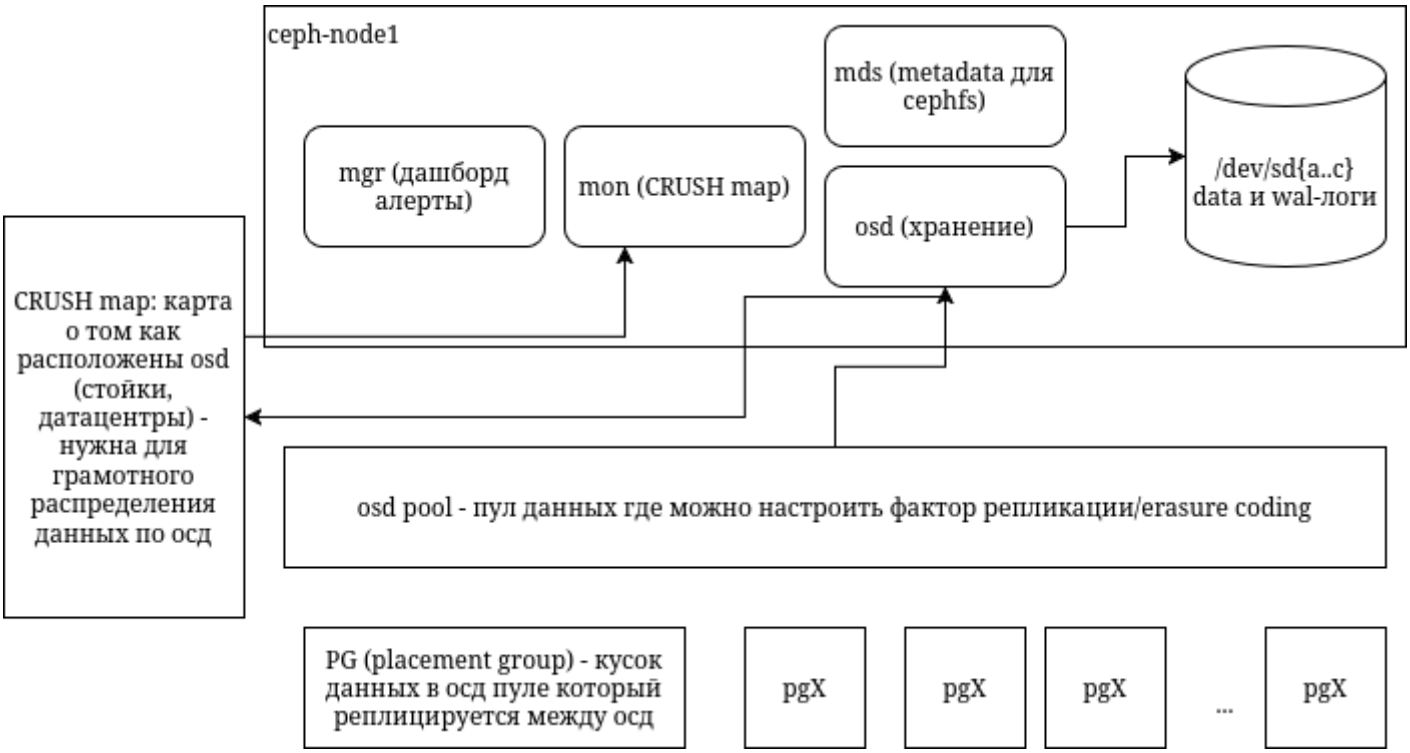
????????

- ceph
- lxd
- Пробрасываем внешний IP адрес к себе домой
- termux
- Intel, дай мне максимальную производительность
- Из socks5 в http
- syslog-ng
- Игровая виртуалка на линуксе без головы
- Яндекс Станция 2 не подключается к линуксу
- virtualbox
    - Склонировал VM (ubuntu) - тот же адрес по DHCP
- Какая у меня архитектура, x86_64, x86_64-v3, x86_64-v4?
- Пересканировать шину SATA (SCSI)
- minecraft
    - systemd unit
    - серваки
    - моды
- fio - тестирование скорости дисковой подсистемы
- debian: компоненты реп, договор
- valheim game server
- openssl check cert
- archlinux gpg nodata
- type c застрял в online
- Проклятый блютус на линуксах

ceph



Если будешь ставить через cephadm, то качай версию pacific, почему-то более новые не работают.

Удобный установщик: https://git.lulzette.ru/lulzette/ceph-installer

???????

- osd: хранение данных
- mon: условно, controlplane, управляет данными, балансирует данные, управляет кворумом, метаданные тоже тут
- mgr: These include the Ceph Dashboard and the cephadm manager module, фигня на питоне которая дает дашборд, прометеус, restful api и всякую такую некритичную для работы кластера фигню

??? ????????????? ?????????

1. Хосты с ролью OSD

2. OSDшки (абстракция на уровне Ceph'а): 1 OSD на 1 Диск хоста, т.е. 10 дисков - 10 OSD на одном хосте
3. PG (Placement Group): Группа размещения, т.е. на каких OSD размещать объекты Ceph'а (не путать с объектом на уровне S3/swift/rgw). Также есть **Карта OSD**, которая ассоциирована с PG, в ней указана одна Primary OSD и одна или несколько Replica OSD. Вся работа происходит с Primary OSD, а синк данных на реплики идет асинхронно.

???? ????? ???????????

Есть 3 разных хранилища (?):

- cephfs
- RBD (Rados Block Device)
- S3/Swift который предоставляется Rados Gateway
- Rados/librados - библиотека для твоего приложения которая может общаться с радосом минуя промежуточные RBD/S3/Swift

?????????:

- Рассинхронизация времени на хоста
- See also:

https://bogachev.biz/2017/08/23/zametki-administratora-ceph-chast-1/

CephFS

Поднять MDS - сервер метаданных

Вообще достаточно просто

```
# mon
ceph fs volume create cephfs


root@microceph:~# ceph config generate-minimal-conf
# minimal ceph.conf for 56db00e1-912c-4ac1-9d1a-1f4194c55834
[global]
```

```
        fsid = 56db00e1-912c-4ac1-9d1a-1f4194c55834
        mon_host = [v2:10.99.99.74:3300/0,v1:10.99.99.74:6789/0]

root@microceph:~# ceph fs authorize cephfs client.foo / rw
[client.foo]
        key = AQCGVs5lyBLmIxAApqSed51BlHOvQlyawvG2Uw==

# client
# может показаться что все что ниже, кроме команды монтирования, не имеет смысла, но не
root@test0:~# mkdir /etc/ceph
root@test0:~# vim /etc/ceph/ceph.conf
root@test0:~# vim /etc/ceph/ceph.client.foo.keyring
root@test0:~# chmod 644 /etc/ceph/ceph.conf
root@test0:~# chmod 600 /etc/ceph/ceph.client.foo.keyring

root@test0:~# mount -t ceph 10.99.99.74:/ /mnt/mycephfs -o
secret=AQCGVs5lyBLmIxAApqSed51BlHOvQlyawvG2Uw== -o name=foo

# fstab:
10.99.99.74:/ /mnt/cephfs ceph
name=foo,secret=AQCGVs5lyBLmIxAApqSed51BlHOvQlyawvG2Uw==,noatime,_netdev    0       2
```

+ k8s

Нам понадобится ключ для доступа к cephfs (а так же к пулу, здесь я указал админского юзера, но можно самому создать, грамотно выделив права), закинуть в наш куб CSI (Container Storage Interface) с указанными параметрами, storageclass с секретом и хранилкой можно пользоваться.

```
# в этом подтоме будут ФСки кластеров
root@microceph:~# ceph fs subvolumegroup create cephfs csi

root@node1:~/cephfs# snap install helm --classic
helm 3.14.1 from Snapcrafters✪ installed

root@node1:~/cephfs# helm repo add ceph-csi https://ceph.github.io/csi-charts
"ceph-csi" has been added to your repositories
```

```
root@node1:~/cephfs# helm inspect values ceph-csi/ceph-csi-cephfs > cephfs.yml
```

valuesы у хельм чарта:

```
---
rbac:
  # Specifies whether RBAC resources should be created
  create: true

serviceAccounts:
  nodeplugin:
    # Specifies whether a ServiceAccount should be created
    create: true
    # The name of the ServiceAccount to use.
    # If not set and create is true, a name is generated using the fullname
    name:
  provisioner:
    # Specifies whether a ServiceAccount should be created
    create: true
    # The name of the ServiceAccount to use.
    # If not set and create is true, a name is generated using the fullname
    name:

# Configuration for the CSI to connect to the cluster
# Ref: https://github.com/ceph/ceph-csi/blob/devel/examples/README.md
# Example:
csiConfig:
  - clusterID: "56db00e1-912c-4ac1-9d1a-1f4194c55834"
    monitors:
      - "10.99.99.74:6789"
#     cephFS:
#       subvolumeGroup: "csi"
#       netNamespaceFilePath: "{{ .kubeletDir }}/plugins/{{ .driverName }}/net"
#csiConfig: []

# Labels to apply to all resources
commonLabels: {}

# Set logging level for csi containers.
# Supported values from 0 to 5. 0 for general useful logs,
```

```
# 5 for trace level verbosity.
# logLevel is the variable for CSI driver containers's log level
logLevel: 5
# sidecarLogLevel is the variable for Kubernetes sidecar container's log level
sidecarLogLevel: 1

nodeplugin:
  name: nodeplugin
  # if you are using ceph-fuse client set this value to OnDelete
  updateStrategy: RollingUpdate
  podSecurityPolicy:
    enabled: true
  # set user created priorityclassName for csi plugin pods. default is
  # system-node-critical which is highest priority
  priorityClassName: system-node-critical

  httpMetrics:
    # Metrics only available for cephcsi/cephcsi => 1.2.0
    # Specifies whether http metrics should be exposed
    enabled: true
    # The port of the container to expose the metrics
    containerPort: 8091

    service:
      # Specifies whether a service should be created for the metrics
      enabled: true
      # The port to use for the service
      servicePort: 8080
      type: ClusterIP

      # Annotations for the service
      # Example:
      # annotations:
      #   prometheus.io/scrape: "true"
      #   prometheus.io/port: "9080"
      annotations: {}

      clusterIP: ""

      ## List of IP addresses at which the stats-exporter service is available
```

```
    ## Ref: https://kubernetes.io/docs/user-guide/services/#external-ips
    ##
    externalIPs: []

    loadBalancerIP: ""
    loadBalancerSourceRanges: []

## Reference to one or more secrets to be used when pulling images
##
imagePullSecrets: []
# - name: "image-pull-secret"

profiling:
  enabled: false

registrar:
  image:
    repository: registry.k8s.io/sig-storage/csi-node-driver-registrar
    tag: v2.9.1
    pullPolicy: IfNotPresent
  resources: {}

plugin:
  image:
    repository: quay.io/cephcsi/cephcsi
    tag: v3.10.2
    pullPolicy: IfNotPresent
  resources: {}

nodeSelector: {}

tolerations: []

affinity: {}

# Set to true to enable Ceph Kernel clients
# on kernel < 4.17 which support quotas
# forcecephkernelclient: true

# common mount options to apply all mounting
```

```
  # example: kernelmountoptions: "recover_session=clean"
  kernelmountoptions: ""
  fusemountoptions: ""

provisioner:
  name: provisioner
  replicaCount: 1
  podSecurityPolicy:
    enabled: true
  strategy:
    # RollingUpdate strategy replaces old pods with new ones gradually,
    # without incurring downtime.
    type: RollingUpdate
    rollingUpdate:
      # maxUnavailable is the maximum number of pods that can be
      # unavailable during the update process.
      maxUnavailable: 50%
  # Timeout for waiting for creation or deletion of a volume
  timeout: 60s
  # cluster name to set on the subvolume
  # clustername: "k8s-cluster-1"

  # set user created priorityclassName for csi provisioner pods. default is
  # system-cluster-critical which is less priority than system-node-critical
  priorityClassName: system-cluster-critical

  # enable hostnetwork for provisioner pod. default is false
  # useful for deployments where the podNetwork has no access to ceph
  enableHostNetwork: false

  httpMetrics:
    # Metrics only available for cephcsi/cephcsi => 1.2.0
    # Specifies whether http metrics should be exposed
    enabled: true
    # The port of the container to expose the metrics
    containerPort: 8081

    service:
      # Specifies whether a service should be created for the metrics
      enabled: true
```

```
    # The port to use for the service
    servicePort: 8080
    type: ClusterIP

    # Annotations for the service
    # Example:
    # annotations:
    #   prometheus.io/scrape: "true"
    #   prometheus.io/port: "9080"
    annotations: {}

    clusterIP: ""

    ## List of IP addresses at which the stats-exporter service is available
    ## Ref: https://kubernetes.io/docs/user-guide/services/#external-ips
    ##
    externalIPs: []

    loadBalancerIP: ""
    loadBalancerSourceRanges: []

## Reference to one or more secrets to be used when pulling images
##
imagePullSecrets: []
# - name: "image-pull-secret"

profiling:
  enabled: false

provisioner:
  image:
    repository: registry.k8s.io/sig-storage/csi-provisioner
    tag: v3.6.2
    pullPolicy: IfNotPresent
  resources: {}
  ## For further options, check
  ## https://github.com/kubernetes-csi/external-provisioner#command-line-options
  extraArgs: []

# set metadata on volume
```

```
  setmetadata: true

  resizer:
    name: resizer
    enabled: true
    image:
      repository: registry.k8s.io/sig-storage/csi-resizer
      tag: v1.9.2
      pullPolicy: IfNotPresent
    resources: {}
    ## For further options, check
    ## https://github.com/kubernetes-csi/external-resizer#recommended-optional-arguments
    extraArgs: []

  snapshotter:
    image:
      repository: registry.k8s.io/sig-storage/csi-snapshotter
      tag: v6.3.2
      pullPolicy: IfNotPresent
    resources: {}
    ## For further options, check
    ## https://github.com/kubernetes-csi/external-snapshotter#csi-external-snapshotter-
sidecar-command-line-options
    extraArgs: []

  nodeSelector: {}

  tolerations: []

  affinity: {}

# readAffinity:
# Enable read affinity for CephFS subvolumes. Recommended to
# set to true if running kernel 5.8 or newer.
# enabled: false
# Define which node labels to use as CRUSH location.
# This should correspond to the values set in the CRUSH map.
# NOTE: the value here serves as an example
# crushLocationLabels:
#   - topology.kubernetes.io/region
```

```
#   - topology.kubernetes.io/zone

# Mount the host /etc/selinux inside pods to support
# selinux-enabled filesystems
selinuxMount: true

storageClass:
  # Specifies whether the Storage class should be created
  create: true
  name: csi-cephfs-sc
  # Annotations for the storage class
  # Example:
  # annotations:
  #   storageclass.kubernetes.io/is-default-class: "true"
  annotations: {}

  # String representing a Ceph cluster to provision storage from.
  # Should be unique across all Ceph clusters in use for provisioning,
  # cannot be greater than 36 bytes in length, and should remain immutable for
  # the lifetime of the StorageClass in use.
  clusterID: 56db00e1-912c-4ac1-9d1a-1f4194c55834
  # (required) CephFS filesystem name into which the volume shall be created
  # eg: fsName: myfs
  fsName: cephfs
  # (optional) Ceph pool into which volume data shall be stored
  # pool: <cephfs-data-pool>
  # For eg:
  # pool: "replicapool"
  pool: "cephfs.cephfs.data"
  # (optional) Comma separated string of Ceph-fuse mount options.
  # For eg:
  # fuseMountOptions: debug
  fuseMountOptions: ""
  # (optional) Comma separated string of Cephfs kernel mount options.
  # Check man mount.ceph for mount options. For eg:
  # kernelMountOptions: readdir_max_bytes=1048576,norbytes
  kernelMountOptions: ""
  # (optional) The driver can use either ceph-fuse (fuse) or
  # ceph kernelclient (kernel).
  # If omitted, default volume mounter will be used - this is
```

```
  # determined by probing for ceph-fuse and mount.ceph
  # mounter: kernel
  mounter: ""
  # (optional) Prefix to use for naming subvolumes.
  # If omitted, defaults to "csi-vol-".
  # volumeNamePrefix: "foo-bar-"
  volumeNamePrefix: ""
  # The secrets have to contain user and/or Ceph admin credentials.
  provisionerSecret: csi-cephfs-secret
  # If the Namespaces are not specified, the secrets are assumed to
  # be in the Release namespace.
  provisionerSecretNamespace: ""
  controllerExpandSecret: csi-cephfs-secret
  controllerExpandSecretNamespace: ""
  nodeStageSecret: csi-cephfs-secret
  nodeStageSecretNamespace: ""
  reclaimPolicy: Delete
  allowVolumeExpansion: true
  mountOptions: []
  # Mount Options
  # Example:
  # mountOptions:
  #   - discard

secret:
  # Specifies whether the secret should be created
  create: true
  name: csi-cephfs-secret
  annotations: {}
  # Key values correspond to a user name and its key, as defined in the
  # ceph cluster. User ID should have required access to the 'pool'
  # specified in the storage class
  adminID: admin
  adminKey: AQDpPctl9T9ZHhAAktyT6vNlGkSE3/rfqnkxKA==

# This is a sample configmap that helps define a Ceph configuration as required
# by the CSI plugins.
# Sample ceph.conf available at
# https://github.com/ceph/ceph/blob/master/src/sample.ceph.conf Detailed
# documentation is available at
```

```
# https://docs.ceph.com/en/latest/rados/configuration/ceph-conf/
cephconf: |
  [global]
    auth_cluster_required = cephx
    auth_service_required = cephx
    auth_client_required = cephx

    # ceph-fuse which uses libfuse2 by default has write buffer size of 2KiB
    # adding 'fuse_big_writes = true' option by default to override this limit
    # see https://github.com/ceph/ceph-csi/issues/1928
    fuse_big_writes = true

# Array of extra objects to deploy with the release
extraDeploy: []

################################################################
# Variables for 'internal' use please use with caution! #
################################################################

# The filename of the provisioner socket
provisionerSocketFile: csi-provisioner.sock
# The filename of the plugin socket
pluginSocketFile: csi.sock
# kubelet working directory,can be set using `--root-dir` when starting kubelet.
kubeletDir: /var/lib/kubelet
# Name of the csi-driver
driverName: cephfs.csi.ceph.com
# Name of the configmap used for state
configMapName: ceph-csi-config
# Key to use in the Configmap if not config.json
# configMapKey:
# Use an externally provided configmap
externallyManagedConfigmap: false
# Name of the configmap used for ceph.conf
cephConfConfigMapName: ceph-config
```

применяем-проверяем

```
root@node1:~/cephfs# helm upgrade -i ceph-csi-cephfs ceph-csi/ceph-csi-cephfs -f cephfs.yml -n
ceph-csi-cephfs --create-namespace
Release "ceph-csi-cephfs" does not exist. Installing it now.
NAME: ceph-csi-cephfs
LAST DEPLOYED: Thu Feb 15 18:58:51 2024
NAMESPACE: ceph-csi-cephfs
STATUS: deployed
REVISION: 1
TEST SUITE: None
NOTES:
Examples on how to configure a storage class and start using the driver are here:
https://github.com/ceph/ceph-csi/tree/v3.10.2/examples/cephfs

#### test

root@node1:~/cephfs# kubectl apply -f cephfs-claim.yml
persistentvolumeclaim/gimme-pvc created

# ypa!
root@node1:~/cephfs# kubectl get pvc
NAME        STATUS   VOLUME                                     CAPACITY   ACCESS MODES
STORAGECLASS    VOLUMEATTRIBUTESCLASS   AGE
gimme-pvc   Bound    pvc-5d0a4e00-1ace-4b1f-83b8-900340e63999   1Gi        RWX            csi-
cephfs-sc   <unset>                 2s

root@node1:~/cephfs# cat cephfs-claim.yml
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: gimme-pvc
spec:
  accessModes:
  - ReadWriteMany
  resources:
    requests:
      storage: 1Gi
  storageClassName: csi-cephfs-sc
```

lxd

Типы инстансов: https://github.com/dustinkirkland/instance-type

Мастхэв:

```
t1.micro:
  cpu: 1.0
  mem: 0.613
t2.2xlarge:
  cpu: 8.0
  mem: 32.0
t2.large:
  cpu: 2.0
  mem: 8.0
t2.medium:
  cpu: 2.0
  mem: 4.0
t2.micro:
  cpu: 1.0
  mem: 1.0
t2.nano:
  cpu: 1.0
  mem: 0.5
t2.small:
  cpu: 1.0
  mem: 2.0
t2.xlarge:
  cpu: 4.0
  mem: 16.0
t3.2xlarge:
  cpu: 8.0
  mem: 32.0
t3.large:
  cpu: 2.0
  mem: 8.0
```

```
t3.medium:
  cpu: 2.0
  mem: 4.0
t3.micro:
  cpu: 2.0
  mem: 1.0
t3.nano:
  cpu: 2.0
  mem: 0.5
t3.small:
  cpu: 2.0
  mem: 2.0
t3.xlarge:
  cpu: 4.0
  mem: 16.0
```

Launch a VM that boots from an ISO

To launch a VM that boots from an ISO, you must first create a VM. Let’s assume that we want to create a VM and install it from the ISO image. In this scenario, use the following command to create an empty VM:

```
lxc init iso-vm --empty --vm
```

The second step is to import an ISO image that can later be attached to the VM as a storage volume:

```
lxc storage volume import <pool> <path-to-image.iso> iso-volume --type=iso
```

Lastly, you need to attach the custom ISO volume to the VM using the following command:

```
lxc config device add iso-vm iso-volume disk pool=<pool> source=iso-volume boot.priority=10
```

The `boot.priority` configuration key ensures that the VM will boot from the ISO first. Start the VM and connect to the console as there might be a menu you need to interact with:

```
lxc start iso-vm --console
```

Once you’re done in the serial console, you need to disconnect from the console using `ctrl+a-q`, and connect to the VGA console using the following command:

```
lxc console iso-vm --type=vga
```

You should now see the installer. After the installation is done, you need to detach the custom ISO volume:

```
lxc storage volume detach <pool> iso-volume iso-vm
```

Now the VM can be rebooted, and it will boot from disk.

?????????? ?????????? ???? ? ????? ? ?????????

```
lxc config device add vm1 disk1 disk source=/dev/sda
```

?????? FreeBSD/OPNSense/pfSense

Добавь в конфиг виртуалки:

```
config:
  raw.qemu: |
    -cpu host
  raw.qemu.conf: |
    [device "dev-qemu_rng"]
```

????????? ?????????? ? ????????

```
sudo snap set lxd ui.enable=true
sudo systemctl reload snap.lxd.daemon
lxc config set core.https_address :8443
lxc config trust add
```

??????? ????? ? ????????? ?? ??????

https://documentation.ubuntu.com/lxd/en/latest/reference/devices_proxy

```
lxc config device set mature-kangaroo eth0 ipv4.address=10.117.170.194
lxc config device add mature-kangaroo ssh-forward proxy listen=tcp:8.8.8.8:2222
```

```
connect=tcp:10.117.170.194:22 nat=true
```

Но можно и так

```
root@example:~# lxc network forward port add lxdbr0 IP-HOST tcp PORT-HOST IP-VM PORT-VM
```

Random

```
# images list
incus image list images: architecture=x86_64 type=virtual-machine

# error
root@huanan:~# incus launch images:almalinux/9 alma1 --vm -c limits.cpu=2 -c
limits.memory=2GiB -d root,size=20GiB
Launching alma1
Error: Failed instance creation: Failed creating instance record: Add instance info to the
database: This "instances" entry already exists
root@huanan:~# incus config device add alma1 agent disk source=agent:config
Device agent added to alma1
```

additional disks

```
lxc storage volume create hdd-storage iscsi-for-vm-export size=3000GiB --type=block

lxc config device add ftp iscsi-for-vm-export disk pool=hdd-storage source=iscsi-for-vm-export
```

road to cluster storage

```
# sdb - iscsi disk
# os - ubuntu 24
iscsiadm -m discovery -t st -p X.X.X.X
iscsiadm -m node --targetname iqn.2023-09.com.example:stor.tgt2-incus --portal X.X.X.X --login
```

```
apt install lvm2-lockd dlm-controld

vim /etc/lvm/lvm.conf:
global {
        use_lvmlockd = 1
        locking_type = 1
}
systemctl status lvmlockd.service lvmlocks.service dlm.service
vgcreate --shared vg0 /dev/sdb
vgchange --lockstart
incus storage create pool1 lvmcluster source=vg0
```

admin init

```
root@host1:~# incus admin init
Would you like to use clustering? (yes/no) [default=no]: yes
What IP address or DNS name should be used to reach this server? [default=10.146.134.217]:
Are you joining an existing cluster? (yes/no) [default=no]:
What member name should be used to identify this server in the cluster? [default=host1]:
Do you want to configure a new local storage pool? (yes/no) [default=yes]: no
Do you want to configure a new remote storage pool? (yes/no) [default=no]: yes
Create a new LVMCLUSTER pool? (yes/no) [default=yes]: yes
Name of the shared LVM volume group: vg0
Would you like to use an existing bridge or host interface? (yes/no) [default=no]: yes
Name of the existing bridge or host interface: enp5s0
Would you like stale cached images to be updated automatically? (yes/no) [default=yes]:
Would you like a YAML "init" preseed to be printed? (yes/no) [default=no]: yes
config:
  core.https_address: 10.146.134.217:8443
networks: []
storage_pools:
- config:
    source: vg0
  description: ""
  name: remote
  driver: lvmcluster
profiles:
- config: {}
  description: ""
  devices:
```

```
    eth0:
      name: eth0
      nictype: macvlan
      parent: enp5s0
      type: nic
    root:
      path: /
      pool: remote
      type: disk
  name: default
projects: []
cluster:
  server_name: host1
  enabled: true
  member_config: []
  cluster_address: ""
  cluster_certificate: ""
  server_address: ""
  cluster_token: ""
  cluster_certificate_path: ""
```

well, thats it. Теперь можно создавать тут виртуалки. Сторадж - remote.

Добавляем остальных челиков

```
# host1
incus cluster add 10.146.134.213
# host2

root@host2:~# cat boot | incus admin init --preseed
root@host2:~# cat boot
cluster:
  enabled: true
  server_address: 10.146.134.213:8443 # адрес host2, который добавляем к кластеру
  cluster_token: eyJzZXJ2ZXJf...yNzY1OVoifQ==
  member_config:
  - entity: storage-pool
    name: default
    key: source
    value: ""
```

```
  - entity: storage-pool
    name: remote
    key: source
    value: "vg0"
  - entity: storage-pool
    name: remote
    key: driver
    value: "lvmcluster"
```

кластеризация работает!

```
root@host2:~# incus ls
+------+---------+----------------------+-----------------------------------------------
+-----------+-----------+----------+
| NAME |  STATE  |         IPV4         |                      IPV6                       |
TYPE    | SNAPSHOTS | LOCATION |
+------+---------+----------------------+-----------------------------------------------
+-----------+-----------+----------+
| u2   | RUNNING | 10.146.134.93 (eth0) | fd42:6c9a:e05f:9e91:216:3eff:fe55:4480 (eth0) |
CONTAINER | 0         | host1    |
+------+---------+----------------------+-----------------------------------------------
+-----------+-----------+----------+
root@host2:~# incus shell u2
root@u2:~# hostname
u2
```

Аналогичным образом добавляем третий хост. Как я понимаю, кворума нет в incus, поэтому достаточно иметь 2+ сервера в кластере.

```
root@host1:~# incus cluster list
+----------------+-----------------------------+-----------------+-------------
+----------------+-------------+--------+-------------------+
|      NAME      |             URL             |      ROLES      | ARCHITECTURE | FAILURE
DOMAIN | DESCRIPTION | STATUS |      MESSAGE      |
+----------------+-----------------------------+-----------------+-------------
+----------------+-------------+--------+-------------------+
| 10.146.134.213 | https://10.146.134.213:8443 | database-standby | x86_64       |
default        |             | ONLINE | Fully operational |
+----------------+-----------------------------+-----------------+-------------
+----------------+-------------+--------+-------------------+
```

```
| host1          | https://10.146.134.217:8443 | database-leader  | x86_64       |
default        |             | ONLINE | Fully operational |
|                |                             | database         |
|                |             |        |                  |
+----------------+-----------------------------+------------------+--------------
+----------------+-------------+--------+------------------+
...
root@host1:~# incus cluster list
+----------------+-----------------------------+-----------------+--------------
+----------------+-------------+--------+------------------+
|      NAME      |             URL             |      ROLES      | ARCHITECTURE | FAILURE
DOMAIN | DESCRIPTION | STATUS |      MESSAGE      |
+----------------+-----------------------------+-----------------+--------------
+----------------+-------------+--------+------------------+
| 10.146.134.213 | https://10.146.134.213:8443 | database        | x86_64       |
default        |             | ONLINE | Fully operational |
+----------------+-----------------------------+-----------------+--------------
+----------------+-------------+--------+------------------+
| host1          | https://10.146.134.217:8443 | database-leader | x86_64       |
default        |             | ONLINE | Fully operational |
|                |                             | database        |
|                |             |        |                  |
+----------------+-----------------------------+-----------------+--------------
+----------------+-------------+--------+------------------+
| host3          | https://10.146.134.183:8443 | database        | x86_64       |
default        |             | ONLINE | Fully operational |
+----------------+-----------------------------+-----------------+--------------
+----------------+-------------+--------+------------------+
```

?????????? ??????? IP ????? ? ???? ?????

src: https://youtu.be/4isitiMSI_w

На интерфейсе vdsки 2 адреса.

На вдску ставим routeros, настраиваем там основной адрес. Генерим всякие серты для SSTP.

Вся суть:

Настраиваем на сервере PPP, указываем "TCP MSS"=yes, upnp=no, mpls=no, compression=yes, encryption=yes. На интерфейсе eth1 включаем proxy_arp.

L2TP

Как сделать это на чистом линухе? С помощью L2TP например.

```
root@host:~# apt install -y xl2tp

root@host:~# cat /etc/ppp/chap-secrets
# Secrets for authentication using CHAP
# client        server  secret                  IP addresses

"user1" l2tpd "parol" *

root@host:~# cat /etc/ppp/options.xl2tpd
+mschap-v2
ipcp-accept-local
ipcp-accept-remote
noccp
auth
mtu 1280
mru 1280
proxyarp
lcp-echo-failure 4
lcp-echo-interval 30
```

```
connect-delay 5000
ms-dns 8.8.8.8
ms-dns 8.8.4.4

root@host:~# cat /etc/xl2tpd/xl2tpd.conf
[global]
port = 1701

[lns default]
ip range = 194.87.56.79 # наш дополнительный адрес, который мы будем вытаскивать на своем
хосте
local ip = 194.87.56.1 # гейтвей дополнительного адреса
require chap = yes
refuse pap = yes
require authentication = yes
name = l2tpd
pppoptfile = /etc/ppp/options.xl2tpd
length bit = yes

root@host:~# bash -c "echo 1 > /proc/sys/net/ipv4/conf/all/proxy_arp"
```

И можно открывать порты, например, чтобы они стали доступны в интернете.

На хосте, где мы хотим получить внешний адрес:

```
root@netplay:~# cat /etc/xl2tpd/xl2tpd.conf
[lac myvpn]
name = username
lns = 1.2.3.4
pppoptfile = /etc/ppp/peers/myvpn.xl2tpd
ppp debug = no
root@netplay:~# cat /etc/ppp/peers/myvpn.xl2tpd
remotename myvpn
user "username"
password "password"
unit 0
nodeflate
nobsdcomp
noauth
persist
```

```
nopcomp
noaccomp
maxfail 5
debug

systemctl restart xl2tp.service
sh -c 'echo "c myvpn" > /var/run/xl2tpd/l2tp-control'

ip r add 1.2.3.4 via вашроутер
ip r add default via гейтвей_хостинга
# ...
# Вы восхитительны! Через несколько минут к вам нагрянут гости из разных стран с попытками
авторизаций по ssh.
```

termux

Конфиг с кнопочками:

```
cat .termux/termux.properties
### After making changes and saving you need to run `termux-reload-settings`
### to update the terminal.  All information here can also be found on the
### wiki: https://wiki.termux.com/wiki/Terminal_Settings

###############
# General
###############

### Allow external applications to execute arbitrary commands within Termux.
### This potentially could be a security issue, so option is disabled by
### default. Uncomment to enable.
# allow-external-apps = true

### Default working directory that will be used when launching the app.
# default-working-directory = /data/data/com.termux/files/home

### Uncomment to disable toasts shown on terminal session change.
# disable-terminal-session-change-toast = true

### Uncomment to not show soft keyboard on application start.
# hide-soft-keyboard-on-startup = true

### Uncomment to let keyboard toggle button to enable or disable software
### keyboard instead of showing/hiding it.
# soft-keyboard-toggle-behaviour = enable/disable

### Adjust terminal scrollback buffer. Max is 50000. May have negative
### impact on performance.
# terminal-transcript-rows = 2000

### Uncomment to use volume keys for adjusting volume and not for the
### extra keys functionality.
# volume-keys = volume
```

```
###############
# Fullscreen mode
###############

### Uncomment to let Termux start in full screen mode.
# fullscreen = true

### Uncomment to attempt workaround layout issues when running in
### full screen mode.
# use-fullscreen-workaround = true

###############
# Cursor
###############

### Cursor blink rate. Values 0, 100 - 2000.
# terminal-cursor-blink-rate = 0

### Cursor style: block, bar, underline.
# terminal-cursor-style = block

###############
# Extra keys
###############

### Settings for choosing which set of symbols to use for illustrating keys.
### Choose between default, arrows-only, arrows-all, all and none
# extra-keys-style = default

### Force capitalize all text in extra keys row button labels.
# extra-keys-text-all-caps = true

### Default extra-key configuration
# extra-keys = [[ESC, TAB, CTRL, ALT, {key: '-', popup: '|'}, DOWN, UP]]

### Two rows with more keys
# extra-keys = [['ESC','/','-','HOME','UP','END','PGUP'], \
#               ['TAB','CTRL','ALT','LEFT','DOWN','RIGHT','PGDN']]
```

```
### Configuration with additional popup keys (swipe up from an extra key)
# extra-keys = [[ \
#   {key: ESC, popup: {macro: "CTRL f d", display: "tmux exit"}}, \
#   {key: CTRL, popup: {macro: "CTRL f BKSP", display: "tmux ←"}}, \
#   {key: ALT, popup: {macro: "CTRL f TAB", display: "tmux →"}}, \
#   {key: TAB, popup: {macro: "ALT a", display: A-a}}, \
#   {key: LEFT, popup: HOME}, \
#   {key: DOWN, popup: PGDN}, \
#   {key: UP, popup: PGUP}, \
#   {key: RIGHT, popup: END}, \
#   {macro: "ALT j", display: A-j, popup: {macro: "ALT g", display: A-g}}, \
#   {key: KEYBOARD, popup: {macro: "CTRL d", display: exit}} \
# ]]

###############
# Colors/themes
###############

### Force black colors for drawer and dialogs
# use-black-ui = true

###############
# HW keyboard shortcuts
###############

### Disable hardware keyboard shortcuts.
# disable-hardware-keyboard-shortcuts = true

### Open a new terminal with ctrl + t (volume down + t)
# shortcut.create-session = ctrl + t

### Go one session down with (for example) ctrl + 2
# shortcut.next-session = ctrl + 2

### Go one session up with (for example) ctrl + 1
# shortcut.previous-session = ctrl + 1

### Rename a session with (for example) ctrl + n
# shortcut.rename-session = ctrl + n
```

```
###############
# Bell key
###############

### Vibrate device (default).
# bell-character = vibrate

### Beep with a sound.
# bell-character = beep

### Ignore bell character.
# bell-character = ignore

###############
# Back key
###############

### Send the Escape key.
# back-key=escape

### Hide keyboard or leave app (default).
# back-key=back

###############
# Keyboard issue workarounds
###############

### Letters might not appear until enter is pressed on Samsung devices
# enforce-char-based-input = true

### ctrl+space (for marking text in emacs) does not work on some devices
# ctrl-space-workaround = true
~ $
```

кнопочки в 2 ряда:

```
### Two rows with more keys
extra-keys = [['ESC','/','-','HOME','UP','END','PGUP', 'INS'], \
              ['TAB','CTRL','ALT','LEFT','DOWN','RIGHT','PGDN', 'DEL']]
```

Intel, ??? ??? ??????????? ???????????????????

Если надо максимальную производительность то никакие performance говернеры не помогут и даже это

```
echo 0 | tee /sys/devices/system/cpu/cpu*/power/energy_perf_bias
```

отруби энергосбережение в биосе

?? socks5 ? http

```
$ gost -L=http://:8080 -F=socks5://127.0.0.1:1080
```

Предположим у тебя есть socks5, но тот же терраформ его не жрет. Используй эту тулзу, а потом прокинь ENVы http_proxy, https_proxy

syslog-ng

Как отфильтровать логи конкретного systemd юнита:

```
filter nginx_service {
  "${.journald._SYSTEMD_UNIT}" eq "nginx.service";
};
```

Другие поля можно узнать посмотрев на лог:

```
journalctl --output json-pretty


{
        "_SYSTEMD_UNIT" : "init.scope",
        "_PID" : "1",
        "_MACHINE_ID" : "XXXXX",
        "INVOCATION_ID" : "ecad364ce7794e13ae6bc35a59ae4ac2",
        "CODE_LINE" : "574",
        "__MONOTONIC_TIMESTAMP" : "3003939058929",
        "_TRANSPORT" : "journal",
        "_UID" : "0",
        "PRIORITY" : "6",
        "_EXE" : "/usr/lib/systemd/systemd",
        "_CAP_EFFECTIVE" : "1fcfdfcffff",
        "_SYSTEMD_CGROUP" : "/init.scope",
        "SYSLOG_FACILITY" : "3",
        "_COMM" : "systemd",
        "MESSAGE" : "Starting The PHP 8.1 FastCGI Process Manager...",
        "CODE_FILE" : "src/core/job.c",
        "_SOURCE_REALTIME_TIMESTAMP" : "1694843970044480",
        "SYSLOG_IDENTIFIER" : "systemd",
        "_CMDLINE" : "/lib/systemd/systemd --system --deserialize 33",
        "_HOSTNAME" : "web",
        "JOB_ID" : "300568",
        "JOB_TYPE" : "start",
        "MESSAGE_ID" : "7d4958e842da4a758f6c1cdc7b36dcc5",
        "__REALTIME_TIMESTAMP" : "1694843970044562",
        "_GID" : "0",
```

```
        "_BOOT_ID" : "04f27e01022c4fe5be6723ddae2991be",
        "UNIT" : "php8.1-fpm.service",
        "_SELINUX_CONTEXT" : "lxc-container-default-cgns (enforce)\n",
        "_SYSTEMD_SLICE" : "-.slice",
        "__CURSOR" :
"s=bf41ba4a89ec44e882c3d8729a337787;i=db8b;b=04f27e01022c4fe5be6723ddae2991be;m=2bb68b874f1;t=
605739cc8c292;x=fe2cf83fe8762ab3",
        "CODE_FUNC" : "job_log_begin_status_message"
}
```

??????? ?????????? ?? ??????? ??? ??????

Если у вас есть сервер с Nvidia RTX A2000 на каком-нибудь проксмоксе, с виртуалкой на линуксе, в которую проброшена видеокарта, могут возникнуть трудности с выводом изображения, т.к. видеокарта без воткнутого dummy адаптера или монитора будет перманентно выводить изображение в 640x480.

```
#xorg.conf:

Section "ServerLayout"
    Identifier     "Default Layout"
    Screen         "Default Screen" 0 0
    InputDevice    "Keyboard0" "CoreKeyboard"
    InputDevice    "Mouse0" "CorePointer"
EndSection

Section "InputDevice"
    # generated from default
    Identifier     "Keyboard0"
    Driver         "kbd"
EndSection

Section "InputDevice"
    # generated from default
    Identifier     "Mouse0"
    Driver         "mouse"
    Option         "Protocol" "auto"
    Option         "Device" "/dev/psaux"
    Option         "Emulate3Buttons" "no"
    Option         "ZAxisMapping" "4 5"
EndSection

Section "Device"
    Identifier     "Device0"
    Driver         "nvidia"
    VendorName     "NVIDIA Corporation"
```

```
    BusID          "PCI:1:0:0"
EndSection

Section "Screen"
    Identifier     "Default Screen"
    Device         "Device0"
    DefaultDepth    24
    Option         "CustomEDID" "GPU-0.DP-0:/home/6DB9E3ADB0D2" # надыбай где-нибудь EDID в
текстовом формате
    Option         "ConnectedMonitor" "DP-0"
    SubSection     "Display"
        Depth       24
        Modes      "nvidia-auto-select"
    EndSubSection
EndSection



Section "Monitor"
    Identifier    "Configured Monitor"
    HorizSync       30.0-62.0
    VertRefresh     50.0-70.0
EndSection
```

Я это провернул на Endevour с проприетарными дровами. Однако почему-то в системе максимальное разрешение при этом - 1600x900. Но через nvidia-xconfig можно задать разрешение 1920x1080, хоть оно и не будет корректно работать в системе (панелька будет растянута на разрешение 900p, но разрешение будет 1080p), но на работу игр это не повлияет.

Советую после установки сразу поставить ssh сервер и x11vnc. VNC сервер можно запускать вот так:

```
[usr@usr-standardpc ~]$ sudo XAUTHORITY=/run/sddm/$(sudo ls /run/sddm/) DISPLAY=:0 x11vnc -
forever
```

После пробросить порт через SSH и подключиться.

?????? ??????? 2 ?? ???????????? ? ???????

Яндекс станция 2 почему-то криво работает с линуксом, когда адаптер в режиме dual (т.е. поддерживает и обычный bredr и Bluetooth LE), решить это можно поставив адаптер в режим bredr:

```
λ grep dual /etc/bluetooth/main.conf
ControllerMode = bredr
λ systemctl restart bluetooth.service
```

Но есть нюанс - скорее всего мыши и клавиатуры используют как раз Bluetooth LE, поэтому они работать перестанут. Либо можно купить другую яндекс станцию, проблемная кажется только вторая.

virtualbox

virtualbox

?????????? VM (ubuntu) - ??? ?? ????? ?? DHCP

Если ты склонировал виртуалку в виртуалбоксе с генерацией MAC адресов сетевух, но адрес выдается все тот же - затри machine-id

```
echo '' | sudo tee /etc/machine-id
sudo rm /var/lib/dbus/machine-id
# Это самый верный способ правильно затереть machine id, генерация через dbus тоже сработает, но это не совсем верно
```

????? ? ???? ??????????, x86_64, x86_64-v3, x86_64-v4?

```
in ~ λ /lib/ld-linux-x86-64.so.2 --help | grep supported
 x86-64-v3 (supported, searched)
 x86-64-v2 (supported, searched)
```

Есть несколько подвидов x86_64, отличаются они набором инструкций. X86_64-v3 сейчас самая распространенная.

https://hackaday.com/2024/02/25/what-is-x86-64-v3/

Redhat в последних версиях собирает ядро под v3, поэтому на простеньких виртуальных процах или старых системах рхел уже не запустить.

?????????????? ???? SATA (SCSI)

Добавил новый диск а он не виден? Пересканируй шину

```
echo "- - -" > /sys/class/scsi_host/hostX/scan
# Либо все сразу
echo "- - -" | tee /sys/class/scsi_host/*/scan
```

Если контроллер глупый, либо в биосе не включена поддержка hot-plug, то никакого тебе hotplug, иди ребутайся

minecraft

minecraft

systemd unit

```
root@v1:~# cat /etc/systemd/system/minecraft@.service
[Unit]
Description=Minecraft server %I
After=local-fs.target network.target

[Service]
WorkingDirectory=/home/ubuntu/%i
User=ubuntu
Group=ubuntu
Type=forking
# Run it as a non-root user in a specific directory

ExecStart=/usr/bin/screen -h 1024 -dmS minecraft ./minecraft_server.sh
# I like to keep my commandline to launch it in a separate file
# because sometimes I want to change it or launch it manually
# If it's in the WorkingDirectory, then we can use a relative path

# Send "stop" to the Minecraft server console
ExecStop=/usr/bin/screen -p 0 -S minecraft -X eval 'stuff \"stop\"\015'
# Wait for the PID to die - otherwise it's killed after this command finishes!
ExecStop=/bin/bash -c "while ps -p $MAINPID > /dev/null; do /bin/sleep 1; done"
# Note that absolute paths for all executables are required!

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

minecraft

???????

Paper: https://papermc.io/

Purpur: https://github.com/PurpurMC/Purpur https://purpurmc.org/

Оптимизация: https://github.com/YouHaveTrouble/minecraft-optimization

Forge+Bukkit: https://github.com/IzzelAliz/Arclight

Остальные форки: https://rubukkit.org/threads/nemnogo-o-forkax-paper-i-ne-tolko.176425/

Почему не стоит использовать серваки с плагинами и модами: https://rubukkit.org/threads/ne-ispolzujte-magma-mohist-catserver-i-podobnye-jadra.185662/

minecraft

????

https://4mforyou.com/blog/forge-vs-fabric-vs-quilt-kakoj-mod-zagruzchik-minecraft-vybrat/

fio - ?????????? ???????? ???????? ??????????

Тут неплохой список тестов: https://docs.oracle.com/en-us/iaas/Content/Block/References/samplefiocommandslinux.htm

Команды:

```
#IOPS

fio --filename=/file --size=5GB --direct=1 --rw=randrw --bs=4k \
--ioengine=libaio --iodepth=256 --runtime=120 --numjobs=4 --time_based \
--group_reporting --name=iops-test-job --eta-newline=1

# B/W
fio --filename=/file --size=5GB --direct=1 --rw=randrw --bs=64k \
--ioengine=libaio --iodepth=64 --runtime=120 --numjobs=4 --time_based \
--group_reporting --name=throughput-test-job --eta-newline=1

# Latency
fio --filename=/file --size=5GB --direct=1 --rw=randrw --bs=4k \
--ioengine=libaio --iodepth=1 --numjobs=1 --time_based --group_reporting \
--name=rwlatency-test-job --runtime=120 --eta-newline=1
```

Вот пример random r/w iops на десктопном Toshiba HDD:

```
iops-test-job: (groupid=0, jobs=4): err= 0: pid=251421: Wed Jul 17 10:37:16 2024
  read: IOPS=76, BW=307KiB/s (314kB/s)(36.1MiB/120567msec)
    slat (usec): min=3, max=1056.9k, avg=25388.41, stdev=70969.08
    clat (msec): min=285, max=10659, avg=6440.46, stdev=987.92
     lat (msec): min=500, max=10763, avg=6465.85, stdev=990.85
    clat percentiles (msec):
     |  1.00th=[ 1653],  5.00th=[ 5671], 10.00th=[ 6007], 20.00th=[ 6141],
```

```
     | 30.00th=[ 6275], 40.00th=[ 6342], 50.00th=[ 6477], 60.00th=[ 6544],
     | 70.00th=[ 6678], 80.00th=[ 6812], 90.00th=[ 7148], 95.00th=[ 7819],
     | 99.00th=[ 8926], 99.50th=[ 9060], 99.90th=[ 9866], 99.95th=[10268],
     | 99.99th=[10671]
   bw (  KiB/s): min=   40, max=  608, per=100.00%, avg=307.70, stdev=26.41, samples=907
   iops        : min=   10, max=  152, avg=76.92, stdev= 6.60, samples=907
  write: IOPS=79, BW=319KiB/s (326kB/s)(37.5MiB/120567msec); 0 zone resets
    slat (usec): min=3, max=975286, avg=25577.27, stdev=70511.46
    clat (msec): min=469, max=9979, avg=6323.55, stdev=931.04
     lat (msec): min=500, max=9979, avg=6349.13, stdev=933.87
    clat percentiles (msec):
     |  1.00th=[ 1569],  5.00th=[ 5604], 10.00th=[ 5940], 20.00th=[ 6074],
     | 30.00th=[ 6208], 40.00th=[ 6275], 50.00th=[ 6342], 60.00th=[ 6409],
     | 70.00th=[ 6544], 80.00th=[ 6611], 90.00th=[ 6879], 95.00th=[ 7684],
     | 99.00th=[ 8658], 99.50th=[ 8792], 99.90th=[ 9194], 99.95th=[ 9597],
     | 99.99th=[10000]
   bw (  KiB/s): min=   64, max=  632, per=100.00%, avg=320.56, stdev=26.98, samples=909
   iops        : min=   16, max=  158, avg=80.14, stdev= 6.74, samples=909
  lat (msec)   : 500=0.02%, 750=0.25%, 1000=0.26%, 2000=0.76%, >=2000=98.71%
  cpu          : usr=0.03%, sys=0.04%, ctx=11061, majf=0, minf=43
  IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.2%, 16=0.3%, 32=0.7%, >=64=98.7%
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.1%
     issued rwts: total=9243,9607,0,0 short=0,0,0,0 dropped=0,0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=256

Run status group 0 (all jobs):
   READ: bw=307KiB/s (314kB/s), 307KiB/s-307KiB/s (314kB/s-314kB/s), io=36.1MiB (37.9MB),
run=120567-120567msec
  WRITE: bw=319KiB/s (326kB/s), 319KiB/s-319KiB/s (326kB/s-326kB/s), io=37.5MiB (39.3MB),
run=120567-120567msec

Disk stats (read/write):
  sda: ios=9257/9607, merge=0/0, ticks=4143068/3114506, in_queue=7257575, util=99.37%
```

Блоками по 512 картина еще унылее.

На 8к вцелом плюс-минус то же

```
iops-test-job: (groupid=0, jobs=4): err= 0: pid=251505: Wed Jul 17 10:42:18 2024
  read: IOPS=77, BW=620KiB/s (635kB/s)(73.0MiB/120481msec)
    slat (usec): min=3, max=387273, avg=25978.96, stdev=69175.81
    clat (msec): min=320, max=8886, avg=6372.20, stdev=876.49
     lat (msec): min=486, max=8886, avg=6398.18, stdev=878.93
    clat percentiles (msec):
     |  1.00th=[ 1552],  5.00th=[ 5738], 10.00th=[ 6007], 20.00th=[ 6208],
     | 30.00th=[ 6275], 40.00th=[ 6409], 50.00th=[ 6477], 60.00th=[ 6544],
     | 70.00th=[ 6678], 80.00th=[ 6812], 90.00th=[ 7013], 95.00th=[ 7215],
     | 99.00th=[ 7684], 99.50th=[ 7819], 99.90th=[ 8288], 99.95th=[ 8423],
     | 99.99th=[ 8926]
   bw (  KiB/s): min=  176, max= 1344, per=99.33%, avg=616.50, stdev=53.06, samples=915
   iops        : min=   22, max=  168, avg=77.06, stdev= 6.63, samples=915
  write: IOPS=80, BW=644KiB/s (660kB/s)(75.8MiB/120481msec); 0 zone resets
    slat (usec): min=3, max=417654, avg=24495.72, stdev=67714.31
    clat (msec): min=305, max=7858, avg=6266.56, stdev=797.47
     lat (msec): min=486, max=7903, avg=6291.05, stdev=799.85
    clat percentiles (msec):
     |  1.00th=[ 1670],  5.00th=[ 5738], 10.00th=[ 5940], 20.00th=[ 6141],
     | 30.00th=[ 6208], 40.00th=[ 6275], 50.00th=[ 6409], 60.00th=[ 6477],
     | 70.00th=[ 6544], 80.00th=[ 6611], 90.00th=[ 6745], 95.00th=[ 6879],
     | 99.00th=[ 7215], 99.50th=[ 7416], 99.90th=[ 7752], 99.95th=[ 7752],
     | 99.99th=[ 7886]
   bw (  KiB/s): min=   96, max= 1264, per=99.48%, avg=641.88, stdev=54.29, samples=918
   iops        : min=   12, max=  158, avg=80.24, stdev= 6.79, samples=918
  lat (msec)   : 500=0.06%, 750=0.25%, 1000=0.17%, 2000=0.88%, >=2000=98.63%
  cpu          : usr=0.03%, sys=0.04%, ctx=11609, majf=0, minf=44
  IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.2%, 16=0.3%, 32=0.7%, >=64=98.7%
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.1%
     issued rwts: total=9340,9704,0,0 short=0,0,0,0 dropped=0,0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=256

Run status group 0 (all jobs):
   READ: bw=620KiB/s (635kB/s), 620KiB/s-620KiB/s (635kB/s-635kB/s), io=73.0MiB (76.5MB),
run=120481-120481msec
```

```
  WRITE: bw=644KiB/s (660kB/s), 644KiB/s-644KiB/s (660kB/s-660kB/s), io=75.8MiB (79.5MB),
run=120481-120481msec

Disk stats (read/write):
  sda: ios=9328/9704, merge=0/0, ticks=4175537/3088888, in_queue=7264425, util=99.67%
```

Иопсы те же, но скорость выросла.

WD RED на 5400 оборотов. Несмотря на то что он на 5400, скорость выше чем у 7200 Toshiba HDD

```
iops-test-job: (groupid=0, jobs=4): err= 0: pid=251543: Wed Jul 17 10:46:32 2024
  read: IOPS=86, BW=348KiB/s (356kB/s)(41.1MiB/120812msec)
    slat (usec): min=2, max=652947, avg=22065.33, stdev=65152.95
    clat (msec): min=411, max=8507, avg=5712.90, stdev=758.82
     lat (msec): min=699, max=8507, avg=5734.97, stdev=760.79
    clat percentiles (msec):
     |  1.00th=[ 2106],  5.00th=[ 4799], 10.00th=[ 5134], 20.00th=[ 5336],
     | 30.00th=[ 5537], 40.00th=[ 5604], 50.00th=[ 5738], 60.00th=[ 5873],
     | 70.00th=[ 6007], 80.00th=[ 6208], 90.00th=[ 6477], 95.00th=[ 6678],
     | 99.00th=[ 7215], 99.50th=[ 7416], 99.90th=[ 7819], 99.95th=[ 7886],
     | 99.99th=[ 8154]
   bw (  KiB/s): min=   56, max=  784, per=100.00%, avg=350.09, stdev=32.50, samples=915
   iops        : min=   14, max=  196, avg=87.52, stdev= 8.12, samples=915
  write: IOPS=90, BW=362KiB/s (371kB/s)(42.7MiB/120812msec); 0 zone resets
    slat (usec): min=2, max=612526, avg=22706.93, stdev=65998.81
    clat (msec): min=542, max=7598, avg=5518.32, stdev=699.72
     lat (msec): min=699, max=7598, avg=5541.03, stdev=702.26
    clat percentiles (msec):
     |  1.00th=[ 1989],  5.00th=[ 4665], 10.00th=[ 5000], 20.00th=[ 5201],
     | 30.00th=[ 5336], 40.00th=[ 5470], 50.00th=[ 5604], 60.00th=[ 5671],
     | 70.00th=[ 5805], 80.00th=[ 5940], 90.00th=[ 6141], 95.00th=[ 6342],
     | 99.00th=[ 6678], 99.50th=[ 6812], 99.90th=[ 7215], 99.95th=[ 7282],
     | 99.99th=[ 7483]
   bw (  KiB/s): min=   56, max=  848, per=100.00%, avg=364.32, stdev=33.56, samples=914
   iops        : min=   14, max=  212, avg=91.08, stdev= 8.39, samples=914
```

```
  lat (msec)   : 500=0.01%, 750=0.05%, 1000=0.08%, 2000=0.85%, >=2000=99.01%
  cpu          : usr=0.04%, sys=0.04%, ctx=13052, majf=0, minf=47
  IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.3%, 32=0.6%, >=64=98.8%
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.1%
     issued rwts: total=10509,10929,0,0 short=0,0,0,0 dropped=0,0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=256

Run status group 0 (all jobs):
   READ: bw=348KiB/s (356kB/s), 348KiB/s-348KiB/s (356kB/s-356kB/s), io=41.1MiB (43.0MB),
run=120812-120812msec
  WRITE: bw=362KiB/s (371kB/s), 362KiB/s-362KiB/s (371kB/s-371kB/s), io=42.7MiB (44.8MB),
run=120812-120812msec

Disk stats (read/write):
  sdc: ios=10532/10929, merge=0/0, ticks=4648821/2648107, in_queue=7296927, util=99.67%
```

Seagate Barracuda 7200.14

```
iops-test-job: (groupid=0, jobs=4): err= 0: pid=251590: Wed Jul 17 10:49:41 2024
  read: IOPS=27, BW=111KiB/s (114kB/s)(13.1MiB/120986msec)
    slat (usec): min=3, max=1103.8k, avg=72458.05, stdev=194381.62
    clat (msec): min=488, max=20448, avg=17135.15, stdev=3593.05
     lat (msec): min=1045, max=20537, avg=17207.61, stdev=3594.34
    clat percentiles (msec):
     |  1.00th=[ 2333],  5.00th=[ 7215], 10.00th=[13758], 20.00th=[17113],
     | 30.00th=[17113], 40.00th=[17113], 50.00th=[17113], 60.00th=[17113],
     | 70.00th=[17113], 80.00th=[17113], 90.00th=[17113], 95.00th=[17113],
     | 99.00th=[17113], 99.50th=[17113], 99.90th=[17113], 99.95th=[17113],
     | 99.99th=[17113]
   bw (  KiB/s): min=   32, max=  336, per=100.00%, avg=131.32, stdev=12.72, samples=697
   iops        : min=    8, max=   84, avg=32.83, stdev= 3.18, samples=697
  write: IOPS=29, BW=116KiB/s (119kB/s)(13.7MiB/120986msec); 0 zone resets
    slat (usec): min=3, max=1236.0k, avg=67783.98, stdev=189096.67
    clat (msec): min=471, max=22212, avg=16727.15, stdev=3523.66
```

```
     lat (msec): min=1045, max=22508, avg=16794.94, stdev=3528.58
    clat percentiles (msec):
     |  1.00th=[ 2265],  5.00th=[ 7215], 10.00th=[13221], 20.00th=[16979],
     | 30.00th=[17113], 40.00th=[17113], 50.00th=[17113], 60.00th=[17113],
     | 70.00th=[17113], 80.00th=[17113], 90.00th=[17113], 95.00th=[17113],
     | 99.00th=[17113], 99.50th=[17113], 99.90th=[17113], 99.95th=[17113],
     | 99.99th=[17113]
   bw (  KiB/s): min=   32, max=  368, per=100.00%, avg=135.84, stdev=13.41, samples=704
   iops        : min=    8, max=   92, avg=33.96, stdev= 3.35, samples=704
  lat (msec)   : 500=0.03%, 1000=0.03%, 2000=0.82%, >=2000=99.13%
  cpu          : usr=0.02%, sys=0.02%, ctx=3369, majf=0, minf=44
  IO depths    : 1=0.1%, 2=0.1%, 4=0.2%, 8=0.5%, 16=0.9%, 32=1.9%, >=64=96.3%
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
     complete  : 0=0.0%, 4=99.9%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.1%
     issued rwts: total=3358,3509,0,0 short=0,0,0,0 dropped=0,0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=256

Run status group 0 (all jobs):
   READ: bw=111KiB/s (114kB/s), 111KiB/s-111KiB/s (114kB/s-114kB/s), io=13.1MiB (13.8MB),
run=120986-120986msec
  WRITE: bw=116KiB/s (119kB/s), 116KiB/s-116KiB/s (119kB/s-119kB/s), io=13.7MiB (14.4MB),
run=120986-120986msec

Disk stats (read/write):
  sdd: ios=3373/3509, merge=0/0, ticks=4327686/2960324, in_queue=7288009, util=99.87%
```

debian: ????????? ???, ???????

У дебиана есть несколько компонентов репозиториев - main, contrib, non-free.

main - тут все понятно

contrib - DFSG-совместимое, но имеющее зависимости не из main ПО

non-free - не совместимое с DFSG ПО.

DFSG это общественный договор дебиана. Подробнее про него - тут

А про репы - тут

valheim game server

debian 12

bash:

```
dpkg --add-architecture i386
vim /etc/apt/sources.list.d/ # добавляем non-free
apt update
apt install steamcmd tmux htop
useradd -m game
chsh -s /bin/bash game
sudo -u game -i
tmux
```

steamcmd:

```
force_install_dir /home/game/valheim
login anonymous
app_update 896660
quit
```

снова bash

```
cd  /home/game/valheim
cp start_server.sh start_server-my.sh
vim start_server-my.sh # меняем name, password, возможно crossplay стоит убрать, мне помогло
```

openssl check cert

```
openssl x509 -in mycert.pem -text -noout


openssl s_client -servername lulzette.ru -connect lulzette.ru:443 2>/dev/null| openssl x509 -text
```

archlinux gpg nodata

```
gpg: error reading key: Нет открытого ключа
pub   rsa4096 2011-09-23 [SC]
      647F28654894E3BD457199BE38DBBDC86092693E
uid         [ неизвестно ] Greg Kroah-Hartman <gregkh@linuxfoundation.org>
uid         [ неизвестно ] Greg Kroah-Hartman <gregkh@kernel.org>
uid         [ неизвестно ] Greg Kroah-Hartman (Linux kernel stable release signing key)
<greg@kroah.com>
sub   rsa4096 2011-09-23 [E]



 :: Ключи PGP, требующие импорта:
 -> ABAF11C65A2970B130ABE3C479BE3E4300411886, требуется пакету: linux-clear
:: Импортировать? [Y/n]
:: Импортирование ключей с помощью gpg...
gpg: сбой при получении с сервера ключей: Нет данных
 -> проблема импортирования ключей

# sudo pacman-key -r тебе не поможет, чувак

# Найди какой-нибудь кейсервер и стащи ключ оттуда, например
gpg --keyserver hkp://pgp.rediris.es --recv-keys "ABAF11C65A2970B130ABE3C479BE3E4300411886"
```

type c ??????? ? online

У меня бывало такое, что зарядку вытащил, а ноутбук продолжает думать что она подключена. Мелочь, но неприятно.

```
~ λ upower -d
...

Device: /org/freedesktop/UPower/devices/line_power_ucsi_source_psy_USBC000o001
 native-path:          ucsi-source-psy-USBC000:001
 power supply:         yes
 updated:              Вт 22 окт 2024 20:23:19 (110 seconds ago)
 has history:          no
 has statistics:       no
 line-power
   warning-level:       none
   online:              yes # <<<<<<<
   icon-name:          'ac-adapter-symbolic'
```

Решается в одну команду (возможно надо подключить потребителя в тупе-с

```
echo source |sudo tee /sys/class/typec/port0/power_role
```

????????? ?????? ?? ????????

Упарываюсь по поводу диагностики работы блютуса. У меня usb блютус 0cf3:e007 (qualcomm) встроенный в сетевку Qualcomm Atheros QCA6174 802.11ac Wireless Network Adapter (rev 32).

Мои CMF Buds pro постоянно рвутся при подключении. При этом у Soundcore Life Dot 2 ситуация лучше, практически идеальная. CMF Buds Pro можно попробовать переподключить раза 3 (не спаривать, а подключать), играться с кодеками, и в какой-то момент они будут работать нормально.

Если честно без понятия куда копать.

Запустил с дебагом блютусДэ, в логах

```
bluetoothd[223642]: profiles/audio/avdtp.c:session_cb()
bluetoothd[223642]: profiles/audio/avdtp.c:avdtp_parse_cmd() Received DELAY_REPORT_CMD
bluetoothd[223642]: profiles/audio/a2dp.c:endpoint_delayreport_ind() Source 0x5d5b9c8aae30:
DelayReport_Ind
```

С саундкорами такого нет. На условном интел вайфае тоже все отлично. При этом Qualcomm кажется официально поддерживает линь.