



European Organisation for Astronomical Research in the Southern Hemisphere

Programme: VLT

Project/WP: VLT Software Infrastructure Maintenance

VLT Common Software CentOS Linux Installation Manual

Document Number: ESO-043702

Document Version: 20

Document Type: Manual (MAN)

Released On: 2023-08-07

Document Classification: Public

Prepared by: W. Pirani

Validated by:

Approved by:

Name



Authors

Name	Affiliation
Werther Pirani	ESO
Sylvie Feyrin	ESO

Change Record from previous Version

Affected Section(s)	Changes / Reason / Remarks
All	VLT-MAN-ESO-17200-2009.doc



Contents

f1. Introduction	5
1.1 Purpose.....	5
1.2 Scope.....	5
1.3 REFERENCE DOCUMENTS	5
1.4 GLOSSARY	6
1.5 STYLISTIC CONVENTIONS	6
1.6 FEEDBACK.....	6
2. PREREQUISITES.....	7
2.1 PREREQUISITE HARDWARE.....	7
2.2 PREREQUISITE COMPUTER MEDIA	7
2.3 PREREQUISITE DOCUMENTATION.....	7
2.4 PREREQUISITE INFORMATION	8
2.4.1 BASICS	8
2.4.2 NETWORK PARAMETERS.....	8
2.4.3 DISK LAYOUT.....	8
3. INSTALLING THE OPERATING SYSTEM	9
3.1 BOOTING from a SAS or SATA disk	9
3.2 BOOTING from a NVME SSD	11
3.3 Verifying the VLT SW installation	12
3.4 Upgrading the VLT SW installation	14
3.5 VERIFICATION NETWORK CONFIGURATION	15
3.6 Configure Linux Bonding	15
3.6.1 Network Script Configuration	15
3.6.2 Module Configuration	16
3.6.3 Bonding test.....	17
4. CONFIGURING THE VLT COMMON SOFTWARE	19
4.1 CONFIGURING ACCOUNTS FOR THE VLT COMMON SOFTWARE.....	19
4.2 CONFIGURING DATA Area.....	19
5. SITE-SPECIFIC CONFIGURATION	22
5.1 FOR CONSORTIA AND OTHER EXTERNAL INSTITUTES ONLY	22
5.1.1 ACCESS TO NFS SERVERS.....	22
5.1.2 ADDING USERS AND USER ENVIRONMENTS	22
6. MISCELLANEOUS CONFIGURATION.....	22
6.1 INSTALLATION OF VIRTUAL MACHINE.....	22

Table of Figures



CentOS Linux Installation Manual

Doc. Number: ESO-303533

Doc. Version: 20

Released on: 2023-07-26

Page: 4 of 23



1. Introduction

1.1 Purpose

The purpose of this manual is to help you to install CentOS 8.5 Operating System on a PC so that the VLT2022 Release of VLT Common Software can be installed.

1.2 Scope

This document covers the following topics in the following order:

- prerequisites - including prerequisite information which you must gather before starting the installation
- configuring the hardware
- installing and configuring the core operating system
- network configuration
- site-specific configuring - principally for within ESO, but giving hints for other sites
- miscellaneous configuration items

This document assumes that you are familiar with basic system management activities for Unix systems. It is *not* a replacement for the standard CentOS installation documentation¹.

The procedure has been tested on the following models:

Dell Poweredge R720, R730 and R740 (server)²

Dell Poweredge 2950 (server)

Dell Blade M620, M630 and M820 (servers)³

NOTE: CentOS 7 has dropped support for PERC 4 RAID controllers. Systems with such card i.e. Dell PowerEdge 2850 **cannot** run starting from VLT2018.

Many more hardware combinations are likely to work as well, the only problematic areas being unsupported graphics cards and network adaptors, for which you will have to download and install up to date drivers. Unfortunately, we're unable to provide details in this document, but you can send email to vltnmgr@eso.org for support and we will try to help as best as we can.

Although a certain level of flexibility in install-time configuration is attempted in this document, it is not feasible to support every conceivable combination of hardware and peripherals, so you are strongly advised to read through this document thoroughly to familiarize yourself with the steps and with your own setup.

1.3 REFERENCE DOCUMENTS

VLT-MAN-ESO 17200-0642 --- VLT Common Software Installation Manual

¹ CentOS installation documentation is available from <https://wiki.centos.org/Documentation>

² R730 and R740 models don't work with previous release

³ Has been tested model with dual and quad core CPU



1.4 GLOSSARY

DNS : Domain Name System.

A distributed hierarchical system for resolving hostnames into IP addresses.

NTP: Network Time Protocol.

A standard defined to facilitate accurate time synchronization across a network, either in hierarchical or peer-to-peer fashion.

1.5 STYLISTIC CONVENTIONS

VLT Common Software-specific abbreviations and acronyms are explained either in footnotes or in the glossary.

Certain sections of the installation procedure are optional. All such sections are identified as such in footnotes, and explanations are provided as the circumstances under which it should be completed.

Courier font is used to indicate text displayed by, or to be entered into the system. Roman font is used to show displayed text, bold font for entered text, selected text or window buttons. Italicized text in angled brackets indicates placeholders or descriptions of fields.

Unix prompts indicate the login of the person a command should be run as. These prompts end with a hash, dollar or percent sign to indicate root login id, Bourne shell or C-shell/T-C-shell shells respectively. e.g.:

```
root# pwd
/sbin/init.d
root# echo "<your_name>"
<your_name>
root#
```

1.6 FEEDBACK

If you experience problems using this document, please send an email to vltnmgr@eso.org or report the problem using the JIRA system (<http://www.jira.eso.org>).



2. PREREQUISITES

Please use this section as a checklist: make sure you have compliant hardware, you have copies of the required computer media, and note down the information that will be required later on, some of which you will be able to decide yourself, and some of which you must decide in cooperation with your users and network administrators.

2.1 PREREQUISITE HARDWARE

A typical configuration would be:

- Pentium or Athlon CPU
- 4096MB of RAM or more (4 GB recommended)
- 128MB graphics card or better (NVIDIA Chipset recommended)
- 20" LCD monitor
- 160GB hard disk or bigger⁴
- USB Mouse
- Ethernet card (or onboard chip)
- DVD-ROM drive (physical or virtual one)

(Check the documentation at <http://www.redhat.com/docs/manuals/linux> for details.)

2.2 PREREQUISITE COMPUTER MEDIA

The computer media consists of two DVD ISO images:

CentOS-8.5-<date>.iso	DVD image for CentOS 8.5 and VLT SW 2022 with VxWorks
CentOS-8.5-<date>-novxworks.iso	DVD image for CentOS 8.5 and VLT SW 2022 without VxWorks

2.3 PREREQUISITE DOCUMENTATION

Besides this document, there are no prerequisite documents.

This document is distributed as part of the VLT Common Software Documentation Kit. However, you should check for late changes, errata, on our web pages at: <https://www.eso.org/sdd/bin/view/SDDPublic/VLTCommonSoftware>

⁴ Direct SATA SCSI SAS or ATA disks are supported but some old RAID controllers (e.g. PERC 4) are not longer supported and the disks might not be visible to the OS.



2.4 PREREQUISITE INFORMATION

2.4.1 BASICS

- What is the system to be named?
- How many hard disks does the system have attached?⁵
- How much physical memory does the system have?⁶
- Is the system to remain at ESO? If yes, then at ESO Headquarters Garching? If yes, then in BASE, DEV, JENKINS or VCM? See option **role=<code>** below.

2.4.2 NETWORK PARAMETERS

- Is the machine to be attached to a network?
- What will be the machine's IP address? Or will it get its IP from a DHCP server?
- What are the subnet masks for the networks on each interface?
- If there is a default gateway available then what is its name and IP address?
- If the system will use DNS, what is the machine's DNS domain name and which is the DNS server IP address and hostname?
- If there is an NTP server available on the network, what is its name and IP address?

For systems remaining at ESO Headquarters Garching the following information should be useful:

- All machines are connected to the network
- IP addresses on Ethernet are assigned by ESO IT Service Desk (servicedesk@eso.org)
- Default gateways are <subnet_address>.254
- Default gateway names can be determined by reverse-lookup with nslookup
- DNS domain is hq.eso.org.
- DNS primary server is 134.171.7.11
- For systems at ESO-Garching ntpvcm(134.171.58.142) is the NTP server

For systems at ESO Paranal or ESO La Silla, please consult local VLT and network administrators.

For systems elsewhere, please consult your local network administrators.

2.4.3 DISK LAYOUT

For your information, this is the disk layout created by the automated installation .

⁵ If the system already has some Linux distribution installed, then a list of attached disk devices with their sizes can be obtained as root with the command **fdisk -l** .

⁶ Use "**grep MemTotal /proc/meminfo**"



disk	filesystem	Paranal-like	Garching Consortia	Recommendation	Your system
1	(swap)	2048 MB	2048 MB	2* <system_memory> but not more than 8192 MB ⁷	
1	/	40000 MB	40000 MB	Not less than 10000MB	
1	/diska	Remaining space	Remaining space	<remaining_space> ⁸	
2	/data	Whole disk	0 Mb	0 Mb	

3. INSTALLING THE OPERATING SYSTEM

We provide the CentOS 8.5 distribution and the VLT Common Software on a DVD ISO image. You can use it to burn your own (double-layer 8.5GB) DVD or use it as it is on machines that accept virtual devices.

Thanks to Linux' *kickstart*, the installation procedure has been largely automated and only little user interaction is required.

WARNING: The installation will wipe out the contents of the primary master (i.e., the first disk attached to the first disk controller) -- even if the disk contains partitions dedicated to different operating systems.

WARNING: Don't use this ISO DVD image to install the OS without following our procedure, the installation may fail because we have removed many packages from the original ISO image.

If you want to install the Linux CentOS for other purpose different from the VLT software installation, please download the original ISO image from "The CentOS Project" page: <https://www.centos.org>

3.1 BOOTING from a SAS or SATA disk

Ensure that the BIOS is set to search for boot media in the sequence: floppies, CD/DVD-ROM, hard disks.

Also, make sure your computer is connected to the network at this stage⁹.

⁷ Any swap bigger than 8192 MB is a waste of disk space, unless you are planning running Virtual Machines.

⁸ Option `diska=<sizeMB>` in the command line of the kickstart can be used to override default value. The remaining disk space can then be used for other purposes, e.g. another OS.

⁹ If the workstation is not connected to the network please refer to the chapter 4



Note: During the POST (Power On Self Test) phase, a single line at the top of the screen tells you which key to press to get to the BIOS settings -- it's usually F2 or DELete, but your mileage may vary. To change the boot sequence as described above, look for a menu entry that reads "Boot Menu", "Boot Sequence", or something similar.

Insert on of the DVDs or attach its ISO image to the DVD-ROM drive.

CentOS-8.5-<date>.iso	DVD image for CentOS 8.5 and VLTSW 2022 with VxWorks
CentOS-8.5-<date>-novxworks.iso	DVD image for CentOS 8.5 and VLTSW 2022 without VxWorks

If the system is off, then power it on, otherwise run the command:

```
root# reboot
```

Shortly you will be presented with the "CentOS Linux 8" menu screen: press the **ESC** key and the "boot:" prompt will appear.

To install a machine with a static IP address, type the following command and then press ENTER:

```
linux ks=cdrom:/isolinux/ks.cfg ip=<IP_ADDR>
hostname=<fully_qualified_name> netmask=255.255.255.0
gateway<IP_ADDR> swap=<sizeMB> ntpserver=<IP_or_hostname>
nameserver=<IP_ADDR> disk=<sizeMB> role=<code>
```

To install a machine with a DHCP, type the following command and then press ENTER:

```
linux ks=cdrom:/isolinux/ks.cfg ip=dhcp
hostname=<fully_qualified_name> swap=<sizeMB>
ntpserver=<IP_or_hostname> nameserver=<IP_ADDR> disk=<sizeMB>
role=<code>
```

where the following entries are mandatory (otherwise the operator will be inquired):

ip=<IPv4_ADDR> dhcp	IP address in IPv4 notation, e.g. 134.171.58.123 or use a DHCP server
hostname=<fully_qualified_name>	Fully qualified name, e.g. te13.hq.eso.org
netmask=255.255.255.0	Netmask in IPv4 notation (ignored if IP=dhcp)
gateway=<IPv4_ADDR>	Gateway in IPv4 notation (ignored if IP=dhcp)

and the following arguments are optional, with defaults:

nameserver=<IPv4_ADDR>	DNS primary nameserver
ntpserver=<IPv4_ADDRE or hostname>	NTP server
swap=<sizeMB>	Size of swap partition in MB. Default: 2xRAM; max. 8192
disk=<sizeMB>	Size of /diska partition in MB. Default: remaining disk space
role=<BASE,DEV,JENKINS,VCM>	Default BASE. Other options valid for ESO-Garching only!

- The first hard disk is partitioned (40 GB for /, the amount you just specified for swap, the rest



for /diska unless you enter also option `diska=<sizeMB>` which could be used to reduce disk partition and leave some disk space unused.

- A few standard services are enabled (e.g. telnet, rlogin, rexec, rsh, ftp)
- The VLT Common Software binaries are installed
- vltmgr, vlt and vx accounts are created with local HOME directories

Note 1: Your system will reboot at the end of the installation process -- be sure to be around to remove the DVD-R from the DVD-ROM drive. If you don't, the installation will start from scratch but, luckily enough, will stop at the very start offering you to check the inserted media before proceeding any further -- just press **CTRL-ALT-DEL** to reboot and remove the DVD (or detach the ISO image) from the DVD-ROM drive to boot your newly installed system.

Note 2: Don't worry if booting will take longer than usual -- this will be fixed once the network has been set up properly.

Note 3: If your system is going to use DHCP, simply enter dummy values for IP and GATEWAY. After the system is installed you can change from STATIC to DHCP with command **nmtui** or editing the network configuration files at `/etc/sysconfig/network-scripts/ifcfg.<nic>`

3.2 BOOTING from a NVME SSD

Modern servers might be provided with a NVME SolidState disk (SDDs with direct PCIe access) which are considerably faster disks than SAS or SATA disks even if they are SDDs.

If you want to boot from a NVME disk the booting disk the sequence does not follow the same boot sequence as on SAS or SATA disks (see 3.1)

First of all, the NVME SDD disk will not be detected as a bootable device unless you change in the **BIOS** setup that your boot mode will be **UEFI** (instead of default BIOS).

With the UEFI boot mode selected and the VLT2022 ISO image mounted on a virtual or real CD/DVD device, the CentOS welcome windows appears a bit different, for instance the ESC key has no reaction.

Select "Troubleshooting" and you will see a new menu of options: between them there are two options prepared to install the VLTSW our two NVME machines in Garching: "wdev44" and "wjen37". Each represents a command line like this:

```
linuxefi /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=CentOS\x207\x20x86_64
inst.ks=cdrom:/isolinux/ks.cfg ip=134.171.12.26 hostname=wdev44.hq.eso.org
netmask=255.255.255.0 gateway=134.171.12.254 nameserver=134.171.7.11
ntpserver=ntpvcn.hq.eso.org role=DEV
```



You have then the option to edit any option. Take either "wdev44" or "wjen37" and the command-line will appear for edition. You can then modify any value you wish: hostname, ip, .etc... By the way role=DEV or role=JENKINS are role specific to ESO Garching. **Just remove "role=XXXX" or make it "role=BASE".**

Once your edited command-line is ready, let it run. The installation of the VLT2022 will continue as on SAS or SATA disks.

3.3 BOOTING from USB

3.3.1 BURNING THE ISO TO THE USB MEMORY STICK

- 1) Make sure the USB memory stick is at least 16 GB in size
- 2) Download Rufus from <https://rufus.ie/>. Available for Windows only, but our tests showed it's the most reliable because it takes care of media formatting and the user don't have to supply the correct Volume label (which is essential for the whole process to work).

3.3.2 BOOTING FROM USB MEMORY STICK

Bootting from the USB memory stick depends on the hardware in use, but typically the user must press a function key (likely F8 or F11) early on in the boot process and select the appropriate entry from the boot menu. Once the "Install CentOS 8.5" screen is displayed, follow the instructions for BIOS-based systems or the instructions for UEFI-based systems.

3.3.3 BIOS SYSTEMS

Press the ESC key at the "Install CentOS 8.5" screen and type the following at the linux prompt (on a single line):

```
linux boot ks=hd:LABEL=CENTOS\x207\x20X8
hostname=the.fully.qualified.hostname ip=nnn.nnn.nnn.nnn netmask=255.255.255.0
gateway=nnn.nnn.nnn.nnn nameserver=nnn.nnn.nnn.nnn
```

Make sure to type the label exactly as shown (CENTOS\x207\x20X8) and to provide the appropriate values for the host being installed.

Alternatively, you could just type:

```
linux boot ks=hd:LABEL=CENTOS\x207\x20X8
```

and provide all the values interactively by answering the questions to collect data.

Either way, the boot process will continue and you will have to select a disk for the installation (see below for the details).



3.3.4 UEFI SYSTEMS

At the "Install CentOS 8.5" screen use the arrow keys to select the entry labelled "Press 'e' to edit this entry and install CentOS Linux 8.5 on UEFI" and, as suggested, press 'e' and edit the following line:

```
linuxefi /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=CentOS\x207\x20x8  
inst.ks=hd:LABEL=CentOS\x207\x20x8:/ks.cfg hostname= ip=  
netmask=255.255.255.0 gateway= nameserver= ntpserver= role=BASE
```

Provide the correct values for ip, hostname, gateway, nameserver and, optionally, ntpserver (or remove this entry altogether if there isn't one) and press Ctrl-x to continue the booting process.

In a short while you will have to select a disk for the installation.

3.3.5 SELECTING A DISK FOR THE INSTALLATION

The "Select Target Disk for Installation" menu will display a list of the disks available (along their respective size) and you will have to pick the desired/correct one by typing its number and pressing ENTER.

WARNING 1: The USB memory stick itself is usually listed as /dev/sda (but due to its small size it's difficult to confuse it with the correct one).

WARNING 2: Make sure to select the correct disk. A couple of confirmation prompts will offer you a last chance to reconsider but, after that, the selected disk will be wiped clean any old data will be gone for good. The installation process will continue with the disk partitioning and the selection of the base RPMS.

From here on, the instructions are the same regardless of the installation media.

3.4 Verifying the VLT SW installation

You can execute the script:

```
root# cd /root/vlt  
root# ./puppet-check
```

It will exit 0 if all is OK, otherwise it will reports the inconsistencies between the machine and the VLT SW release.



Note: the script will produce errors if it does not have access to internet, as it tries to connect with the YUM repository

3.5 Upgrading the VLT SW installation

You can execute the script:

```
root# cd /root/vlt
root# ./puppet-force-align
```

It will update the current machine with new updates available in the YUM repository. Once it has finished it is recommended to execute the verification:

```
root# cd /root/vlt
root# ./puppet-check
```

Note: both scripts will produce errors if it does not have access to internet, as it tries to connect with the YUM repository



3.6 VERIFICATION NETWORK CONFIGURATION

Note: If you want to change or correct some of the network parameters given at the Kickstart command line keep reading this section.

Log in as root and (use **CTRL-ALT-F1** to work in text mode if you wish) and do the following:

- Edit **/etc/hostname** and make sure **HOSTNAME** is set to the correct hostname without the domain name (e.g. **tel3** instead of **tel3.hq.eso.org**).
- Edit **/etc/sysconfig/network-scripts/ifcfg-eth0** and verify that **IPADDR**, **NETMASK**, **GATEWAY** and **DNS1** (if you entered **nameserver=**) is set to the IP addresses assigned to this computer in the kickstart command-line. For instance, if you want to change the Kickstart-given static IP for a DHCP one, change the value of **BOOTPROTO** from “static” to “dhcp”.
- Edit **/etc/resolv.conf** and verify that **search** contains your domain name (e.g. **hq.eso.org**) and **nameserver** contains the correct IP address for your DNS server.

Finally, connect your system to the network and issue the command:

```
root# service network restart
```

After this step, your system should be able to access the network. You can try pinging its own IP address first, then the default gateway, and then some other system using both its IP address and its hostname.

If something doesn't work as expected, double check the files above to verify you entered the correct data and issue again the restart sequence.

Also very useful are the commands “ifconfig” and “ethtool”. The command “**ethtool <NIC>**” will tell you with a “Link detected: yes” if the NIC is alive. Example:

```
root# ethtool eth0
```

3.7 Configure Linux Bonding

In a very simple words, the Linux bonding driver (a kernel module) provides a method for aggregating multiple network interfaces into a single logical bonded interface. The behavior of the bonded interfaces depends upon the mode; generally speaking, modes provide either hot standby or load balancing services.

Additionally, link integrity monitoring may be performed.

Connect the 2 network interfaces in 2 different switches.

3.7.1 Network Script Configuration

In our Linux distribution network scripts resides under

/etc/sysconfig/network-scripts directory .



First create bond0 network interface script as root:

```
root# cd /etc/sysconfig/network-scripts
root# vi ifcfg-bond0
```

Add the following lines

(Example of wg0acc.hq.eso.org machine locate in VCM Garching):

```
DEVICE=bond0
IPADDR=134.171.58.1
NETMASK=255.255.255.0
NETWORK=134.171.58.0
BROADCAST=134.171.58.255
ONBOOT=yes
BOOTPROTO=none
USERCTL=no
GATEWAY=134.171.58.254
```

All the interfaces that are part of the trunk, must have SLAVE and MASTER definitions.

For example, if you wish to make eth0 and eth1 (or other interfaces) a part of the bonding interface

bond0 , their config files (ifcfg-eth0 , ifcfg-eth1 , etc) have to look like this:

```
root# vi ifcfg-eth0

DEVICE=eth0
USERCTL=no
ONBOOT=yes
MASTER=bond0
SLAVE=yes
BOOTPROTO=no
```

3.7.2 Module Configuration

Linux modules configuration file resides under /etc and it's called **modprobe.conf** :

As root



```
root# cd /etc
root# vi modprobe.conf
```

Add the following lines:

```
alias bond0 bonding
options bond0 mode-1 miimon=100 downdelay=200 updelay=200
```

3.7.3 Bonding test

Restart network services:

```
root# service network restart
```

Verify bond0 proc entry, output should be similar to:

```
wg0acc# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.2.4 (January 28, 2008)
```

```
Bonding Mode: fault-tolerance (active-backup)
Primary Slave: None
Currently Active Slave: eth0
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 200
Down Delay (ms): 200
```

```
Slave Interface: eth0
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:1e:4f:35:2a:af
```

```
Slave Interface: eth1
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:1e:4f:35:2a:b1
```

Verify bond0 interfaces with ifconfig command, output should be something similar to:

```
wg0acc# ifconfig bond0
bond0      Link encap:Ethernet  HWaddr 00:1E:4F:35:2A:AF
            inet addr:134.171.58.1  Bcast:134.171.58.255
Mask:255.255.255.0
            inet6 addr: fe80::21e:4fff:fe35:2aaf/64 Scope:Link
            UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
            RX packets:39978655 errors:0 dropped:0 overruns:0 frame:0
            TX packets:34446723 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:973435310 (928.3 MiB)  TX bytes:1558003684 (1.4 GiB)
```

```
wg0acc#
```



Now ping any WS to check if bond0 interface work well.

Now ping any WS and watch the output, during the ping disconnect one of the 2 network cable and check the output, if ping continue OK, then bond0 interface is working as expected.

In other case, re-check configurations files and test again



4. CONFIGURING THE VLT COMMON SOFTWARE¹⁰

Be sure that the network has been configured, the system is working properly, and the system has been rebooted at least once. Use the **CTRL-ALT-F1** combination to switch to the first text-only console and log in as root.

```
root# cd /etc/pecs/releases/000/etc/locality
root# vi apps-`hostname`.env
```

Then add the following line at the very bottom of `apps-`hostname`.env` (the shell will substitute `hostname` with the actual hostname of your computer, so the file will be called something like `apps-te101.env`):

```
export SOFTWARE_ROOTS=
export ACC_HOST=`hostname`
```

4.1 CONFIGURING ACCOUNTS FOR THE VLT COMMON SOFTWARE

To add new user see chapter 8.1.2

Use the **CTRL-ALT-F1** combination to switch to the first text-only console, and log in as a user of the VLT Common Software

```
-bash-3.00$ /etc/pecs/bin/pecssh mklinks -i
PECS_ROOTDIR [/etc/pecs]: (just press ENTER)
PECS_RELEASE [000]: (just press ENTER)
[...]
Do you wish to install VUE support files? [y]: (type n and press ENTER)
```

Repeat the above steps for any user of the VLT Common Software or remember to use the **mklinks** command each time a new user has been added: this will place a few startup scripts in his HOME directory.

This completes a fairly basic setup of VLT SW 2019 VLT Common Software on your computer. However, it's strongly recommended that you refer to the latest VLT Common Software Installation Manual (VLT-MAN-ESO-17200-0642). While the section about compiling the software itself is not relevant for you, the chapters about verification and configuration are likely to require your attention.

4.2 CONFIGURING DATA Area

To configure /data there are two ways, depend if you have one or two disks.

If you have 1 disk:

as root type

```
root# ln -s /diska/data /data
```

If you have 2 disks:

¹⁰**pecsmgr** , **vltmgr** , **vlt** and **vx** accounts have already been configured during the installation



as root type

```
root# fdisk -l
```

```
Disk /dev/sda: 73.2 GB, 73274490880 bytes
255 heads, 63 sectors/track, 8908 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	1	1275	10241406	83	Linux
/dev/sda2		1276	2189	7341705	82	Linux swap
/dev/sda3		2190	8908	53970367+	83	Linux

```
Disk /dev/sdb: 146.6 GB, 146695782400 bytes
255 heads, 63 sectors/track, 17834 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

Device	Boot	Start	End	Blocks	Id	System
--------	------	-------	-----	--------	----	--------

```
root#
```

As you notice the first disk (/dev/sda) it's already use
but the second disk not (/dev/sdb)

So as root type:

```
wvgvlti# fdisk /dev/sdb
```

The number of cylinders for this disk is set to 17834.

There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:

- 1) software that runs at boot time (e.g., old versions of LILO)
- 2) booting and partitioning software from other OSs
(e.g., DOS FDISK, OS/2 FDISK)

Command (m for help): **m**

Command action

- a toggle a bootable flag
- b edit bsd disklabel
- c toggle the dos compatibility flag
- d delete a partition
- l list known partition types
- m print this menu
- n add a new partition
- o create a new empty DOS partition table
- p print the partition table
- q quit without saving changes
- s create a new empty Sun disklabel
- t change a partition's system id
- u change display/entry units
- v verify the partition table
- w write table to disk and exit
- x extra functionality (experts only)



Command (m for help): **n**
Command action
e extended
p primary partition (1-4)

type the letter **p**

Partition number (1-4): 1
First cylinder (1-17834, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-17834, default 17834):
Using default value 17834

Command (m for help): type the letter **w**
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.

```
root#  
Now you are out of the fdisk utility and you have to create the xfs filesystem  
root# mkfs.xfs /dev/sdb1
```

Add the label on the disk so will be more easy to add the entry on the file **/etc/fstab**

```
root# tune2fs -L data /dev/sdb1
```

edit the file **/etc/fstab** and add the following line

```
LABEL=data          /data          xfs  defaults    1 2
```

```
root# mkdir /data  
root# chown -R vltmgr:vlt /data  
root# chmod -R 775 /data  
root# mount -a
```

edit the file **/etc/exports** and add the line
/data *(sync,rw,insecure_locks)

```
root# exportfs -av
```

That's all !



5. SITE-SPECIFIC CONFIGURATION

5.1 FOR CONSORTIA AND OTHER EXTERNAL INSTITUTES ONLY

Site-specific configuration for external sites is obviously somewhat beyond the scope of this document. However, what follows in this section and the sections dedicated to site-specific configuration within ESO may be useful. If you are unsure what to do, then seek help from your local Unix systems administrators.

5.1.1 ACCESS TO NFS SERVERS

If there are central software repositories available on your network then you will need to alter either `/etc/fstab` for permanent NFS mounts, or `/etc/auto.master` and `/etc/auto_staging` for automounted NFS mounts. These are probably best copied and carefully merged from a similar system within the same group at your site.

Remember to update remote systems' exports lists and re-export everything.

5.1.2 ADDING USERS AND USER ENVIRONMENTS

Since NIS has not been configured, the easiest way to create a large amount of accounts is to copy the relevant files from another system. This you can do either with `ftp` or `remsh` with a suitable `.rhosts` entry on the other system.

Files to consider are:

- `/etc/shadow`
- `/etc/passwd`
- `/etc/group`
- `/etc/profile`
- `/etc/csh.login`
- `/etc/PATH`
- `/etc/MANPATH`
- `/etc/auto.master`
- `/etc/auto.home`

Typically, the files cannot be simply installed as replacements, but must be carefully merged with existing files (e.g. to preserve root's password). So it may be useful to consider the approach used in the following example:

```
root# rcp <remote_host>:/etc/passwd /etc/passwd.new
root# diff /etc/passwd /etc/passwd.new | more
```

6. MISCELLANEOUS CONFIGURATION

6.1 INSTALLATION OF VIRTUAL MACHINE

With this release VLT2022 we distribute the Virtual Machine using the Open Virtualization Format



OVA.

You can download any of the VM images from the ESO ftp site and follow the instruction of the INSTALL.txt, see the page <https://www.eso.org/sdd/bin/view/SDDPublic/Download2022>

--- End of document ---