

**LINBIT**

**Software  
Defined Storage  
(SDS)**

**ALWAYS** ON  
AVAILABLE  
EFFICIENT





# LINBIT SDS

The obvious approach to enterprise grade SDS

# LINBIT Software-Defined Storage



**LINBIT HA**

- NFS / CIFS / iSCSI
- KVM / VMWare / Xen
- Databases
- Fileservers
- Webservers
- NagiosXI
- Messaging (MQ)
- Nearly any other app

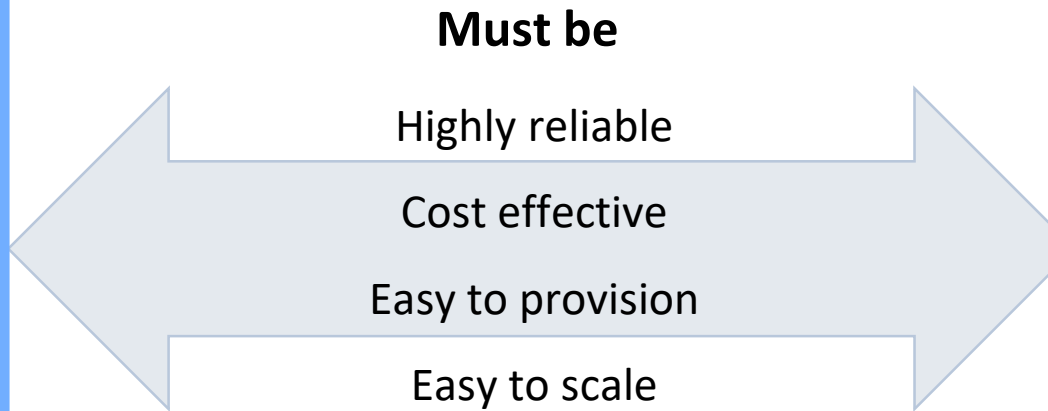
**LINBIT SDS**

**Container-native**

- OpenShift
- Kubernetes
- Docker

**Cloud-native**

- OpenNebula
- OpenStack
- Proxmox VE



# OS-Based Storage Technology



- Linux already provides several storage gems:
  - LVM
  - RAID
  - SSD cache tiers
  - De-duplication
  - Targets & initiators



Native Storage Management Capabilities

# Container Storage



## LINUX BLOCK STORAGE MANAGEMENT FOR CONTAINERS

**LIN**STOR

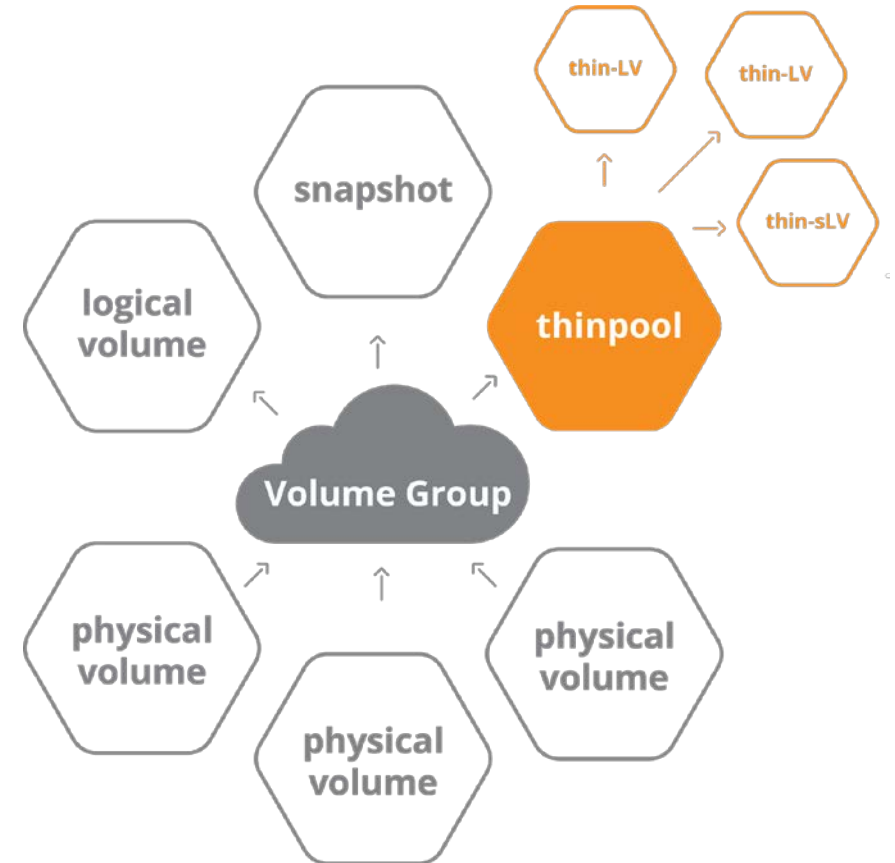


# LVM

## Capabilities

- Based on device mapper
- Original objects
  - PVs, VGs, LVs, snapshots
  - LVs can scatter over PVs in multiple segments
- thinlv
  - thinpools = LVs
  - Thin LVs live in thinpools
  - Multiple snapshots are efficient!

## Linux LVM




# RAID

## Capabilities

- Original MD code
  - `mdadm` command
  - Raid Levels: 0,1,4,5,6,10
- Now available in LVM as well
  - Device mapper interface for MD code
  - Do not call it 'dmraid'; that is software for hardware fake-raid
  - `lvcreate --type raid6 --size 100G VG_name`

## Linux RAID



# ZFS on Linux

- Ubuntu eco-system only
- Has its own
  - logic volume manager (zVols)
  - thin provisioning
  - RAID (RAIDz)
  - caching for SSDs (ZIL, SLOG)
  - and a file system!





# Cache Devices

- dm-cache
  - device mapper module
  - accessible via LVM tools
- bcache
  - generic Linux block device
  - slightly ahead in the performance game



# Linux's Inline Deduplication



- Virtual Data Optimizer (VDO) since RHEL 7.5
  - Red hat acquired Permabit and is GPLing VDO
- Linux upstreaming is in preparation
- In-line data deduplication
- Kernel part is a device mapper module
- Indexing service runs in user-space
- Asynch or synchronous writeback
- Recommended to be used below LVM

# Targets & Initiators

## Capabilities

- Open-ISCSI initiator
- letd, STGT, SCST
  - mostly historical
- LIO
  - iSCSI, iSER, SRP, FC, FCoE
  - SCSI pass through, block IO, file IO, user-specific-IO
- NVMe-OF
  - target & initiator

## Targets & initiators

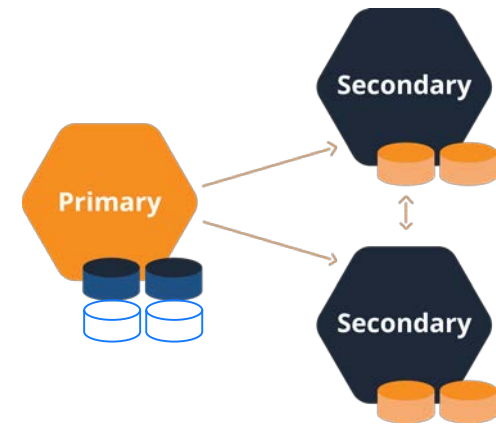
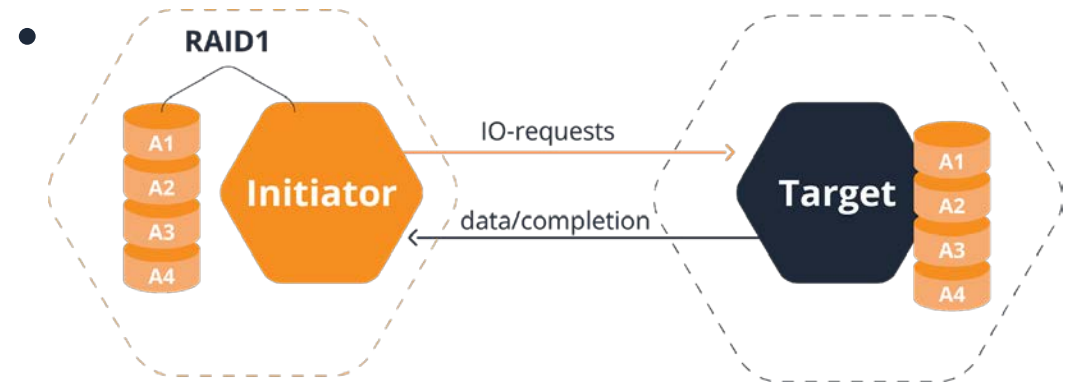


# DRBD – Mainline Linux Kernel

## Capabilities

- 1000's of Nodes
  - Up to 32 Synchronous or async replicas per volume
  - Automatic partial resync after connection outage
  - Multiple resources per node possible (1000s)
- Diskless nodes
  - Intentional diskless (no change tracking bitmap)
  - Disks can fail
- Reliable
  - A node knows the version of the data it exposes
  - Checksum-based verify & resync
  - Split brain detection & resolution policies
  - Fencing
  - Quorum
  - Dual Primary for live migration of VMs only!

## DRBD



**LIN**  **STOR**

**Capabilities → Products → Solution**

**LIN**  **BIT**


# LINSTOR - Goals

## Goals

### Build storage from generic (x86) nodes

- Serve SDS consumers (OpenStack Cinder, Kubernetes, Custom platforms)
- Allow multi-tenancy
- Enable multiple Deployment architectures
  - Distinct storage nodes
  - Hyperconverged with hypervisors / container hosts
- Don't recreate the wheel
  - Use existing Linux storage components

## Approach

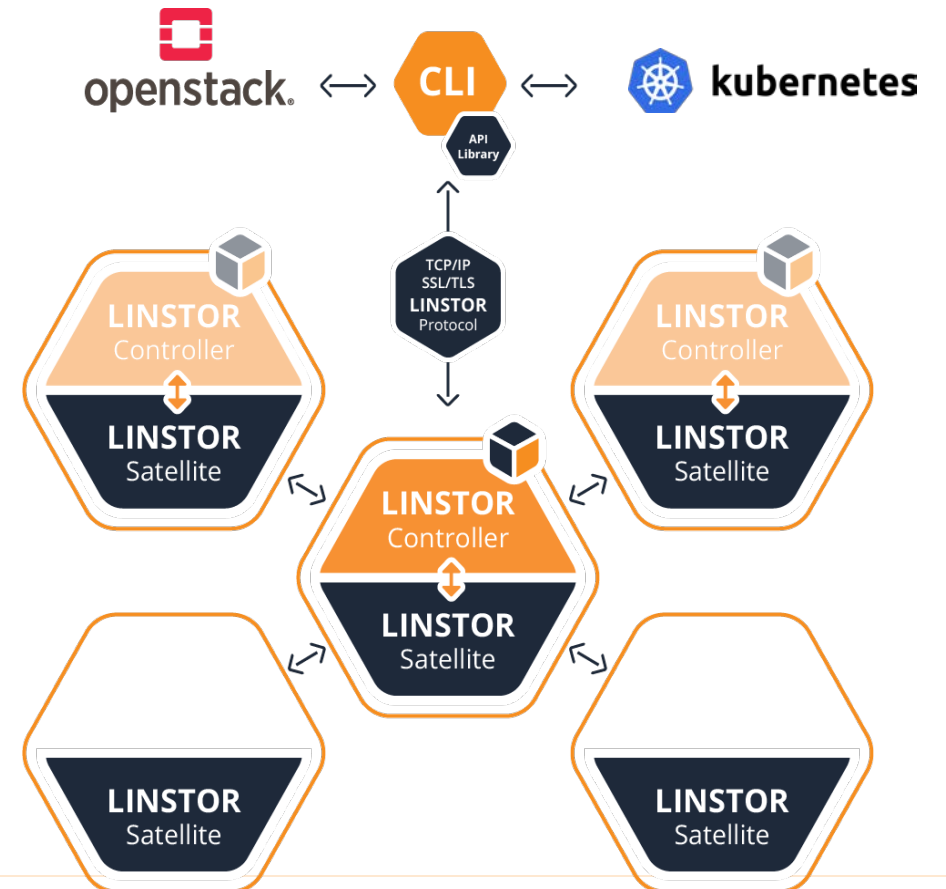


# LINSTOR

## Capabilities

- Controls LVM/ZFS
  - Snapshots
  - Thin
- Multiple VGs
  - For caching SSDs
  - Different pools
- Controls DRBD

## LINSTOR



# New LINSTOR Features

## Complete

- Snapshot support
- Multiple geo-diverse sites with DRBD Proxy
- File-based backing storage (via “loop”)
- Secure REST API (HTTPS)
- Multi-user REST API support (LDAP)
- PMEM backing storage for DRBD metadata
- Support for several orchestrators and cloud platforms
- Swordfish API
  - Can manage NVMe-oF targets and initiators

## Roadmap

- iSCSI Targets: creating and attaching (Q4 2019)
- VDO deduplication (2020)
- DRBD 10 (already in alpha)
  - Performance Improvements
  - PMEM caching + journaling
  - Erasure Coding
  - Request Forwarding (“resource chaining”)



# Resources

## High Availability

- <https://www.linbit.com/en/resources/documentation/>
- <https://www.linbit.com/en/drbd-community/drbd-download/>

## Disaster Recovery

- Video Demo (6 minutes)  
<https://www.youtube.com/watch?v=Sf0lPClIDWk>
- Blog post with LINSTOR + LINBIT DR Video demo:  
<https://www.linbit.com/en/demo-extending-linstor-managed-drbd-volume-linbit-dr-node/>

## Software-Defined Storage

- Container Storage:  
<https://www.linbit.com/en/linbit-sds-container-storage/>
- Private Cloud:  
<https://www.linbit.com/en/linbit-sds-private-cloud/>
- Public Cloud:  
<https://www.linbit.com/en/linbit-sds-public-cloud/>
- LINSTOR:  
<https://www.linbit.com/en/linstor/>