **-p** switch:

```
[root@centos8 ~]# sar -p -d 5 3
Linux 4.18.0-305.3.1.el8.x86_64 (centos8.ittraining.loc)      30/06/21      _x86_64_      (8 CPU)

07:48:32      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
07:48:37      sda      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:37      sdb      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:37      sr0      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:37  cl_centos8-root      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:37  cl_centos8-swap      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00

07:48:37      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
07:48:42      sda      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:42      sdb      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:42      sr0      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:42  cl_centos8-root      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:42  cl_centos8-swap      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00

07:48:42      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
07:48:47      sda      0.40      0.00      0.40      1.00      0.02      56.00      56.50      2.26
07:48:47      sdb      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:47      sr0      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
07:48:47  cl_centos8-root      0.40      0.00      0.80      2.00      0.02      56.00      56.50      2.26
07:48:47  cl_centos8-swap      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00

Average:      DEV      tps      kB/s      kB/s      areq-sz      aqu-sz      await      svctm      %util
Average:      sda      0.13      0.00      0.13      1.00      0.01      56.00      56.50      0.75
Average:      sdb      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
Average:      sr0      0.00      0.00      0.00      0.00      0.00      0.00      0.00      0.00
Average:  cl_centos8-root      0.13      0.00      0.27      2.00      0.01      56.00      56.50      0.75
```

Average:	cl_centos8-swap	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
----------	-----------------	------	------	------	------	------	------	------	------

The command line switches of this command are:

```
[root@centos8 ~]# sar --help
Usage: sar [ options ] [ <interval> [ <count> ] ]
Main options and reports (report name between square brackets):
  -B      Paging statistics [A_PAGE]
  -b      I/O and transfer rate statistics [A_IO]
  -d      Block devices statistics [A_DISK]
  -F [ MOUNT ]
          Filesystems statistics [A_FS]
  -H      Hugepages utilization statistics [A_HUGE]
  -I { <int_list> | SUM | ALL }
          Interrupts statistics [A_IRQ]
  -m { <keyword> [,...] | ALL }
          Power management statistics [A_PWR_...]
          Keywords are:
          CPU      CPU instantaneous clock frequency
          FAN      Fans speed
          FREQ     CPU average clock frequency
          IN       Voltage inputs
          TEMP     Devices temperature
          USB      USB devices plugged into the system
  -n { <keyword> [,...] | ALL }
          Network statistics [A_NET_...]
          Keywords are:
          DEV      Network interfaces
          EDEV     Network interfaces (errors)
          NFS      NFS client
          NFSD     NFS server
          SOCK     Sockets (v4)
          IP       IP traffic          (v4)
          EIP      IP traffic          (v4) (errors)
```

```

    ICMP    ICMP traffic      (v4)
    EICMP   ICMP traffic      (v4) (errors)
    TCP     TCP traffic       (v4)
    ETCP    TCP traffic       (v4) (errors)
    UDP     UDP traffic       (v4)
    SOCK6   Sockets (v6)
    IP6     IP traffic        (v6)
    EIP6    IP traffic        (v6) (errors)
    ICMP6   ICMP traffic      (v6)
    EICMP6  ICMP traffic      (v6) (errors)
    UDP6    UDP traffic       (v6)
    FC      Fibre channel HBAs
    SOFT    Software-based network processing
-q        Queue length and load average statistics [A_QUEUE]
-r [ ALL ]
          Memory utilization statistics [A_MEMORY]
-S        Swap space utilization statistics [A_MEMORY]
-u [ ALL ]
          CPU utilization statistics [A_CPU]
-v        Kernel tables statistics [A_KTABLES]
-W        Swapping statistics [A_SWAP]
-w        Task creation and system switching statistics [A_PCSW]
-y        TTY devices statistics [A_SERIAL]

```

USB Modules

The Universal Serial Bus can offer data transfer speeds of upto 480Mb/s for version 2.0 and upto 4.8 Gb/s for version 3.0. Under Linux the USB modules are:

USB Version	Module	Name
1.0\1.1	UHCI	<i>Universal Controller Host Interface</i>
	OHCI	<i>Open Controller Host Interface</i>

USB Version	Module	Name
2.0	EHCI	<i>Enhanced Host Controller Interface</i>
3.0	XHCI	<i>Extensible Host Controller Interface</i>

The following table shows a list of commonly used USB modules:

Module	Peripheral
usb_storage	Block devices
usbhid	Human Interface Device
snd-usb-audio	Sound cards
usbvideo	Video acquisition cards
irda-usb	IR peripherals
usbnet	NICs

These modules can be loaded by any one of the following:

- Initramfs,
- The init process,
- kmod by using the **/lib/modules/2.6.32-358.23.2.el6.i686/modules.usbmap** file,
- udev,
- manually.

udev

Since the 2.6 Kernel series, Linux capable of **hotplugging**. Linux uses three components to manage hotplugging:

- Udev,
- HAL,
- Dbus.

The roles of each component are as follows:

- Udev dynamically creates/deletes nodes in the **/dev** directory,
- HAL obtains information from udev, creates a files in XML format representing the peripheral and then informs Nautilus by using Dbus,
- Dbus acts as a system bus and is used for inter-process communication.

When Linux is booted, udev plays an important role:

- at boot **tmpfs** is mounted on **/dev**,
- udev copies any statically configured nodes from **/lib/udev/devices** to **/dev**,
- the **udev** daemon collects **uevents** from the kernel and looks for an appropriate rule in the **/lib/udev/rules.d/** directory,
- udev creates the nodes and any required symbolic links specified in the rule previously identified,
- udev stores in RAM the rules from **/lib/udev/rules.d/*.rules**,
- when a change occurs udev updates the the rules in the RAM.

udev uses the **sysfs** filesystem mounted on **/sys** which renders the peripherals visible to udev in user space. For example when a USB stick is inserted, udev creates **/dev/sdb1** automatically

The main configuration file for udev is **/etc/udev/udev.conf** :

```
[root@centos8 ~]# cat /etc/udev/udev.conf
# see udev.conf(5) for details
#
# udevd is also started in the initrd.  When this file is modified you might
# also want to rebuild the initrd, so that it will include the modified configuration.

#udev_log="info"
```

Rules files can be found in **/lib/udev/rules.d/** :

```
[root@centos8 ~]# ls /lib/udev/rules.d/
01-md-raid-creating.rules      70-uaccess.rules
10-dm.rules                    70-wacom.rules
11-dm-lvm.rules                71-biosdevname.rules
11-dm-mpath.rules              71-nvmf-iopolicy-netapp.rules
11-dm-parts.rules              71-prefixdevname.rules
```

```
13-dm-disk.rules
39-usbmuxd.rules
40-elevator.rules
40-libgphoto2.rules
40-redhat.rules
40-usb-blacklist.rules
40-usb_modeswitch.rules
50-udev-default.rules
60-alias-kmsg.rules
60-block.rules
60-cdrom_id.rules
60-drm.rules
60-evdev.rules
60-fido-id.rules
60-input-id.rules
60-libfprint-2-autosuspend.rules
60-net.rules
60-persistent-alsa.rules
60-persistent-input.rules
60-persistent-storage.rules
60-persistent-storage-tape.rules
60-persistent-v4l.rules
60-raw.rules
60-rdma-ndd.rules
60-rdma-persistent-naming.rules
60-sensor.rules
60-serial.rules
60-tpm-udev.rules
61-gdm.rules
61-gnome-bluetooth-rfkill.rules
61-gnome-settings-daemon-rfkill.rules
61-scsi-sg3_id.rules
62-multipath.rules
63-fc-wwpn-id.rules
71-seat.rules
73-idrac.rules
73-seat-late.rules
75-net-description.rules
75-probe_mtd.rules
75-rdma-description.rules
77-mm-cinterion-port-types.rules
77-mm-dell-port-types.rules
77-mm-ericsson-mbm.rules
77-mm-fibocom-port-types.rules
77-mm-haier-port-types.rules
77-mm-huawei-net-port-types.rules
77-mm-longcheer-port-types.rules
77-mm-mtk-port-types.rules
77-mm-nokia-port-types.rules
77-mm-pcmcia-device-blacklist.rules
77-mm-quectel-port-types.rules
77-mm-sierra.rules
77-mm-simtech-port-types.rules
77-mm-telit-port-types.rules
77-mm-ublox-port-types.rules
77-mm-usb-device-blacklist.rules
77-mm-usb-serial-adapters-greylist.rules
77-mm-x22x-port-types.rules
77-mm-zte-port-types.rules
78-sound-card.rules
80-drivers.rules
80-iio-sensor-proxy.rules
80-libinput-device-groups.rules
80-mm-candidate.rules
80-net-setup-link.rules
80-udisks2.rules
81-kvm-rhel.rules
84-nm-drivers.rules
```

```
63-md-raid-arrays.rules      85-nm-unmanaged.rules
63-scsi-sg3_symlink.rules   85-regulatory.rules
64-btrfs.rules              90-alsa-restore.rules
64-md-raid-assembly.rules   90-bolt.rules
65-libwacom.rules           90-fwupd-devices.rules
65-md-incremental.rules     Limiter les Ressources 90-iprutils.rules
65-sane-backends.rules      90-libinput-fuzz-override.rules
66-kpartx.rules             90-nm-thunderbolt.rules
68-del-part-nodes.rules     90-pulseaudio.rules
69-btattach-bcm.rules       90-rdma-hw-modules.rules
69-cd-sensors.rules         90-rdma-ulp-modules.rules
69-dm-lvm-metad.rules       90-rdma-umad.rules
69-libmtp.rules             90-vconsole.rules
69-md-clustered-confirm-device.rules 91-drm-modeset.rules
70-hypervfcopy.rules        95-cd-devices.rules
70-hypervkvp.rules          95-dm-notify.rules
70-hypervvss.rules          95-upower-csr.rules
70-joystick.rules           95-upower-hid.rules
70-mouse.rules              95-upower-wup.rules
70-nvmf-autoconnect.rules   98-kexec.rules
70-power-switch.rules       99-qemu-guest-agent.rules
70-printers.rules           99-systemd.rules
70-spice-vdagentd.rules     99-vmware-scsi-udev.rules
70-touchpad.rules
```



Important : You can create your own rules by putting them in the **99-local.rules** file.

The default udev rule file is **50-udev-default.rules**:

```
[root@centos8 ~]# cat /lib/udev/rules.d/50-udev-default.rules | more
# do not edit this file, it will be overwritten on update
```

```
# run a command on remove events
ACTION=="remove", ENV{REMOVE_CMD}!="", RUN+="$env{REMOVE_CMD}"
ACTION=="remove", GOTO="default_end"

SUBSYSTEM=="virtio-ports", KERNEL=="vport*", ATTR{name}=="?*", SYMLINK+="virtio-ports/${attr{name}}"

# select "system RTC" or just use the first one
SUBSYSTEM=="rtc", ATTR{hctosys}=="1", SYMLINK+="rtc"
SUBSYSTEM=="rtc", KERNEL=="rtc0", SYMLINK+="rtc", OPTIONS+="link_priority=-100"

SUBSYSTEM=="usb", ENV{DEVTYPE}=="usb_device", IMPORT{builtin}="usb_id", IMPORT{builtin}="hwdb --subsystem=usb"
ENV{MODALIAS}!="", IMPORT{builtin}="hwdb --subsystem=$env{SUBSYSTEM}"

ACTION!="add", GOTO="default_end"

SUBSYSTEM=="tty", KERNEL=="ptmx", GROUP="tty", MODE="0666"
SUBSYSTEM=="tty", KERNEL=="tty", GROUP="tty", MODE="0666"
SUBSYSTEM=="tty", KERNEL=="tty[0-9]*", GROUP="tty", MODE="0620"
SUBSYSTEM=="tty", KERNEL=="sclp_line[0-9]*", GROUP="tty", MODE="0620"
SUBSYSTEM=="tty", KERNEL=="ttysclp[0-9]*", GROUP="tty", MODE="0620"
SUBSYSTEM=="tty", KERNEL=="3270/tty[0-9]*", GROUP="tty", MODE="0620"
SUBSYSTEM=="vc", KERNEL=="vcs*|vcsa*", GROUP="tty"
KERNEL=="tty[A-Z]*[0-9]|tymxc[0-9]*|pppox[0-9]*|ircomm[0-9]*|noz[0-9]*|rfcomm[0-9]*", GROUP="dialout"

SUBSYSTEM=="mem", KERNEL=="mem|kmem|port", GROUP="kmem", MODE="0640"

SUBSYSTEM=="input", GROUP="input"
SUBSYSTEM=="input", KERNEL=="js[0-9]*", MODE="0664"

SUBSYSTEM=="video4linux", GROUP="video"
SUBSYSTEM=="graphics", GROUP="video"
SUBSYSTEM=="drm", KERNEL!="renderD*", GROUP="video"
SUBSYSTEM=="dvb", GROUP="video"
SUBSYSTEM=="media", GROUP="video"
```

```

SUBSYSTEM=="cec", GROUP="video"

SUBSYSTEM=="drm", KERNEL=="renderD*", GROUP="render", MODE="0666"
SUBSYSTEM=="kfd", GROUP="render", MODE="0666"

SUBSYSTEM=="sound", GROUP="audio", \
  OPTIONS+="static_node=snd/seq", OPTIONS+="static_node=snd/timer"
--More--

```

Each rule has the following format:

KEY, [KEY, ...] NAME [, SYMLINK]

The Key is a **type=value** pair which uniquely identifies a peripheral. The **type** can be one of the following:

Type	Description	Examples
BUS	Bus type	usb, scsi, ide
KERNEL	The default name given to the peripheral by the kernel	hda, ttyUSB0, lp0
SUBSYSTEM	The default sub-system name given by the Kernel, generally identical to the BUS value	usb, scsi
DRIVER	The name of the module used by the peripheral	usb-storage
ID	The position of the peripheral on its bus	PCI bus id, USB id
PLACE	The topological position of a USB oeripheral on its bus.	S/O
SYSFS{filename}	The name of the peripheral file in /sys. This file contains the manufacturer's name, the label, the serial number and the UUID of the peripheral. Each rule can contains upto five references to files.	S/O
PROGRAM	An eventual external program to be called in order to identify the peripheral	S/O
RESULT	Value expected from PROGRAM	S/O

NAME and SYMLINK are used to tell udev what to do with the peripheral:

Type	Description
NAME	The name of the peripheral in /dev
SYMLINK	The eventual symbolic links that point to NAME

The udevadm Command

To obtain information from udev on a particular peripheral, you can use the **udevadm** command which has replaced the **udevinfo** command available in Red Hat/CentOS 5:

```
[root@centos8 ~]# udevadm info --query=all -n /dev/sda
P: /devices/pci0000:00/0000:00:07.0/ata3/host2/target2:0:0/2:0:0:0/block/sda
N: sda
S: disk/by-id/ata-QEMU_HARDDISK_QM00005
S: disk/by-id/scsi-0ATA_QEMU_HARDDISK_QM00005
S: disk/by-id/scsi-1ATA_QEMU_HARDDISK_QM00005
S: disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00005
S: disk/by-path/pci-0000:00:07.0-ata-1
E: DEVLINKS=/dev/disk/by-path/pci-0000:00:07.0-ata-1 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00005
/dev/disk/by-id/ata-QEMU_HARDDISK_QM00005 /dev/disk/by-id/scsi-0ATA_QEMU_HARDDISK_QM00005 /dev/disk/by-
id/scsi-1ATA_QEMU_HARDDISK_QM00005
E: DEVNAME=/dev/sda
E: DEVPATH=/devices/pci0000:00/0000:00:07.0/ata3/host2/target2:0:0/2:0:0:0/block/sda
E: DEVTYPEDISK
E: ID_ATA=1
E: ID_ATA_FEATURE_SET_SMART=1
E: ID_ATA_FEATURE_SET_SMART_ENABLED=1
E: ID_ATA_SATA=1
E: ID_ATA_WRITE_CACHE=1
E: ID_ATA_WRITE_CACHE_ENABLED=1
E: ID_BUS=ata
E: ID_MODEL=QEMU_HARDDISK
E: ID_MODEL_ENC=QEMU\x20HARDDISK\x20\x20\x20
E: ID_PART_TABLE_TYPE=dos
E: ID_PART_TABLE_UUID=b39ec5c8
E: ID_PATH=pci-0000:00:07.0-ata-1
E: ID_PATH_TAG=pci-0000_00_07_0-ata-1
E: ID_REVISION=2.5+
```

```
E: ID SCSI=1
E: ID SCSI_INQUIRY=1
E: ID_SERIAL=QEMU_HARDDISK_QM00005
E: ID_SERIAL_SHORT=QM00005
E: ID_TYPE=disk
E: ID_VENDOR=ATA
E: ID_VENDOR_ENC=ATA\x20\x20\x20\x20\x20
E: MAJOR=8
E: MINOR=0
E: SCSI_IDENT_LUN_ATA=QEMU_HARDDISK_QM00005
E: SCSI_IDENT_LUN_T10=ATA_QEMU_HARDDISK_QM00005
E: SCSI_IDENT_LUN_VENDOR=QM00005
E: SCSI_IDENT_SERIAL=QM00005
E: SCSI_MODEL=QEMU_HARDDISK
E: SCSI_MODEL_ENC=QEMU\x20HARDDISK\x20\x20\x20
E: SCSI_REVISION=2.5+
E: SCSI_TPGS=0
E: SCSI_TYPE=disk
E: SCSI_VENDOR=ATA
E: SCSI_VENDOR_ENC=ATA\x20\x20\x20\x20\x20
E: SUBSYSTEM=block
E: TAGS=:systemd:
E: USEC_INITIALIZED=8735808
```

The command line switches of this command are:

```
[root@centos8 ~]# udevadm --help
udevadm [--help] [--version] [--debug] COMMAND [COMMAND OPTIONS]
```

Send control commands or test the device manager.

Commands:

info	Query sysfs or the udev database
trigger	Request events from the kernel

```
settle      Wait for pending udev events
control     Control the udev daemon
monitor     Listen to kernel and udev events
test        Test an event run
test-builtin Test a built-in command
```

```
[root@centos8 ~]# udevadm info --help
udevadm info [OPTIONS] [DEVPATH|FILE]
```

Query sysfs or the udev database.

```
-h --help          Print this message
-V --version       Print version of the program
-q --query=TYPE    Query device information:
    name           Name of device node
    symlink        Pointing to node
    path           sysfs device path
    property       The device properties
    all            All values
-p --path=SYSPATH sysfs device path used for query or attribute walk
-n --name=NAME     Node or symlink name used for query or attribute walk
-r --root          Prepend dev directory to path names
-a --attribute-walk Print all key matches walking along the chain
                   of parent devices
-d --device-id-of-file=FILE Print major:minor of device containing this file
-x --export        Export key/value pairs
-P --export-prefix Export the key name with a prefix
-e --export-db     Export the content of the udev database
-c --cleanup-db   Clean up the udev database
```

The /sys Filesystem

The virtual filesystem **/sys** was introduced with the 2.6 Kernel. Its role is to identify and describe peripherals for udev:

```
[root@centos8 ~]# ls -l /sys
total 0
drwxr-xr-x.  2 root root 0 Jul 12 08:15 block
drwxr-xr-x. 33 root root 0 Jul 12 08:15 bus
drwxr-xr-x. 57 root root 0 Jul 12 08:15 class
drwxr-xr-x.  4 root root 0 Jul 12 08:15 dev
drwxr-xr-x. 14 root root 0 Jul 12 08:15 devices
drwxr-xr-x.  6 root root 0 Jul 12 08:15 firmware
drwxr-xr-x.  9 root root 0 Jul 12 08:15 fs
drwxr-xr-x.  2 root root 0 Jul 12 08:15 hypervisor
drwxr-xr-x. 15 root root 0 Jul 12 08:15 kernel
drwxr-xr-x. 153 root root 0 Jul 12 08:15 module
drwxr-xr-x.  2 root root 0 Jul 12 08:15 power
```

Each directory contains specific information:

- **block**
 - information concerning block devices
- **bus**
 - information concerning buses
- **class**
 - information concerning classes
- **devices**
 - information concerning the position of devices on their bus
- **firmware**
 - information concerning APCI
- **module**
 - information concerning kernel modules
- **power**
 - information concerning power management
- **fs**
 - information concerning file systems

For example:

```
[root@centos ~]# cat /sys/block/sda/sda1/size
2097152
```

The figure returned is in sectors.

LAB #4 - Limiting Resources

4.1 - ulimit

Resources available to users can be limited by using the **ulimit** command.

The **ulimit** command manages two limits:

- a *hard* hard limit by specifying the **-H** switch,
- a *soft* soft limit by specifying the **-S** switch.

The soft limit is the limit imposed on the user whilst the hard limit is the limit that a user can obtain once he has gone over the soft limit.

Only root can position a hard limit and only if the limit does not exceed real resource levels.

Root can define limits by editing the **/etc/security/limits.conf** file:

```
[root@centos8 ~]# cat /etc/security/limits.conf
# /etc/security/limits.conf
#
#This file sets the resource limits for the users logged in via PAM.
#It does not affect resource limits of the system services.
#
#Also note that configuration files in /etc/security/limits.d directory,
#which are read in alphabetical order, override the settings in this
```

```
#file in case the domain is the same or more specific.
#That means for example that setting a limit for wildcard domain here
#can be overridden with a wildcard setting in a config file in the
#subdirectory, but a user specific setting here can be overridden only
#with a user specific setting in the subdirectory.
#
#Each line describes a limit for a user in the form:
#
#<domain>          <type> <item> <value>
#
#Where:
#<domain> can be:
#     - a user name
#     - a group name, with @group syntax
#     - the wildcard *, for default entry
#     - the wildcard %, can be also used with %group syntax,
#         for maxlogin limit
#
#<type> can have the two values:
#     - "soft" for enforcing the soft limits
#     - "hard" for enforcing hard limits
#
#<item> can be one of the following:
#     - core - limits the core file size (KB)
#     - data - max data size (KB)
#     - fsize - maximum filesize (KB)
#     - memlock - max locked-in-memory address space (KB)
#     - nofile - max number of open file descriptors
#     - rss - max resident set size (KB)
#     - stack - max stack size (KB)
#     - cpu - max CPU time (MIN)
#     - nproc - max number of processes
#     - as - address space limit (KB)
#     - maxlogins - max number of logins for this user
```

```
# - maxsyslogins - max number of logins on the system
# - priority - the priority to run user process with
# - locks - max number of file locks the user can hold
# - sigpending - max number of pending signals
# - msgqueue - max memory used by POSIX message queues (bytes)
# - nice - max nice priority allowed to raise to values: [-20, 19]
# - rtprio - max realtime priority
#
#<domain>      <type> <item>      <value>
#
#*              soft   core        0
#*              hard   rss         10000
#@student       hard   nproc       20
#@faculty       soft   nproc       20
#@faculty       hard   nproc       50
#ftp            hard   nproc       0
#@student       -      maxlogins   4

# End of file
```



Important : The limit can be a number or the word **unlimited**.

For example if root adds the two following lines to `/etc/security/limits.conf`:

```
...
trainee         soft   nofile      1024
trainee         hard   nofile      4096
...
```

the number of open files for trainee is limited to 1024. However trainee can increase this limit to 4 096 by using the following command:

```
$ ulimit -n 4096
```

To see the list of the current limits use the **-a** switch:

```
[root@centos8 ~]# ulimit -a
core file size          (blocks, -c) unlimited
data seg size           (kbytes, -d) unlimited
scheduling priority     (-e) 0
file size               (blocks, -f) unlimited
pending signals         (-i) 14702
max locked memory       (kbytes, -l) 64
max memory size         (kbytes, -m) unlimited
open files              (-n) 1024
pipe size               (512 bytes, -p) 8
POSIX message queues    (bytes, -q) 819200
real-time priority      (-r) 0
stack size              (kbytes, -s) 8192
cpu time                (seconds, -t) unlimited
max user processes      (-u) 14702
virtual memory          (kbytes, -v) unlimited
file locks              (-x) unlimited
```

The command line switches of this command are:

```
[root@centos8 ~]# help ulimit
ulimit: ulimit [-SHabcdefiklmnpqrstuvxPT] [limit]
  Modify shell resource limits.
  Provides control over the resources available to the shell and processes
  it creates, on systems that allow such control.
  Options:
    -S      use the `soft' resource limit
    -H      use the `hard' resource limit
    -a      all current limits are reported
    -b      the socket buffer size
```

- c the maximum size of core files created
- d the maximum size of a process's data segment
- e the maximum scheduling priority (`nice')
- f the maximum size of files written by the shell and its children
- i the maximum number of pending signals
- k the maximum number of kqueues allocated for this process
- l the maximum size a process may lock into memory
- m the maximum resident set size
- n the maximum number of open file descriptors
- p the pipe buffer size
- q the maximum number of bytes in POSIX message queues
- r the maximum real-time scheduling priority
- s the maximum stack size
- t the maximum amount of cpu time in seconds
- u the maximum number of user processes
- v the size of virtual memory
- x the maximum number of file locks
- P the maximum number of pseudoterminals
- T the maximum number of threads

Not all options are available on all platforms.

If LIMIT is given, it is the new value of the specified resource; the special LIMIT values `soft', `hard', and `unlimited' stand for the current soft limit, the current hard limit, and no limit, respectively. Otherwise, the current value of the specified resource is printed. If no option is given, then -f is assumed.

Values are in 1024-byte increments, except for -t, which is in seconds, -p, which is in increments of 512 bytes, and -u, which is an unscaled number of processes.

Exit Status:

Returns success unless an invalid option is supplied or an error occurs.

4.2 - CGroups

Control Groups, also called **CGroups** are another, more modern way, of limiting resources.

CGroups are organised hierarchially just like processes. However as opposed to processes, CGroups are organised into **multiple** hierarchies called **Resource Controllers** or simply **Controllers**.

To consult the cgroup hierarchies, use the **lssubsys** command:

```
[root@centos8 ~]# lssubsys -am
bash: lssubsys: command not found...
Install package 'libcgroup-tools' to provide command 'lssubsys'? [N/y] y

* Waiting in queue...
The following packages have to be installed:
 libcgroup-0.41-19.el8.x86_64  Library to control and monitor control groups
 libcgroup-tools-0.41-19.el8.x86_64  Command-line utility programs, services and daemons for libcgroup
Proceed with changes? [N/y] y

* Waiting in queue...
* Waiting for authentication...
* Waiting in queue...
* Downloading packages...
* Requesting data...
* Testing changes...
* Installing packages...
cpuset /sys/fs/cgroup/cpuset
cpu,cpuacct /sys/fs/cgroup/cpu,cpuacct
blkio /sys/fs/cgroup/blkio
memory /sys/fs/cgroup/memory
devices /sys/fs/cgroup/devices
freezer /sys/fs/cgroup/freezer
```

```
net_cls,net_prio /sys/fs/cgroup/net_cls,net_prio
perf_event /sys/fs/cgroup/perf_event
hugetlb /sys/fs/cgroup/hugetlb
pids /sys/fs/cgroup/pids
rdma /sys/fs/cgroup/rdma
```

```
[root@centos8 ~]# lssubsys -am
cpuset /sys/fs/cgroup/cpuset
cpu,cpuacct /sys/fs/cgroup/cpu,cpuacct
blkio /sys/fs/cgroup/blkio
memory /sys/fs/cgroup/memory
devices /sys/fs/cgroup/devices
freezer /sys/fs/cgroup/freezer
net_cls,net_prio /sys/fs/cgroup/net_cls,net_prio
perf_event /sys/fs/cgroup/perf_event
hugetlb /sys/fs/cgroup/hugetlb
pids /sys/fs/cgroup/pids
rdma /sys/fs/cgroup/rdma
```

Systemd organises the processes in each cgroup. As a result, the CGroup resource management is closely aligned to Systemd's units.

At the top of the hierarchy we can see the root slice **--.slice**, under which we can find:

- **system.slice** - system services,
- **user.slice** - user sessions,
- **machine.slice** - virtual machines and containers.

Under these slices can be found:

- **scopes** - processes created by a **Fork**,
- **services** - processes created by a **Unit**.

Slices can be consulted with the following command:

```
[root@centos8 ~]# systemctl list-units --type=slice
```

UNIT	LOAD	ACTIVE	SUB	DESCRIPTION
-.slice	loaded	active	active	Root Slice
machine.slice	loaded	active	active	Virtual Machine and Container Slice
system-getty.slice	loaded	active	active	system-getty.slice
system-lvm2\x2dpvscan.slice	loaded	active	active	system-lvm2\x2dpvscan.slice
system-sshd\x2dkeygen.slice	loaded	active	active	system-sshd\x2dkeygen.slice
system-systemd\x2dfsck.slice	loaded	active	active	system-systemd\x2dfsck.slice
system-systemd\x2dhibernate\x2dresume.slice	loaded	active	active	system-systemd\x2dhibernate\x2dresume.slice
system-user\x2druntime\x2ddir.slice	loaded	active	active	system-user\x2druntime\x2ddir.slice
system-vncserver.slice	loaded	active	active	system-vncserver.slice
system.slice	loaded	active	active	System Slice
user-1000.slice	loaded	active	active	User Slice of UID 1000
user-42.slice	loaded	active	active	User Slice of UID 42
user.slice	loaded	active	active	User and Session Slice

LOAD = Reflects whether the unit definition was properly loaded.

ACTIVE = The high-level unit activation state, i.e. generalization of SUB.

SUB = The low-level unit activation state, values depend on unit type.

13 loaded units listed. Pass --all to see loaded but inactive units, too.

To show all installed unit files use 'systemctl list-unit-files'.

CGroup hierarchies can be seen by using the **systemd-cgls** command:

```
[root@centos8 ~]# systemd-cgls
Control group /:
-.slice
├─user.slice
│   └─user-42.slice
│       └─session-c1.scope
│           ├──1317 gdm-session-worker [pam/gdm-launch-environment]
│           ├──1459 /usr/libexec/gdm-wayland-session --register-session gnome-session --autostart
│           └─/usr/share/gdm/greeter/autostart
│               └─1856 /usr/libexec/gnome-session-binary --autostart /usr/share/gdm/greeter/autostart
```

```
├─1882 /usr/bin/gnome-shell
├─2059 /usr/bin/Xwayland :1024 -rootless -terminate -accessx -core -listen 4 -listen 5 -displayfd 6
├─2132 ibus-daemon --xim --panel disable
├─2135 /usr/libexec/ibus-dconf
├─2138 /usr/libexec/ibus-x11 --kill-daemon
├─2251 /usr/libexec/gsd-xsettings
├─2261 /usr/libexec/gsd-ally-settings
├─2268 /usr/libexec/gsd-clipboard
├─2271 /usr/libexec/gsd-color
├─2272 /usr/libexec/gsd-datetime
├─2273 /usr/libexec/gsd-housekeeping
├─2274 /usr/libexec/gsd-keyboard
├─2275 /usr/libexec/gsd-media-keys
├─2280 /usr/libexec/gsd-mouse
├─2281 /usr/libexec/gsd-power
├─2283 /usr/libexec/gsd-print-notifications
├─2284 /usr/libexec/gsd-rfkill
├─2285 /usr/libexec/gsd-screensaver-proxy
├─2290 /usr/libexec/gsd-sharing
├─2321 /usr/libexec/gsd-smartcard
├─2328 /usr/libexec/gsd-sound
├─2333 /usr/libexec/gsd-wacom
├─2432 /usr/libexec/ibus-engine-simple
└─user@42.service
  ├─xdg-permission-store.service
  │ └─2170 /usr/libexec/xdg-permission-store
  ├─pulseaudio.service
  │ └─1455 /usr/bin/pulseaudio --daemonize=no --log-target=journal
  ├─init.scope
  │ ├─1357 /usr/lib/systemd/systemd --user
  │ └─1377 (sd-pam)
  ├─at-spi-dbus-bus.service
  │ ├─2090 /usr/libexec/at-spi-bus-launcher
  │ └─2095 /usr/bin/dbus-daemon --config-file=/usr/share/defaults/at-spi2/accessibility.conf --nofork --
```



```
| | | | 1969 /usr/bin/dbus-daemon --config-file=/usr/share/defaults/at-spi2/accessibility.conf --nofork --
print-address 3
| | | |   1972 /usr/libexec/at-spi2-registryd --use-gnome-session
| | | |   dbus.service
| | | |   1786 /usr/bin/dbus-daemon --session --address=systemd: --nofork --nopidfile --systemd-activation --
syslog-only
| | | |   2183 /usr/libexec/ibus-portal
| | | |   2201 /usr/libexec/gnome-shell-calendar-server
| | | |   2225 /usr/libexec/goa-daemon
| | | |   2397 /usr/libexec/goa-identity-service
| | | |   2721 /usr/libexec/dconf-service
| | | |   evolution-addressbook-factory.service
| | | |   2727 /usr/libexec/evolution-addressbook-factory
| | | |   2771 /usr/libexec/evolution-addressbook-factory-subprocess --factory all --bus-name
org.gnome.evolution.dataserver.Subprocess.Backend.AddressBookx2727x2 --own-path /org/gnome/evo>
| | | |   gvfs-mtp-volume-monitor.service
lines 44-86
```

Using Systemd, several resources can be limited:

- **CPUShares** - default value = 1024,
- **MemoryLimit** - expressed in MB or GB, there is no default value,
- **BlockIOWeight** - expressed as a value between 10 and 1000, there is no default value,
- **StartupCPUShares** - the same as CPUShares but only used during startup,
- **StartupBlockIOWeight** - the same as BlockIOWeight but only used during startup,
- **CPUQuota** - used to limit CPU usage even when the system is doing nothing.



Important : You can check the `systemd.resource-control(5)` manual to see which CGroup parameters can be passed to `systemctl`.

Limiting Memory Usage

Start by creating the **hello-world.sh** script that will be used to work with CGroups :

```
[root@centos8 ~]# vi hello-world.sh
[root@centos8 ~]# cat hello-world.sh
#!/bin/bash
while [ 1 ]; do
    echo "hello world"
    sleep 360
done
```

Make the script runnable and test it:

```
[root@centos8 ~]# chmod u+x hello-world.sh
[root@centos8 ~]# ./hello-world.sh
hello world
^C
```

Now create a CGroup in the **memory** sub-system called **helloworld** :

```
[root@centos8 ~]# mkdir /sys/fs/cgroup/memory/helloworld
```

By default, this CGroup inherits the use of the **all** of the available memory. To avoid that, create a **40,000,000** Byte limit for this CGroup:

```
[root@centos8 ~]# echo 40000000 > /sys/fs/cgroup/memory/helloworld/memory.limit_in_bytes
[root@centos8 ~]# cat /sys/fs/cgroup/memory/helloworld/memory.limit_in_bytes
39997440
```



Important - Note that the 40,000,000 requested have become 39,997,440 which corresponds to a round number of Kernel Memory Pages, each of 4KB size (39,997,440 /


```
[1]+  Killed                ./hello-world.sh
```

The cgroup Command

This command creates a CGroup:

```
[root@centos8 ~]# cgroupcreate -g memory:helloworld2
[root@centos8 ~]# ls -l /sys/fs/cgroup/memory/helloworld2/
total 0
-rw-rw-r--. 1 root root 0 Jul 13 10:39 cgroup.clone_children
--w--w----. 1 root root 0 Jul 13 10:39 cgroup.event_control
-rw-rw-r--. 1 root root 0 Jul 13 10:39 cgroup.procs
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.failcnt
--w--w----. 1 root root 0 Jul 13 10:39 memory.force_empty
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.kmem.failcnt
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.kmem.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.kmem.max_usage_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:39 memory.kmem.slabinfo
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.kmem.tcp.failcnt
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.kmem.tcp.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.kmem.tcp.max_usage_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:39 memory.kmem.tcp.usage_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:39 memory.kmem.usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.max_usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.memsw.failcnt
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.memsw.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.memsw.max_usage_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:39 memory.memsw.usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.move_charge_at_immigrate
-r--r--r--. 1 root root 0 Jul 13 10:39 memory.numa_stat
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.oom_control
-----
1 root root 0 Jul 13 10:39 memory.pressure_level
```

```
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.soft_limit_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:39 memory.stat
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.swappiness
-r--r--r--. 1 root root 0 Jul 13 10:39 memory.usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:39 memory.use_hierarchy
-rw-rw-r--. 1 root root 0 Jul 13 10:39 notify_on_release
-rw-rw-r--. 1 root root 0 Jul 13 10:39 tasks
```

Set a memory limit for **helloworld2**:

```
[root@centos8 ~]# echo 40000000 > /sys/fs/cgroup/memory/helloworld2/memory.limit_in_bytes
```

The cgexec Command

This command inserts the limit into the CGroup **and** launches the script:

```
[root@centos8 ~]# cgexec -g memory:helloworld2 ./hello-world.sh &
[1] 37670

[root@centos8 ~]# hello world
[Enter]

[root@centos8 ~]#
```

The cgdelete Command

Once the script has terminated, this command deletes the CGroup:

```
[root@centos8 ~]# ps aux | grep *.sh
root      37670  0.0  0.0 12724  3112 pts/1    S   10:41   0:00 /bin/bash ./hello-world.sh
root      37685  0.0  0.0 12136  1148 pts/1    S+  10:42   0:00 grep --color=auto hello-world.sh
```

```
[root@centos8 ~]# kill 37670

root@centos8 ~]# ps aux | grep *.sh
root      37726  0.0  0.0 12136 1156 pts/1    R+   10:43   0:00 grep --color=auto hello-world.sh
[1]+  Terminated                  cgexec -g memory:helloworld2 ./hello-world.sh

[root@centos8 ~]# cgdelete memory:helloworld2

[root@centos8 ~]# ls -l /sys/fs/cgroup/memory/helloworld2/
ls: cannot access '/sys/fs/cgroup/memory/helloworld2/': No such file or directory
```

The `/etc/cgconfig.conf` File

In order for limits to be persistent, it is necessary to edit the `/etc/cgconfig.conf` file:

```
[root@centos8 ~]# vi /etc/cgconfig.conf
[root@centos8 ~]# cat /etc/cgconfig.conf
#
# Copyright IBM Corporation. 2007
#
# Authors:      Balbir Singh <balbir@linux.vnet.ibm.com>
# This program is free software; you can redistribute it and/or modify it
# under the terms of version 2.1 of the GNU Lesser General Public License
# as published by the Free Software Foundation.
#
# This program is distributed in the hope that it would be useful, but
# WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
#
# By default, we expect systemd mounts everything on boot,
# so there is not much to do.
# See man cgconfig.conf for further details, how to create groups
```

```
# on system boot using this file.
group helloworld2 {
    cpu {
        cpu.shares = 100;
    }
    memory {
        memory.limit_in_bytes = 40000;
    }
}
```



Important - Here you have created **two** limits : a memory limit of 40,000 Bytes and a **cpu.shares** limit of **100**. This latter corresponds to about 9,77% of the CPU.

Now create the **two** CGroups:

```
[root@centos8 ~]# cgcreate -g memory:helloworld2
[root@centos8 ~]# ls -l /sys/fs/cgroup/memory/helloworld2/
total 0
-rw-rw-r--. 1 root root 0 Jul 13 10:46 cgroup.clone_children
--w--w----. 1 root root 0 Jul 13 10:46 cgroup.event_control
-rw-rw-r--. 1 root root 0 Jul 13 10:46 cgroup.procs
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.failcnt
--w--w----. 1 root root 0 Jul 13 10:46 memory.force_empty
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.kmem.failcnt
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.kmem.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.kmem.max_usage_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:46 memory.kmem.slabinfo
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.kmem.tcp.failcnt
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.kmem.tcp.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.kmem.tcp.max_usage_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:46 memory.kmem.tcp.usage_in_bytes
```

```
-r--r--r--. 1 root root 0 Jul 13 10:46 memory.kmem.usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.max_usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.memsw.failcnt
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.memsw.limit_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.memsw.max_usage_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:46 memory.memsw.usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.move_charge_at_immigrate
-r--r--r--. 1 root root 0 Jul 13 10:46 memory.numa_stat
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.oom_control
-----r--. 1 root root 0 Jul 13 10:46 memory.pressure_level
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.soft_limit_in_bytes
-r--r--r--. 1 root root 0 Jul 13 10:46 memory.stat
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.swappiness
-r--r--r--. 1 root root 0 Jul 13 10:46 memory.usage_in_bytes
-rw-rw-r--. 1 root root 0 Jul 13 10:46 memory.use_hierarchy
-rw-rw-r--. 1 root root 0 Jul 13 10:46 notify_on_release
-rw-rw-r--. 1 root root 0 Jul 13 10:46 tasks
```

```
[root@centos8 ~]# cgcreate -g cpu:helloworld2
[root@centos8 ~]# ls -l /sys/fs/cgroup/cpu/helloworld2/
total 0
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cgroup.clone_children
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cgroup.procs
-r--r--r--. 1 root root 0 Jul 13 10:47 cpuacct.stat
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cpuacct.usage
-r--r--r--. 1 root root 0 Jul 13 10:47 cpuacct.usage_all
-r--r--r--. 1 root root 0 Jul 13 10:47 cpuacct.usage_percpu
-r--r--r--. 1 root root 0 Jul 13 10:47 cpuacct.usage_percpu_sys
-r--r--r--. 1 root root 0 Jul 13 10:47 cpuacct.usage_percpu_user
-r--r--r--. 1 root root 0 Jul 13 10:47 cpuacct.usage_sys
-r--r--r--. 1 root root 0 Jul 13 10:47 cpuacct.usage_user
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cpu.cfs_period_us
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cpu.cfs_quota_us
```

```
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cpu.rt_period_us
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cpu.rt_runtime_us
-rw-rw-r--. 1 root root 0 Jul 13 10:47 cpu.shares
-r--r--r--. 1 root root 0 Jul 13 10:47 cpu.stat
-rw-rw-r--. 1 root root 0 Jul 13 10:47 notify_on_release
-rw-rw-r--. 1 root root 0 Jul 13 10:47 tasks
```

The cgconfigparser Command

Apply the contents of the **/etc/cgconfig.conf** file using the **cgconfigparser** command:

```
[root@centos8 ~]# cgconfigparser -l /etc/cgconfig.conf
[root@centos8 ~]# cat /sys/fs/cgroup/memory/helloworld2/memory.limit_in_bytes
36864
[root@centos8 ~]# cat /sys/fs/cgroup/cpu/helloworld2/cpu.shares
100
```