

15. Proxmox

Proxmox

: Proxmox ISO

Proxmox ISO

<https://www.proxmox.com/en/downloads/proxmox-virtual-environment/iso>

1. ISO Install Proxmox VE (Graphical)

Proxmox VE 8.3 (iso release 1) - <https://www.proxmox.com/>



Install Proxmox VE (Graphical)
Install Proxmox VE (Terminal UI)
Advanced Options

enter: select, arrow keys: navigate, e: edit entry, esc: back

Graphical(GUI)

2. EULA



Proxmox VE Installer

END USER LICENSE AGREEMENT (EULA)

END USER LICENSE AGREEMENT (EULA) FOR PROXMOX VIRTUAL ENVIRONMENT (PROXMOX VE)

By using Proxmox VE software you agree that you accept this EULA, and that you have read and understand the terms and conditions. This also applies for individuals acting on behalf of entities. This EULA does not provide any rights to Support Subscriptions Services as software maintenance, updates and support. Please review the Support Subscriptions Agreements for these terms and conditions. The EULA applies to any version of Proxmox VE and any related update, source code and structure (the Programs), regardless of the delivery mechanism.

1. License. Proxmox Server Solutions GmbH (Proxmox) grants to you a perpetual, worldwide license to the Programs pursuant to the GNU Affero General Public License V3. The license agreement for each component is located in the software component's source code and permits you to run, copy, modify, and redistribute the software component (certain obligations in some cases), both in source code and binary code forms, with the exception of certain binary only firmware components and the Proxmox images (e.g. Proxmox logo). The license rights for the binary only firmware components are located within the components. This EULA pertains solely to the Programs and does not limit your rights under, or grant you rights that supersede, the license terms of any particular component.

2. Limited Warranty. The Programs and the components are provided and licensed "as is" without warranty of any kind, expressed or implied, including the implied warranties of merchantability, non-infringement or fitness for a particular purpose. Neither Proxmox nor its affiliates warrants that the functions contained in the Programs will meet your requirements or that the operation of the Programs will be entirely error free, appear or perform precisely as described in the accompanying documentation, or comply with regulatory requirements.

3. Limitation of Liability. To the maximum extent permitted under applicable law, under no circumstances will Proxmox, its affiliates, any Proxmox authorized distributor, or the licensor

Abort

Previous

I agree

3. Proxmox OS가 Target Harddisk , Next .



Proxmox VE Installer

Proxmox Virtual Environment (PVE)

The Proxmox Installer automatically partitions your hard disk. It installs all required packages and makes the system bootable from the hard disk. All existing partitions and data will be lost.

Press the Next button to continue the installation.

- **Please verify the installation target**
The displayed hard disk will be used for the installation.
Warning: All existing partitions and data will be lost.
- **Automatic hardware detection**
The installer automatically configures your hardware.
- **Graphical user interface**
Final configuration will be done on the graphical user interface, via a web browser.

1

Target Harddisk /dev/sda (100.00GiB, QEMU HARDDISK)

Options

Abort

Previous

Next

2

4. Country Time zone .



Proxmox VE Installer

Location and Time Zone selection

The Proxmox Installer automatically makes location-based optimizations, like choosing the nearest mirror to download files from. Also make sure to select the correct time zone and keyboard layout.

Press the Next button to continue the installation.

- **Country:** The selected country is used to choose nearby mirror servers. This will speed up downloads and make updates more reliable.
- **Time Zone:** Automatically adjust daylight saving time.
- **Keyboard Layout:** Choose your keyboard layout.

Country	South Korea
Time zone	Asia/Seoul
Keyboard Layout	U.S. English

5. Email . Email Proxmox



Proxmox VE Installer

Administration Password and Email Address

Proxmox Virtual Environment is a full featured, highly secure GNU/Linux system, based on Debian.

In this step, please provide the *root* password.

- **Password:** Please use a strong password. It should be at least 8 characters long, and contain a combination of letters, numbers, and symbols.
- **Email:** Enter a valid email address. Your Proxmox VE server will send important alert notifications to this email account (such as backup failures, high availability events, etc.).

Press the Next button to continue the installation.

Password	••••••••
Confirm	••••••••
Email	proxmox@hostway.co.kr

6. NIC, Hostname, IP, DNS



Proxmox VE Installer

Management Network Configuration

Please verify the displayed network configuration. You will need a valid network configuration to access the management interface after installing.

After you have finished, press the Next button. You will be shown a list of the options that you chose during the previous steps.

- **IP address (CIDR):** Set the main IP address and netmask for your server in CIDR notation.
- **Gateway:** IP address of your gateway or firewall.
- **DNS Server:** IP address of your DNS server.

Management Interface: ens18 -

Hostname (FQDN): pve-node01.hostway.co.kr

IP Address (CIDR): .133 / 24

Gateway: .1

DNS Server:

Buttons: Abort, Previous, Next

Management Interface: Proxmox , API ,
NIC . NIC

Hostname (FQDN): Hostname + domain .

ex) hostname: pve-node01 / domain: hostway.co.kr , pve-
node01.hostway.co.kr

IP Address (CIDR): Proxmox Node IP Prefix, Gateway, DNS

7. Summary , Install OS



Proxmox VE Installer

Summary

Please confirm the displayed information. Once you press the **Install** button, the installer will begin to partition your drive(s) and extract the required files.

Option	Value
Filesystem:	ext4
Disk(s):	/dev/sda
Country:	South Korea
Timezone:	Asia/Seoul
Keymap:	en-us
Email:	support@hostway.co.kr
Management Interface:	ens18
Hostname:	pve-node01
IP CIDR:	172.16.16.133/24
Gateway:	172.16.16.1
DNS:	8.8.8.8

☒ Automatically reboot after successful installation

Abort

Previous

Install

8. , Proxmox VE GNU/Linux (,
Default .)

GNU GRUB version 2.06-13+pmx2

```
*Proxmox VE GNU/Linux
Advanced options for Proxmox VE GNU/Linux
Memory test (memtest86+x64.bin)
Memory test (memtest86+x64.bin, serial console)
```

Use the ↑ and ↓ keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the commands before booting or 'c'
for a command-line.
The highlighted entry will be executed automatically in 5s.

9. shell, ssh,

```
-----
Welcome to the Proxmox Virtual Environment. Please use your web browser to
configure this server - connect to:

https://172.16.16.133:8006/
-----

Hint: Num Lock on
pve-node01 login:
```

ssh: 22 port(default)

web console & API: 8006 port

open), (default

05. Proxmox Authenitcation

Proxmox Authenitcation

(VM, ,)
.
.
(API) ()
.

Proxmox VE Linux PAM(Linux), Proxmox VE(Web)
, LDAP, Microsoft Active Directory, OpenID Connect

: [Hostway Proxmox](#)
: Proxmox@hostway.co.kr

04. Proxmox Clustering

Proxmox Clstering

Proxmox Virtual Environment
(32)

Proxmox (**pmxcfs**)

Proxmox VE Proxmox가
[Proxmox](#) ([pmxcfs](#))

Proxmox (“pmxcfs”)
, corosync
Proxmox VE

RAM . 128MiB 가 .
가 .

-
- VM ID
- 가
- Corosync
-

/

/

가

Proxmox VE

Live Migration

가 가

- VM 가 ,
PCI USB

- Proxmox VE
- ()
- Proxmox VE

- 가 CPU 가
- VM CPU

KVM

Linux

. GUI

VM

가

: Hostway Proxmox
: Proxmox@hostway.co.kr

GlusterFS - dispersed Type Volume

GFS Volume Features - Dispersed Volume

GFS Volume type

HW RAID

- Distribute Volume (RAID 0)
- Replicate Volume (RAID 1)
- Distribute + Replicate Volume (RAID 0+1)
- **Dispersed Volume (RAID 5,6)**
- Dispersed + Replicate Volume (RAID 50 , 60)

Dispersed Volume?

HW RAID Level

RAID 5 6

3

Brick

가

가

(Redundancy)

RAID 5

Replicate 가

Distribute

?

Distribute

Brick

Dispersed erasure coding ()

Brick

Create Command

e.x)


```
#
gluster volume info
```

```
# . stop
gluster volume start
volume start: data: success
```

```
#
gluster volume status
Status of volume: data
Gluster process TCP Port RDMA
Port Online Pid
-----
-----
Brick 211.239.151.196:/gfs_node 49152 0
Y 65695
Brick 211.239.151.197:/gfs_node 49152 0
Y 13261
Brick 211.239.151.198:/gfs_node 49152 0
Y 13513
Self-heal Daemon on localhost N/A N/A
Y 65712
Self-heal Daemon on 211.239.151.198 N/A N/A
Y 13530
Self-heal Daemon on 211.239.151.197 N/A N/A
Y 13278
```

```
Task Status of Volume data
-----
-----
```

```
#
(All Node)
brick 2 1 redundancy( ) .
```

```
mount -t glusterfs gfs-node01:/data /mnt
```

```
df -h /mnt
```

```
Filesystem      Size  Used Avail Use% Mounted on
gfs-node01:/data 200G  2.1G  198G   2% /mnt
```

```
# fstab
```

```
echo "gfs-node01:data /mnt glusterfs defaults,_netdev 0 0" >>
/etc/fstab
```

Node Fail-over TEST

```
( )
```

```
** 1 node down
```

```
# Node          Poweroff      1          가
```

```
gluster status volume
```

```
** 2 node down
```

```
# Redundancy 가 1          2 가
```

```
[root@node01 mnt]# ll
```

```
ls: cannot open directory .: Transport endpoint is not
connected
```

```
( )
```

```
** 1 node up
```

```
2 가
```

```
up
```

```
[root@node01 mnt]# ll
```

```
ls: cannot open directory .: Transport endpoint is not
connected
```

```
[root@node01 mnt]# ll
```

```
total 0
```

```
-rw-r--r--. 1 root root 0 Jan 13 10:42 a
```

```
-rw-r--r--. 1 root root 0 Jan 13 10:44 b
```

```
-rw-r--r--. 1 root root 0 Jan 13 10:44 c
```

```
[root@node01 mnt]# gluster volume status
```

```
Status of volume: data
```

```
Gluster process
```

```
TCP Port  RDMA
```

```
Port  Online  Pid
```

```
-----
```

```
-----
```

```
Brick 211.239.151.196:/gfs_node      49152      0
Y      67046
```

```
Brick 211.239.151.197:/gfs_node      49153      0
```

Y	3551		
Self-heal Daemon on localhost		N/A	N/A
Y	67063		
Self-heal Daemon on 211.239.151.197		N/A	N/A
Y	3108		

Task Status of Volume data

** 2 node up

** 3 (all) node up

online .

Volume Expansion (ADD-BRICK)

```
# node03    100GB Brick 3EA    가    ..
Dispersed          3(1    )    Brick
6(2    )    가    .
-    GFS Volume
-
```

Check now

```
[root@node02 ~]# gluster volume info | egrep -iE "brick|t
ype"
```

Type: Disperse

Number of Bricks: 1 x (2 + 1) = 3

Transport-type: tcp

Bricks:

Brick1: 211.239.151.196:/gfs_node

Brick2: 211.239.151.197:/gfs_node

Brick3: 211.239.151.198:/gfs_node

```
#    가    BRICK    ( sdc , sdd , sde 3    )
```

```
fdisk /dev/sdc <<EOF
```

n

p

```

w
EOF
mkfs.xfs -i size=512 /dev/sdc1
mkdir -p /add_brick1
mount /dev/sdc1 /add_brick1

echo "/dev/sdc1          /add_brick1          xfs          defaults
0 0" >> /etc/fstab

# check
[root@node03 ~]# df -h | grep brick
/dev/sdc1          100G    33M   100G    1% /add_bric
1
/dev/sdd1          100G    33M   100G    1% /add_bric
2
/dev/sde1          100G    33M   100G    1% /add_bric
3

# Volume ADD issue
3          가          가
[root@node03 ~]# gluster volume add-brick data gfs-
node03:/add_brick1
volume add-brick: failed: Incorrect number of bricks supplied
1 with count 3

Peer (          ) 가          Brick
[root@node03 ~]# gluster volume add-brick data gfs-
node03:/add_brick1      gfs-node03:/add_brick2      gfs-
node03:/add_brick
3
volume add-brick: failed: Multiple bricks of a disperse volume
are present on the same server. This setup is not optimal.
Bricks should be on different nodes to have best fault
tolerant configuration. Use 'force' at the end of the command
if you want to override this behavior.

force
volume add-brick: success

# Volume brick check
6          2 가
[root@node03 ~]# gluster volume info | egrep -iE "brick"

```

Number of Bricks: $2 \times (2 + 1) = 6$

Bricks:

Brick1: 211.239.151.196:/gfs_node

Brick2: 211.239.151.197:/gfs_node

Brick3: 211.239.151.198:/gfs_node

Brick4: gfs-node03:/add_brick1

Brick5: gfs-node03:/add_brick2

Brick6: gfs-node03:/add_brick3

Volume Size Check

brick 가 volume size .

[root@node03 ~]# df -h | grep mnt

gfs-node01:data	400G	4.2G	396G	2%	/mnt
-----------------	------	------	------	----	------

Rebalance

가 Brick

gluster volume rebalance data start

volume rebalance: data: success: Rebalance on data has been started successfully. Use rebalance status command to check status of the rebalance process.

ID: 9782e187-6cda-4e2b-aae9-0b78746d69fa

check

1) status

Task Status of Volume data

Task	: Rebalance
------	-------------

ID	: 9782e187-6cda-4e2b-aae9-0b78746d69fa
----	--

Status	: completed
--------	-------------

2) 가 brick 가

[root@node03 add_brick3]# ll ()

total 0

[root@node03 add_brick3]# ll ()

total 8

-rw-r--r--. 2 root root 512 Jan 13 14:57 d

[root@node03 add_brick3]#

[AWS] EC2

- Amazon Linux 2

가 AWS

.

Nitro Hypervisor

T3

AWS Console

가가

GRUB Timeout

가

.

Amazon Linux 2 OS GRUB

.

2가

가

.

1) GRUB TIMEOUT

2) mount nofail 가 , 가

IAM

IAM Administrator Group ,
가 .

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "ec2:GetSerialConsoleAccessStatus",  
        "ec2:EnableSerialConsoleAccess",  
        "ec2:DisableSerialConsoleAccess"  
      ],  
    },  
  ],  
}
```

```

        "Resource": "*"
    }
}

```

GRUB TIMEOUT

```

# root Default GRUB
Console Serial root
1 10( )

sed -i 's/GRUB_TIMEOUT=0/GRUB_TIMEOUT=10/g' /etc/default/grub
sed -i 's/GRUB_TERMINAL="ec2-console"/GRUB_TERMINAL="console
serial"/g' /etc/default/grub
echo -e GRUB_SERIAL_COMMAND="\serial --speed=115200\" >>
/etc/default/grub

# check
cat /etc/default/grub

GRUB_CMDLINE_LINUX_DEFAULT="console=tty0
console=ttyS0,115200n8 net.ifnames=0 biosdevname=0
nvme_core.io_timeout=4294967295 rd.emergency=poweroff
rd.shell=0"
GRUB_TIMEOUT=10
GRUB_DISABLE_RECOVERY="true"
GRUB_TERMINAL="console serial"
GRUB_X86_USE_32BIT="true"
GRUB_SERIAL_COMMAND="serial --speed=115200"

# GRUB
grub2-mkconfig -o /boot/grub2/grub.cfg

```

EC2 Reboot

10 가 e



single 가 Ctrl + x



FSTAB nofail

가 nofail



[AWS](#) ()

[AWS] Cloud-init SSH RootPassword

AWS Amazon-Linux OS

 Password

, ssh

Cloudinit

.

2가

.

1. Advance User-data

 root

2. Cloudinit Config

가 .

sshd Permitrootlogin

sed -i 's/#PermitRootLogin yes/PermitRootLogin yes/g'
/etc/ssh/sshd_config

sed -i 's/PasswordAuthentication no/PasswordAuthentication
yes/g' /etc/ssh/sshd_config

Key root account permit

sed -i 's/^. *10" //g' /root/ssh/authorized_keys

```
# set root password
echo "P@ssw0rd" | passwd root --stdin
```

```
# sshd restart
systemctl restart sshd
```

User-data 가



root password 가 .



Cloudinit

```
disable_root : true --> false ( 0 )
ssh_pwauth: false --> true ( 1 )
```



Apache vhost

<Directory> .

1. <Directory>
- 2.

Directory .

1. <Directory>
2. <DirectoryMatch>
3. <Files> <FilesMatch>
4. <Location>

< >

Document /home/web VHOST .

<Directory> All grants 가 .

```
<Directory /home/web>
    AllowOverride All
    Require all granted
</Directory>
```

```
<VirtualHost *:80>
    ServerAdmin admin-bmi
    DocumentRoot "/home/web"
    ServerName bmitest.co.kr
    ServerAlias www.bmitest.co.kr
    ErrorLog "/home/web/logs/bmitest.co.kr-error_log"
    CustomLog "/home/web/logs/bmitest.co.kr" common
</VirtualHost>
```

#

[root@CentOS7-CT01 conf]# curl http://localhost/

BMI!

/home/web/logs/accesstest File

```
[root@CentOS7-CT01 conf]# curl
http://localhost/logs/accesstest
oops. This Page must No view!
```

Directory DirectoryMatch

```
<Directory "/home/web/">
    AllowOverride All
    Require all granted
</Directory>
```

```
<DirectoryMatch "^/.*logs/"> # DocumentRoot/Logs
    Order deny,allow
    Deny from all
</DirectoryMatch>
```

```
<VirtualHost *:80>
```

```

    ServerAdmin admin-bmi
    DocumentRoot "/home/web"
    ServerName bmitest.co.kr
    ServerAlias www.bmitest.co.kr
    ErrorLog "/home/web/logs/bmitest.co.kr-error_log"
    CustomLog "/home/web/logs/bmitest.co.kr" common
</VirtualHost>

#                                403
[root@CentOS7-CT01                conf]#                                curl
http://localhost/logs/accesstest
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>403 Forbidden</title>
</head><body>
<h1>Forbidden</h1>
<p>You don't have permission to access this resource.</p>
</body></html>

#
<Directory>                        가.                        URL      가      .

ErrorDocument 403 "No No NO!"

#
[root@CentOS7-CT01                conf]#                                curl
http://localhost/logs/accesstest
No No NO!

```

Ubuntu 22.04

Zabbix-agent

Zabbix Server Proxy 4.x Syntax 가

가 .

```

#
Ubuntu 18.04      Default Repository      Zabbix-Agent

```

3.0.12

Ubuntu 20.04 Default Repository Zabbix-Agent

4.0.17

22.04 Default Repository Zabbix-Agent

5.0.17

.

```
#
sudo apt-cache madison zabbix-agent

#                               Repo deb
apt remove zabbix-agent      #      repo
apt autoremove

wget
https://repo.zabbix.com/zabbix/4.0/ubuntu/pool/main/z/zabbix-r
elease/zabbix-release_4.0-4+ubuntu22.04_all.deb

dpkg -i zabbix-release_4.0-4+ubuntu22.04_all.deb

apt update
apt install zabbix-agent

zabbix_agentd -V      #
apt-mark hold zabbix-agent      #
```

growpart

() 가 .

/home (80G) (50G) 가

.

가 / growpart

가 .

OS Disk 가 /

.

cloud-utils-growpart

FlexCloud

요약

서버 관리

방화벽 그룹

스냅샷

모니터링

서버 관리:

초기 비밀번호 받기

볼륨 관리

아이피 관리

하드웨어 스펙 관리

동일하게 신규 서버 설치

소프트웨어

백업 관리

하드웨어 관리

CPU Core: 1 Core (23원/시간 또는 15000원/월)

RAM: 1024 MB (22원/시간 또는 10000원/월)

HDD 용량: 100 GB (20원/시간 또는 10000원/월)

* HDD 용량은 축소는 데이터 보호를 위해 불가능 합니다.

하드웨어 스펙 변경

요약

서버 관리

방화벽 그룹

스냅샷

모니터링

서버 관리:

초기 비밀번호 받기

볼륨 관리

아이피 관리

하드웨어 스펙 관리

동일하게 신규 서버 설치

소프트웨어

백업 관리

하드웨어 관리

CPU Core: 1 Core (23원/시간 또는 15000원/월)

RAM: 1024 MB (22원/시간 또는 10000원/월)

HDD 용량: 150 GB (30원/시간 또는 15000원/월)

* HDD 용량은 축소는 데이터 보호를 위해 불가능 합니다.

하드웨어 스펙 변경

HDD 100G -> 150G .

1. 가 / .

(> HDD 가 ,)

```
# lsblk
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
------	---------	----	------	----	------	------------

fd0	2:0	1	4K	0	disk	
-----	-----	---	----	---	------	--

sda	8:0	0	150G	0	disk	
-----	-----	---	------	---	------	--

sda1	8:1	0	20G	0	part	/
------	-----	---	-----	---	------	---

sda2	8:2	0	80G	0	part	/home
------	-----	---	-----	---	------	-------

```
# fdisk -l
```

Disk /dev/sda: 161.1 GB, 161061273600 bytes, 314572800 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 4096 bytes

I/O size (minimum/optimal): 4096 bytes / 4096 bytes

Disk label type: dos

Disk identifier: 0x000822ad

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	2048	41943039	20970496	83	Linux
/dev/sda2		41943040	209715199	83886080	83	Linux

2.

growpart /home 가 .

```
[root@localhost ~]# growpart /dev/sda 2
```

```
CHANGED: partition=2 start=41943040 old: size=167772160
end=209715200 new: size=272629727 end=314572767
```

```
[root@localhost ~]#
```

lsblk 가 .

```
[root@localhost ~]# lsblk
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
fd0	2:0	1	4K	0	disk	
sda	8:0	0	150G	0	disk	
sda1	8:1	0	20G	0	part	/
sda2	8:2	0	130G	0	part	/home

df -h /home 가 ..

```
[root@localhost ~]# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda1	20G	1.5G	19G	8%	/
devtmpfs	905M	0	905M	0%	/dev
tmpfs	915M	0	915M	0%	/dev/shm

tmpfs	915M	8.7M	906M	1%	/run
tmpfs	915M	0	915M	0%	/sys/fs/cgroup
/dev/sda2	80G	35M	80G	1%	/home

3.

blkid

```
[root@localhost ~]# blkid
```

```
/dev/sda1:      UUID="67966418-a410-4461-8aba-fb45131b4737"
TYPE="xfs"
```

```
/dev/sda2:      UUID="ab5b8818-676a-448f-a827-0686fb4e7122"
TYPE="xfs"
```

```
ext4            resize2fs
```

```
xfs             xfs_growfs
```

```
                xfs                xfs_growfs                .
```

```
[root@localhost ~]# xfs_growfs -d /home
```

```
meta-data=/dev/sda2            isize=512    agcount=4,
agsize=5242880 blks
```

```
                =                sectsz=4096   attr=2,
projid32bit=1
```

```

=                                crc=1          finobt=0
spinodes=0

data      =                      bsize=4096    blocks=20971520,
imaxpct=25

=                                sunit=0        swidth=0 blks

naming    =version 2            bsize=4096    ascii-ci=0
ftype=1

log        =internal            bsize=4096    blocks=10240,
version=2

=                                sectsz=4096   sunit=1 blks,
lazy-count=1

realtime  =none                 extsz=4096    blocks=0,
rtextents=0

data blocks changed from 20971520 to 34078715

[root@localhost ~]#

```

4. /home .

```
[root@localhost ~]# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda1	20G	1.5G	19G	8%	/
devtmpfs	905M	0	905M	0%	/dev
tmpfs	915M	0	915M	0%	/dev/shm
tmpfs	915M	8.7M	906M	1%	/run
tmpfs	915M	0	915M	0%	/sys/fs/cgroup

/dev/sda2	130G	35M	130G	1%	/home
tmpfs	183M	0	183M	0%	/run/user/0

LVM

LVM () 가 .
 (20G) 가 /home (70G) .

1.

(sdb) 가 (20G)
 가 가

```
# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
fd0                                2:0    1    4K  0 disk
sda                                8:0    0  100G  0 disk
sda1                              8:1    0    1G   0 part /boot
sda2                              8:2    0   30G  0 part /
sda3                              8:3    0   69G  0 part
vg--home-lv--home 253:0    0   69G  0 lvm  /home
sdb                                8:16   0   20G  0 disk
```

2.

```
/dev/sdb sdb1 Linux LVM .
# fdisk /dev/sdb
```

Welcome to fdisk (util-linux 2.23.2).
 Changes will remain in memory only, until you decide to write them.
 Be careful before using the write command.
 Device does not contain a recognized partition table
 Building a new DOS disklabel with disk identifier 0x49a98a24.
 The device presents a logical sector size that is smaller than

the physical sector size. Aligning to a physical sector (or optimal

I/O) size boundary is recommended, or performance may be impacted.

Command (m for help):

Command (m for help): n

Partition type:

p primary (0 primary, 0 extended, 4 free)

e extended

Select (default p): p

Partition number (1-4, default 1):

First sector (2048-41943039, default 2048):

Using default value 2048

Last sector, +sectors or +size{K,M,G} (2048-41943039, default 41943039):

Using default value 41943039

Partition 1 of type Linux and of size 20 GiB is set

Command (m for help): p

Disk /dev/sdb: 21.5 GB, 21474836480 bytes, 41943040 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 4096 bytes

I/O size (minimum/optimal): 4096 bytes / 4096 bytes

Disk label type: dos

Disk identifier: 0x49a98a24

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		2048	41943039	20970496	83	Linux

Command (m for help):

Command (m for help): t

Selected partition 1

Hex code (type L to list all codes): 8e

Changed type of partition 'Linux' to 'Linux LVM'

Command (m for help): p

Disk /dev/sdb: 21.5 GB, 21474836480 bytes, 41943040 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 4096 bytes

I/O size (minimum/optimal): 4096 bytes / 4096 bytes

Disk label type: dos

Disk identifier: 0x49a98a24

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		2048	41943039	20970496	8e	Linux LVM

Command (m for help): w

The partition table has been altered!

Calling ioctl() to re-read partition table.

Syncing disks.

3. PV

```
pvccreate                                pvdisplay                                .
[root@localhost ~]# pvccreate /dev/sdb1
Physical volume "/dev/sdb1" successfully created.
```

```
[root@localhost ~]# pvdisplay
```

--- Physical volume ---

```
PV Name          /dev/sda3
VG Name          vg-home
PV Size          <69.00 GiB / not usable 3.00 MiB
Allocatable      yes (but full)
PE Size          4.00 MiB
Total PE         17663
Free PE          0
Allocated PE     17663
PV UUID          dFUFjg-ETp6-aCkV-UHpw-HpP7-e40T-7L00p3
"/dev/sdb1" is a new physical volume of "<20.00 GiB"
```

--- NEW Physical volume ---

```
PV Name          /dev/sdb1
VG Name
PV Size          <20.00 GiB
Allocatable      NO
```

PE Size	0
Total PE	0
Free PE	0
Allocated PE	0
PV UUID	VjYTMM-ECCZ-GWrd-WpNS-0xMI-qYns-pnEr9x

4. VG

vgextend

```
# vgdisplay
```

--- Volume group ---

VG Name	vg-home
System ID	
Format	lvm2
Metadata Areas	1
Metadata Sequence No	2
VG Access	read/write
VG Status	resizable
MAX LV	0
Cur LV	1
Open LV	1
Max PV	0
Cur PV	1
Act PV	1
VG Size	<69.00 GiB
PE Size	4.00 MiB
Total PE	17663
Alloc PE / Size	17663 / <69.00 GiB
Free PE / Size	0 / 0
VG UUID	MESw3F-j0gs-yz0o-bwR9-34nn-u0dP-qH0HEC

가

```
[root@localhost ~]# vgextend vg-home /dev/sdb1
```

Volume group "vg-home" successfully extended

```
[root@localhost ~]# vgdisplay
```

```
--- Volume group ---
```

VG Name	vg-home
System ID	
Format	lvm2
Metadata Areas	2
Metadata Sequence No	3
VG Access	read/write
VG Status	resizable
MAX LV	0
Cur LV	1
Open LV	1
Max PV	0
Cur PV	2
Act PV	2
VG Size	88.99 GiB
PE Size	4.00 MiB
Total PE	22782
Alloc PE / Size	17663 / <69.00 GiB
Free PE / Size	5119 / <20.00 GiB
VG UUID	MESw3F-j0gs-yz0o-bwR9-34nn-u0dP-qH0HEC

5. LV (lvextend)

```
lvdisplay .
```

```
# lvdisplay
```

```
--- Logical volume ---
```

LV Path	/dev/vg-home/lv-home
LV Name	lv-home
VG Name	vg-home
LV UUID	AY0b7g-L0pN-DrTm-Lhgb-vqU5-T2W1-0cFe4e
LV Write Access	read/write
LV Creation host, time	localhost.localdomain, 2022-05-11 16:20:59 +0900
LV Status	available
# open	1

```

LV Size          <69.00 GiB
Current LE       17663
Segments         1
Allocation       inherit
Read ahead sectors auto
- currently set to 8192
Block device     253:0

```

6. (resize2fs)

```
[root@localhost ~]# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda2	30G	1.3G	29G	5%	/
devtmpfs	484M	0	484M	0%	/dev
tmpfs	493M	0	493M	0%	/dev/shm
tmpfs	493M	19M	474M	4%	/run
tmpfs		493M		0	493M 0%
/sys/fs/cgroup					
/dev/sda1	1014M	140M	875M	14%	/boot
/dev/mapper/vg--home-lv--home	68G	55M	65G	1%	/home

```

xfs      -> xfs_growfs
ext4     -> resize2fs

```

```
[root@localhost ~]# blkid /dev/vg-home/lv-home
```

```

/dev/vg-home/lv-home:                                UUID="c29e7a45-
dc2f-46b0-9f9a-62aa7e17ad07" TYPE="ext4"

```

```
[root@localhost ~]# resize2fs /dev/vg-home/lv-home
```

```
resize2fs 1.42.9 (28-Dec-2013)
```

```

Filesystem at /dev/vg-home/lv-home is mounted on /home; on-
line resizing required

```

```
old_desc_blocks = 9, new_desc_blocks = 12
The filesystem on /dev/vg-home/lv-home is now 23328768 blocks
long.
```

```
[root@localhost ~]# resize2fs /dev/vg-home/lv-home
```

```
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/vg-home/lv-home is mounted on /home; on-
line resizing required
old_desc_blocks = 9, new_desc_blocks = 12
The filesystem on /dev/vg-home/lv-home is now 23328768 blocks
long.
```

7.

```
# /home          가          .
```

```
[root@localhost ~]# df -h
Filesystem                                Size  Used Avail Use% Mounted
on
/dev/sda2                                30G   1.3G   29G    5% /
devtmpfs                                484M     0   484M    0% /dev
tmpfs                                    493M     0   493M    0% /dev/shm
tmpfs                                    493M    19M   474M    4% /run
tmpfs                                    493M     0   493M    0%
/sys/fs/cgroup
/dev/sda1                               1014M   140M   875M   14% /boot
/dev/mapper/vg--home-lv--home            88G     59M   84G    1% /home
```