Managing Storage Using Partitions, LVM, and Stratis

By: Sean Twiehaus

About Me

System Administrator for several years

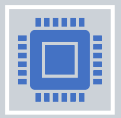

Saw the transition to Cloud and DevOps

‘Noped’ out to Software Development

Current Roles:

Tech Lead Software Engineer at Southwest Airlines
Adjunct Faculty at St. Charles Community College

This content was only tested on **Red Hat Enterprise Linux 9.3**

Partitions

Split Hard Drive into sections

Pro

- Isolation
- Prevent partition with root (/) from filling up!

Con

- Inflexible
- Must plan ahead!

Two Partition Types

MBR – MASTER BOOT RECORD

GPT – GUID PARTITION TABLE

Master Boot Record (MBR)

MBR is an older partitioning scheme

Total disk size is limited to 2 TiB

Single point of failure

Used with Basic Input/Output System (BIOS)

Only four “primary” partitions.

An extended partition allows for more partitions to be created.

GUID Partition Table (GPT)

- More Modern than MBR
- Up to 128 partitions
- Up to 8 ZiB drives
- GUID - Globally Unique Identifier (128-bit)
- Two copies of GPT for redundancy
- Compatible with UEFI firmware
  - Unified Extensible Firmware Interface

Disk Size Units

Symbol	Name	Value	Symbol	Name	Value
KB	Kilobyte	1000^1	KiB	Kibibyte	1024^1
MB	Megabyte	1000^2	MiB	Mebibyte	1024^2
GB	Gigabyte	1000^3	GiB	Gibibyte	1024^3
TB	Terabyte	1000^4	TiB	Tebibyte	1024^4
PB	Petabyte	1000^5	PiB	Pebibyte	1024^5
EB	Exabyte	1000^6	EiB	Exbibyte	1024^6
ZB	Zettabyte	1000^7	ZiB	Zebibyte	1024^7
YB	Yottabyte	1000^8	YiB	Yobibyte	1024^8

Common Device Names

Device Name	Description
/dev/sda	A hard disk that uses the SCSI driver. Used for SCSI and SATA disk devices. Common on physical servers but also in VMware virtual machines.
/dev/nvme0n1	The first hard disk on an NVM Express (NVMe) interface. NVMe is a server-grade method to address advanced SSD devices. Note at the end of the device name that the first disk in this case is referred to as *n1* instead of *a* (as is common with the other types).
/dev/hda	The (legacy) IDE disk device type. You will seldom see this device type on modern computers.
/dev/vda	A disk in a KVM virtual machine that uses the virtio disk driver. This is the common disk device type for KVM virtual machines.
/dev/xvda	A disk in a Xen virtual machine that uses the Xen virtual disk driver. You see this when installing RHEL as a virtual machine in Xen virtualization. RHEL 9 cannot be used as a Xen hypervisor, but you might see RHEL 9 virtual machines on top of the Xen hypervisor using these disk types.

Partitioning Components

- 'lsblk' command
- 'parted’ utility
- Storage Devices '/dev/sd[a-z]*'

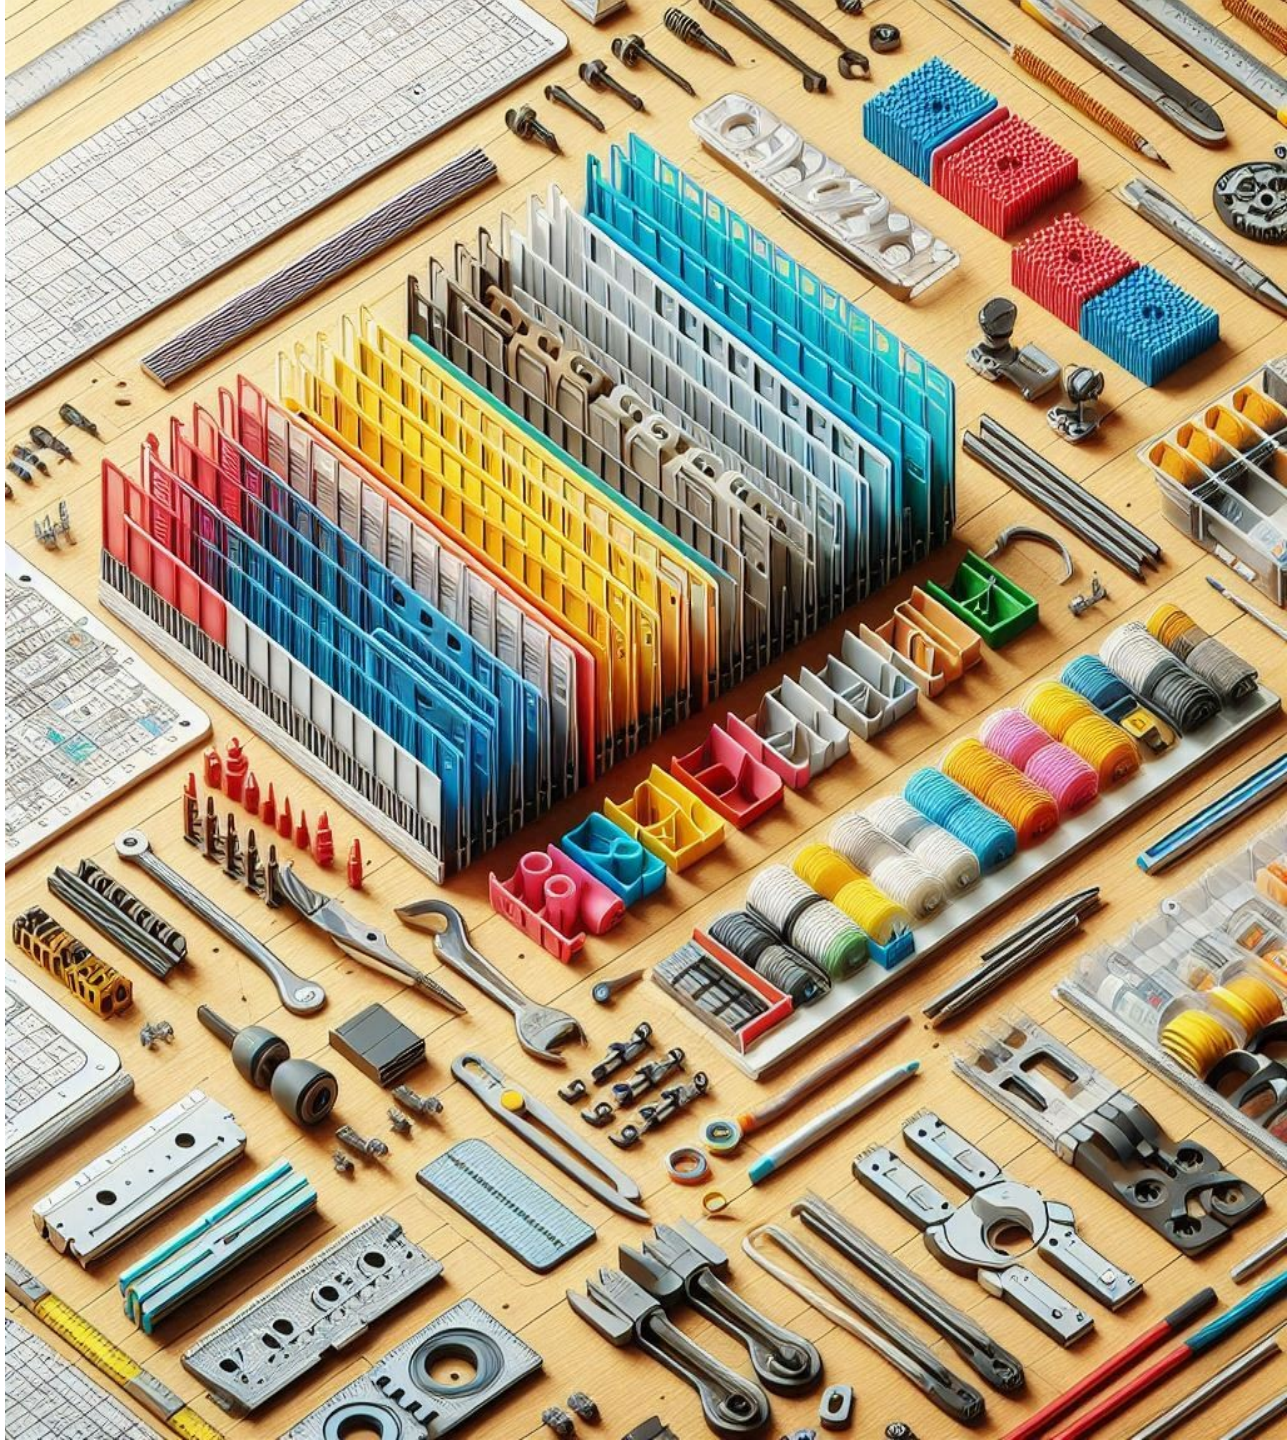

Create Five MBR Partitions on /dev/sdb

1. View storage devices
2. Inspect /dev/sdb
3. Create msdos disk label
4. Create five partitions
5. Verify the partitions

https://asciinema.org/a/702306

File Systems

- Controls how files are stored on a storage device
- Many choices
  - **EXT4**, BTRFS, ZFS, **XFS**
- Red Hat recommends **XFS**

XFS (eXtensible File System)

Journaling File System

Supports up to 1024 TiB Devices

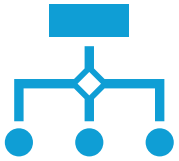

XFS file systems can only grow, not shrink

EXT4 supports resizing in both directions

Create File Systems on /dev/sdb

1. Inspect /dev/sdb
2. Create XFS File System on /dev/sdb1
3. Create EXT4 File System on /dev/sdb2
4. Verify the file systems

https://asciinema.org/a/702307

Mount File System Temporarily

Mount command

Similar sytax as *cp*

mount /from /to

Mount File System automatically at Boot

/etc/fstab

- F-Stab or FS-tab(le), you decide! 😆

Entries are translated to Systemd.Mount units

/etc/fstab fields

Field	Description
Device	The device that must be mounted. A device name, UUID, or label can be used.
Mount Point	The directory or kernel interface where the device needs to be mounted.
File System	The file system type.
Mount Options	Mount options.
Dump Support	Use 1 to enable support to back up using the **dump** utility. This may be necessary for some backup solutions.
Automatic Check	This field specifies whether the file system should be checked automatically when booting. Use 0 to disable automated check, 1 if this is the root file system and it has to be checked automatically, and 2 for all other file systems that need automatic checking while booting. Network file systems should have this option set to 0.

Mount /dev/sdb1 file system on /mnt/records

1. Inspect /dev/sdb
2. Temporarily mount /dev/sdb1 on /mnt/records
3. Verify temporary mount
4. Unmount /dev/sdb1
5. Persistently mount /dev/sdb1 on /mnt/records
6. Verify mount
7. Reboot and verify

https://asciinema.org/a/702308

Logical Volume Management (LVM)

- Abstraction over storage devices
- Provides flexibility
  - Resize volumes without stopping applications

Layers

- Physical Volumes
- Volume Groups
- Logical Volumes

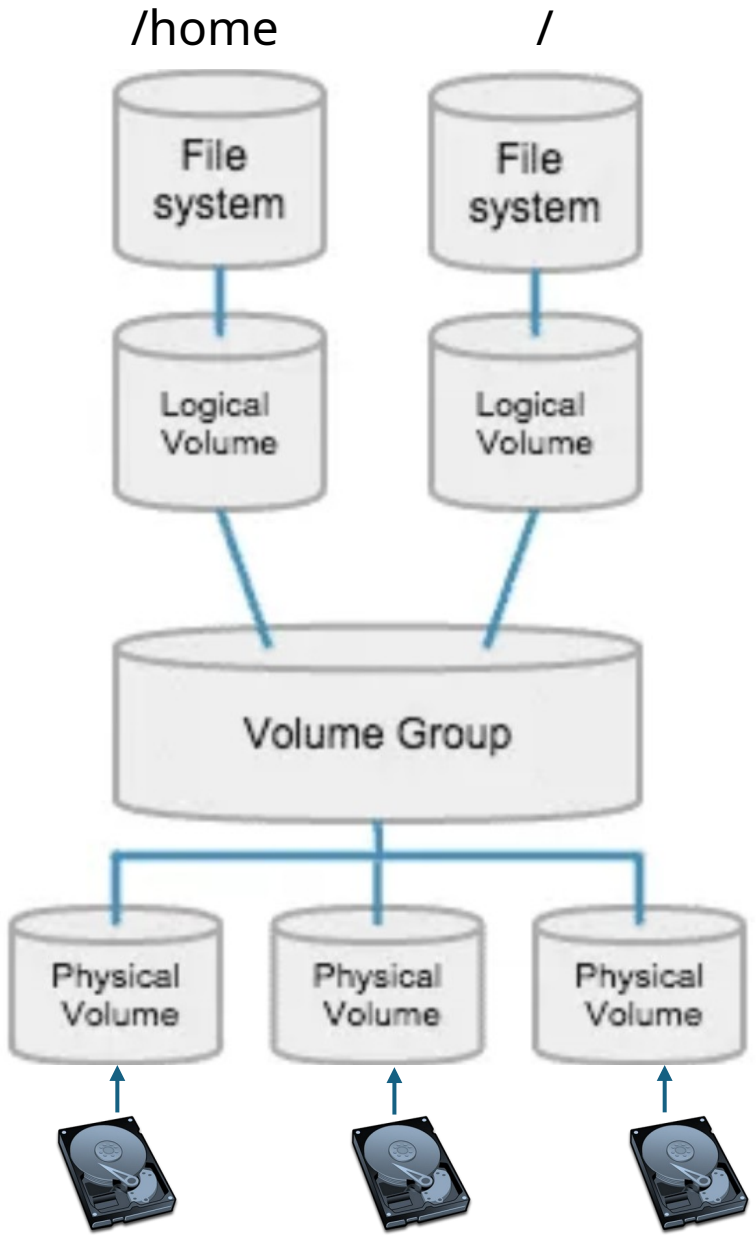

Create Logical Volume

- /dev/sdc1 → Physical Volume
- Physical Volume → Volume Group (vg1)
- Volume Group → Logical Volume (lv1)
- Mount Logical Volume

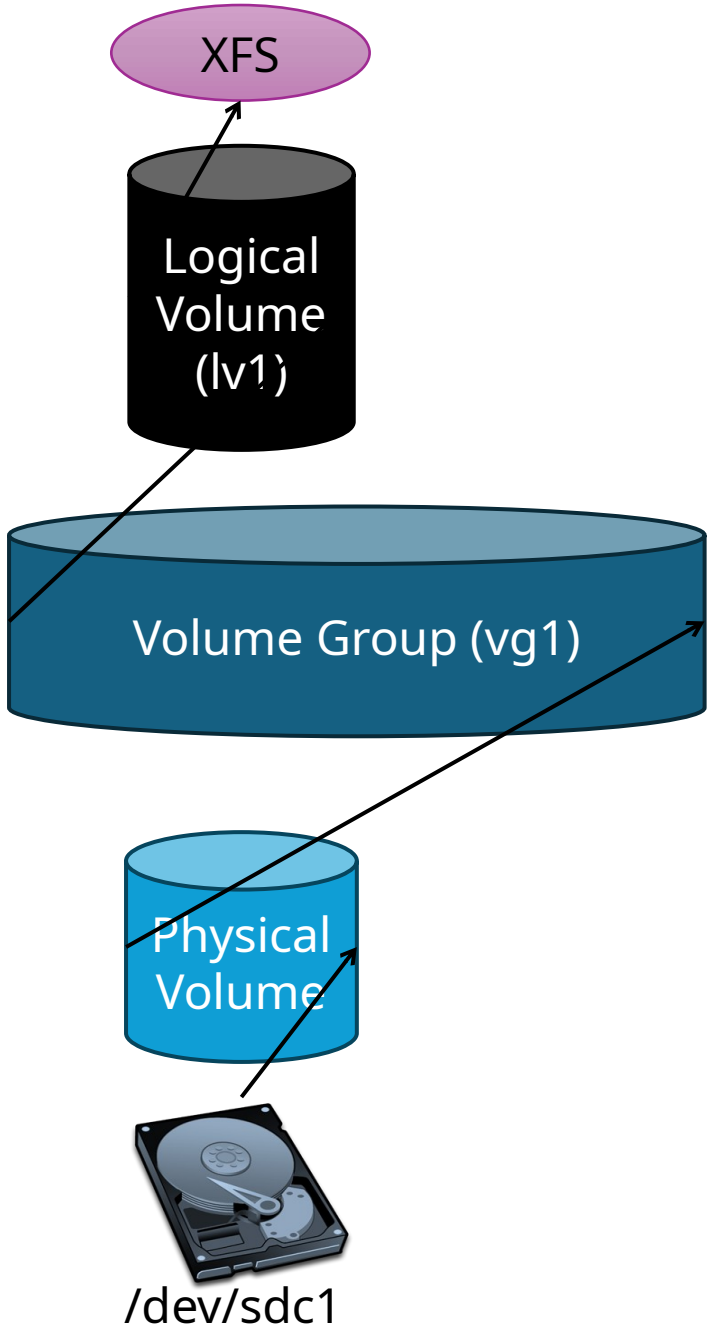

LVM Components - Create

Physical Volume	pvcreate, pvs, pvdisplay, pvremove
Volume Group	vgcreate, vgs, vgdisplay, vgremove
Logical Volume	lvcreate, lvs, lvdisplay, lvremove

Create a Logical Volume

1. Identify and inspect a storage device
2. Create /dev/sdc1 partition to be used in LVM
3. Create a Physical Volume
4. Create a Volume Group
5. Create a Logical Volume
6. Locate the Logical Volume
7. Create a file system on /dev/vg1/lv1
8. Mount file system to /mnt/evidence
9. Verify the file system

https://asciinema.org/a/702649

Extend Logical Volume

- /dev/sdd1 → Physical Volume
- Physical Volume → Volume Group (vg1)
- Expand Logical Volume (lv1)

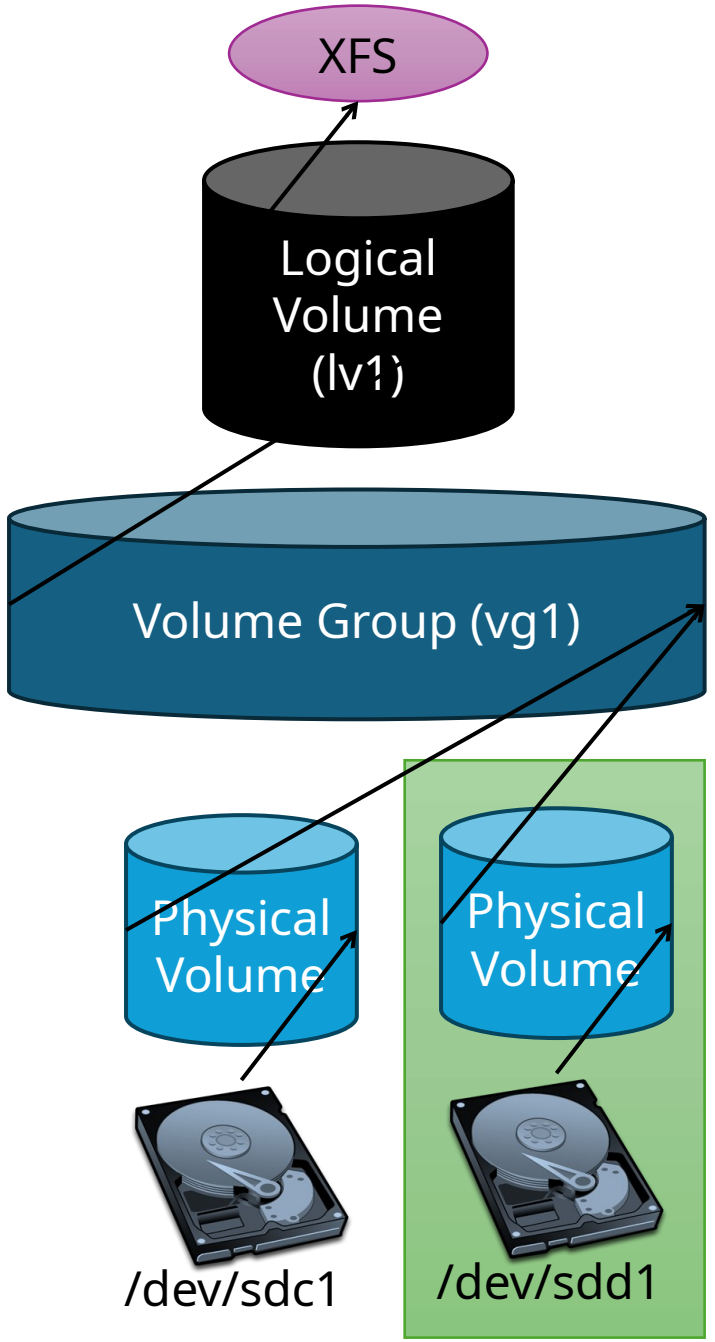

LVM Components - Extend

Volume Group	vgextend
Logical Volume	lvextend
File System	xfs_growfs

Extend Logical Volume lv1

1. Inspect lv1 Logical Volume
2. Identify and inspect a storage device
3. Create /dev/sdd1 partition to be used in LVM
4. Create a Physical Volume
5. Extend the vg1 Volume Group
6. Extend the lv1 Logical Volume
7. Extend the file system
8. Verify the file system

https://asciinema.org/a/702661

Stratis

- Brand New
- Abstraction over storage devices
- Advanced Features out of the box
  - Thin provisioning
  - Snapshots
  - Encryption
- XFS Only

Create Stratis File System

- Create Storage Pool from block devices
- Create file systems from the pool
  - Thin provisioned

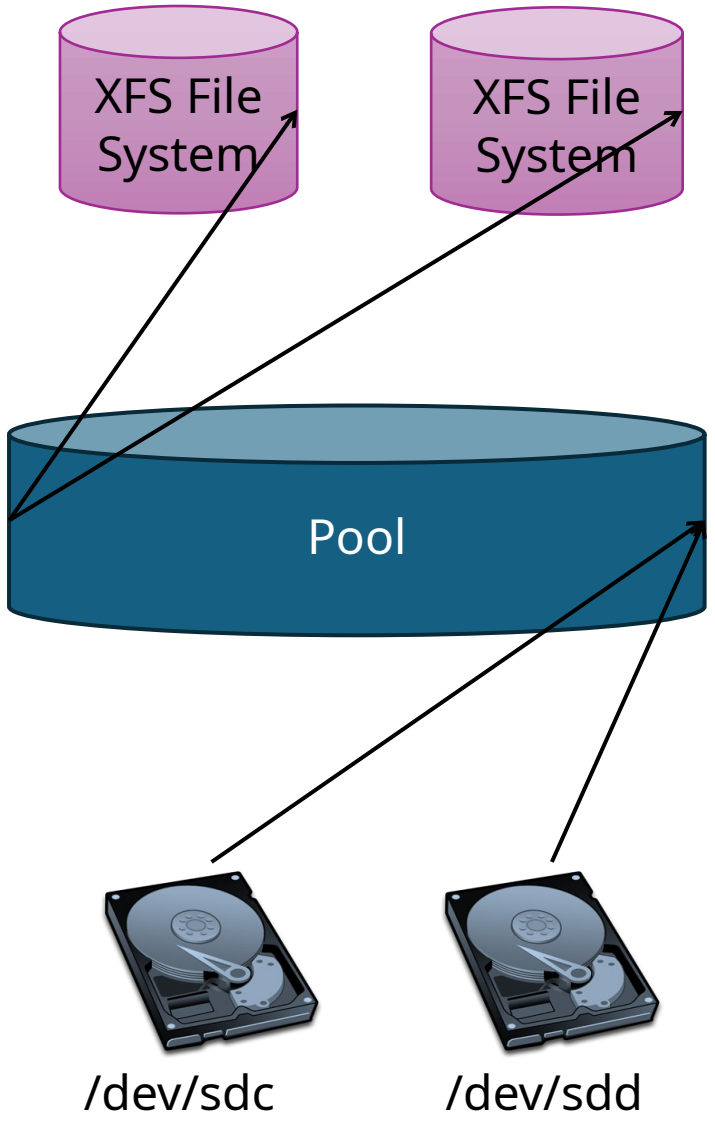

Stratis Components

Packages

stratisd – Systemd service
Stratis-cli – Command line interface

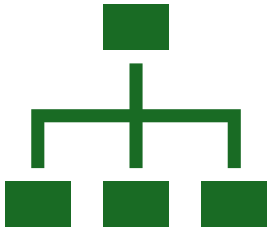

Command

stratis [pool | filesystem | blockdev | ...]

Create a File System on Stratis

1. Install the statis packages
2. Start and enable the stratisd service
3. Identify and inspect a storage device
4. Inspect stratus-cli
5. Create stratis pool
6. Create stratis file system
7. Locate file system
8. Mount file system
9. Inspect thin provisioning

https://asciinema.org/a/702866

Stratis Snapshots

- Take a “moment in time” backup
- Easily revert or rollback changes

Restore data using a stratis snapshot

1. Take a snapshot of the *fs1* file system
2. Catastrophically destroy data
3. Mount the snapshot
4. Restore data from snapshot
5. Unmount the snapshot
6. Destroy the snapshot

https://asciinema.org/a/7###