

The LINBIT logo is positioned at the top center of the slide. It features the word "LIN" in a bold, orange, sans-serif font, followed by a stylized colon consisting of two vertically aligned circles, and the word "BIT" in a bold, black, sans-serif font. The background of the slide is a dark, blue-toned image of Earth from space, overlaid with a complex network of glowing white lines representing global data connections.

LIN:BIT

A larger version of the LINBIT logo is centered in the upper half of the slide. It uses the same color scheme and font as the smaller logo above, with "LIN" in orange and "BIT" in black, separated by a stylized colon of two circles. The background is the same space-themed network visualization of Earth.

LIN:BIT

Keeping the Digital World Running

**We specialize in Open Source OS based
software defined storage**



DRBD's use cases – Cloud, High Availability and Geo Clustering



Cloud Ready

Storage is one of the most critical components in any cloud environment. It **has to be easy to provision, highly reliable, cost effective** and ideally running on commodity hardware.



High Availability

The principal goal of a high availability solution is to **minimize or mitigate the impact of downtime**.



Disaster Recovery

The principal goal of disaster recovery is to **restore your systems and data to a previous acceptable state** in the event of a failure/loss of a data center.

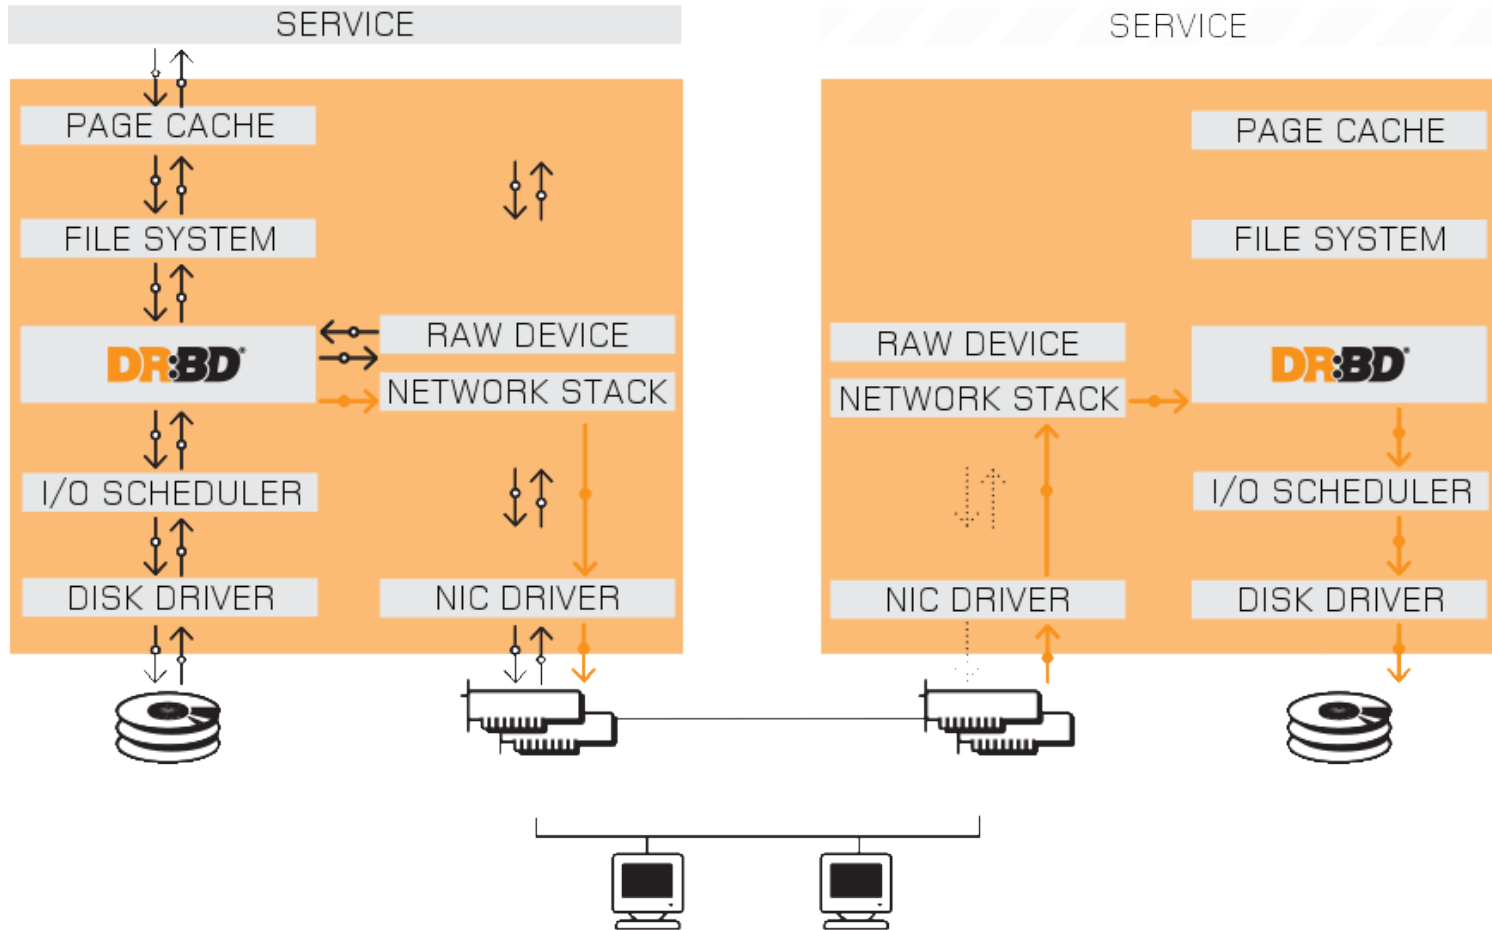
DRBD9 is in OpenStack, so it integrates **seamlessly** with your Private Cloud Environment and **runs on commodity hardware**, other cloud OSs and virtualization platforms.

DRBD **seamlessly replicates data** transparent to your applications and databases, **eliminating single points of failure** within your IT infrastructure.

DRBD can **mirror data asynchronously** over long distances, forming an important building block of your **disaster recovery plan**. Geo-Clustering with automatic fail-over is possible.



DRBD 8.x in 3 minutes





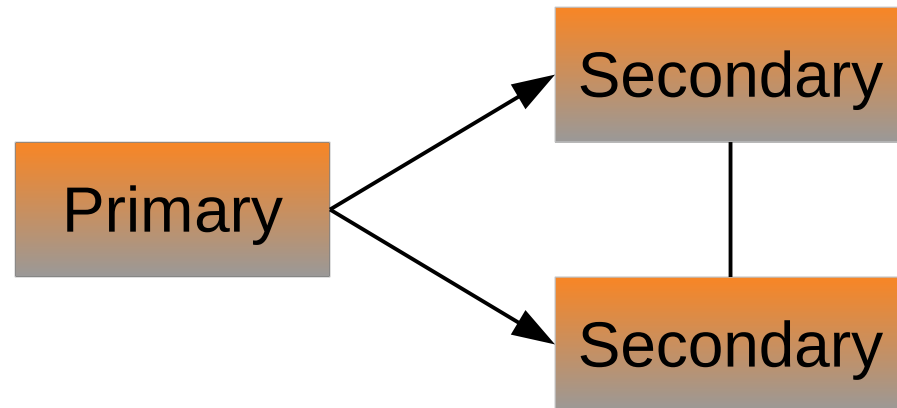
DRBD 8.x key features

- : automatic resync after node or connectivity failure
 - direction, amount, no full resync needed
- : performant in Linux kernel implementation
 - 160k IOPs measured (on SSDs of course)
- : multiple volumes per resource (replication group)
 - write order fidelity within resource
- : comes with Pacemaker integration
- : synchronous and asynchronous repl. LAN and WAN
- : In Linux upstream since 2.6.33 (Released 2010)



New features of DRBD9

- :Up to 31 connections per Resource
 - Fixes the drawbacks of *stacking*
- :Auto promote
- :Transport abstraction (TCP, SCTP, RDMA)
- :drbdmanage





New features of DRBD9

:8.x

```
drbdadm primary <res>
```

```
mount /dev/drbdX ...
```

```
umount /dev/drbdX
```

```
drbdadm secondary <res>
```

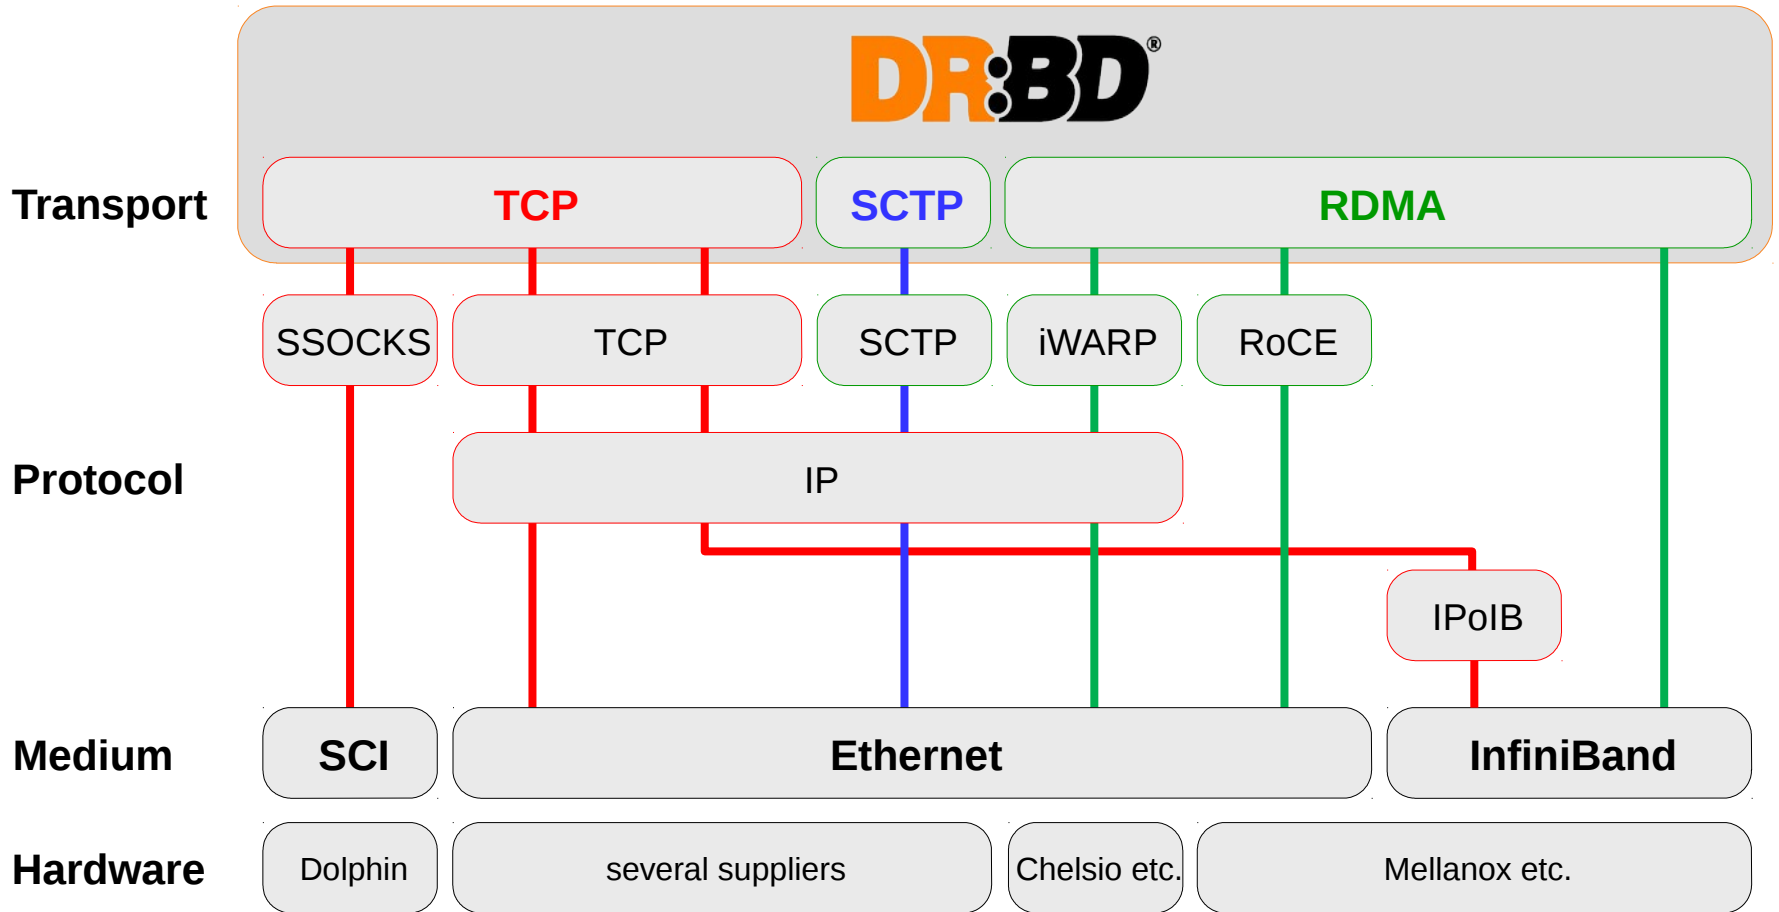
:9.x automatic promote

```
mount /dev/drbdX ...
```

```
umount /dev/drbdX
```



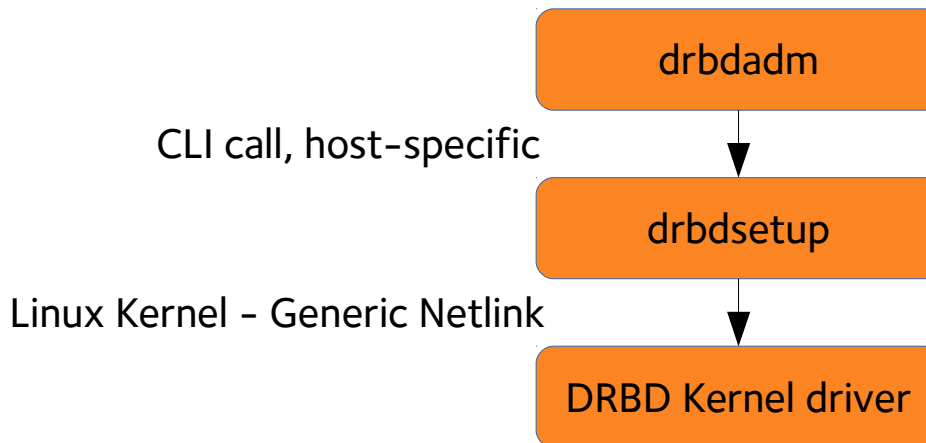
New features of DRBD9



DRBD control plane 8.4

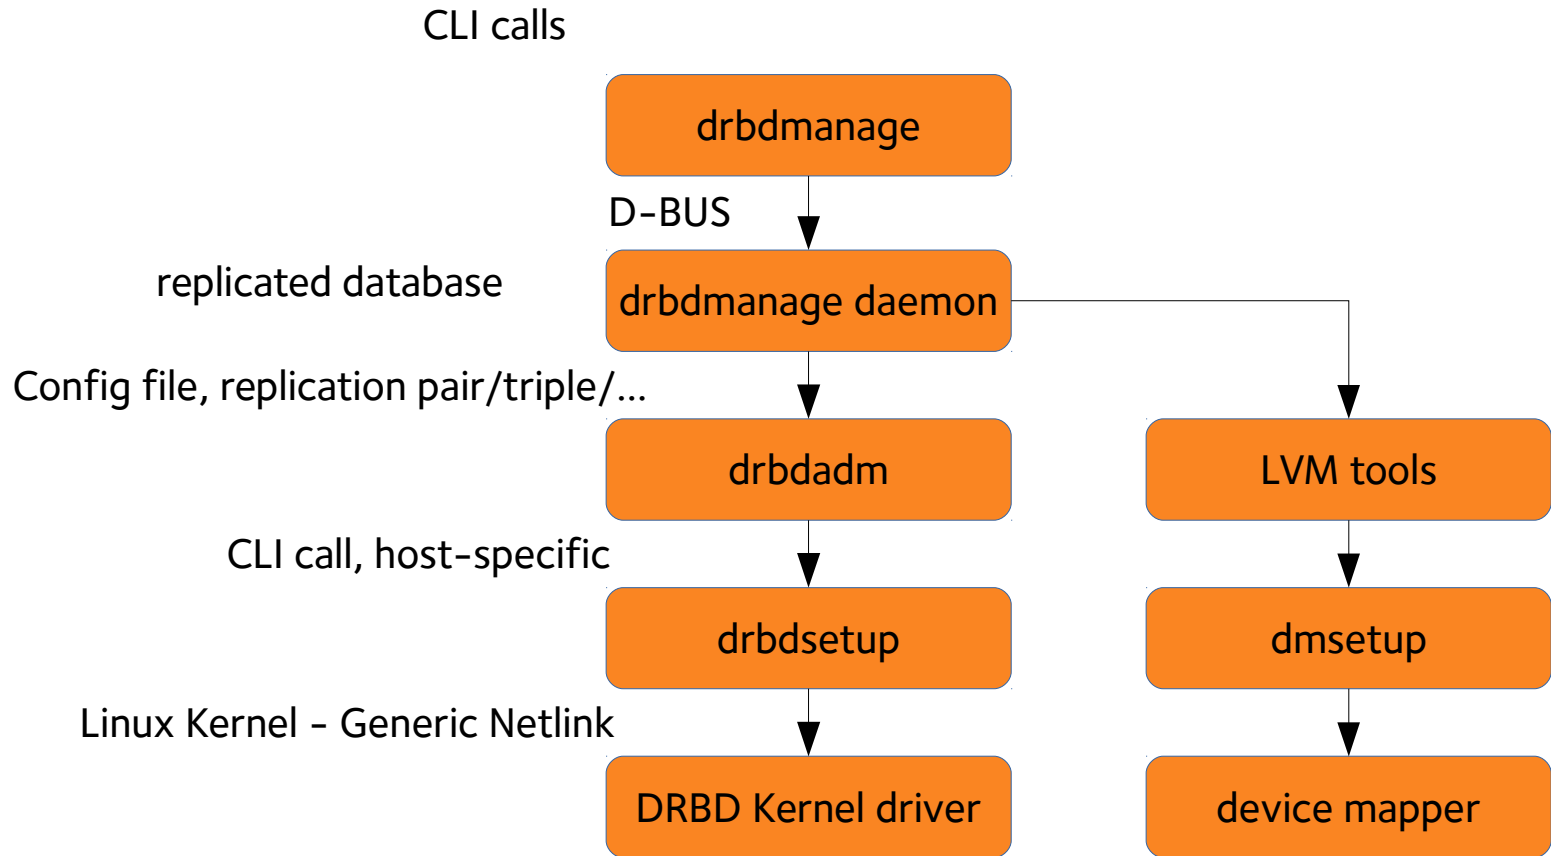
- :You need to create/provide block devices for DRBD
- :You need to distribute DRBD config files among your nodes

Config file, replication pair/triple/...





DRBD control plane 9 – drbdmanage





drbdmanage

:drbdmanage needs

- nodes, with storage in a LVM VG

:you can have

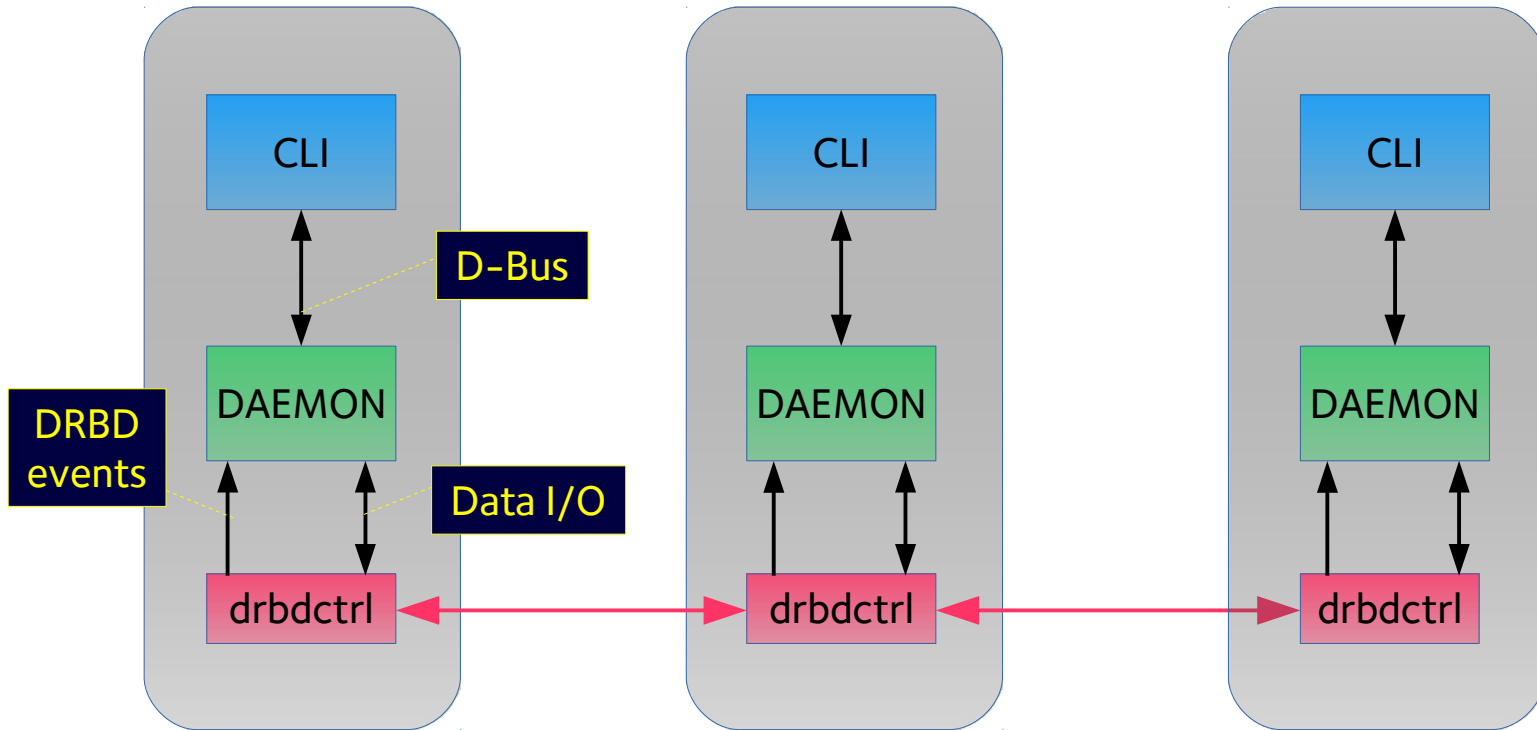
- DRBD resources with ...
 - name, size and **replica count**

:drbdmanage does for you

- calls lvmttools (lvcreate, lvresize, ...)
- distributes DRBD config, activates it with drbdadm



drbdmanage software architecture

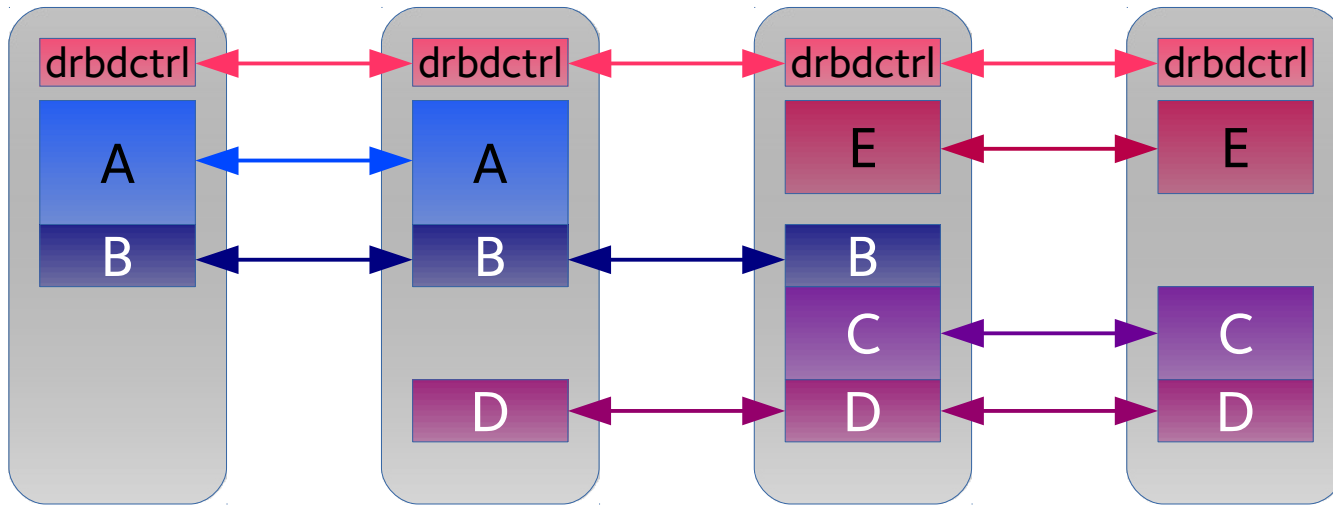




drbdmanage, volume management example

■ control volume – replicated across all nodes

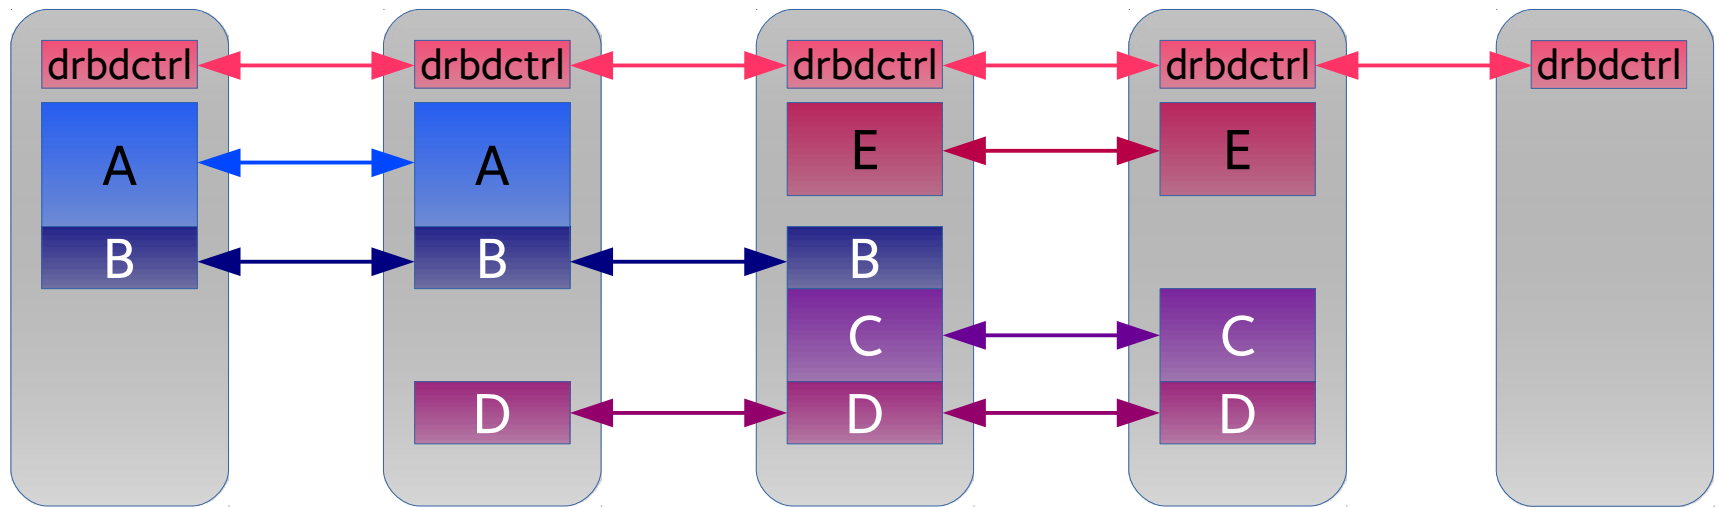
automatically managed replicated volumes





drbdmanage

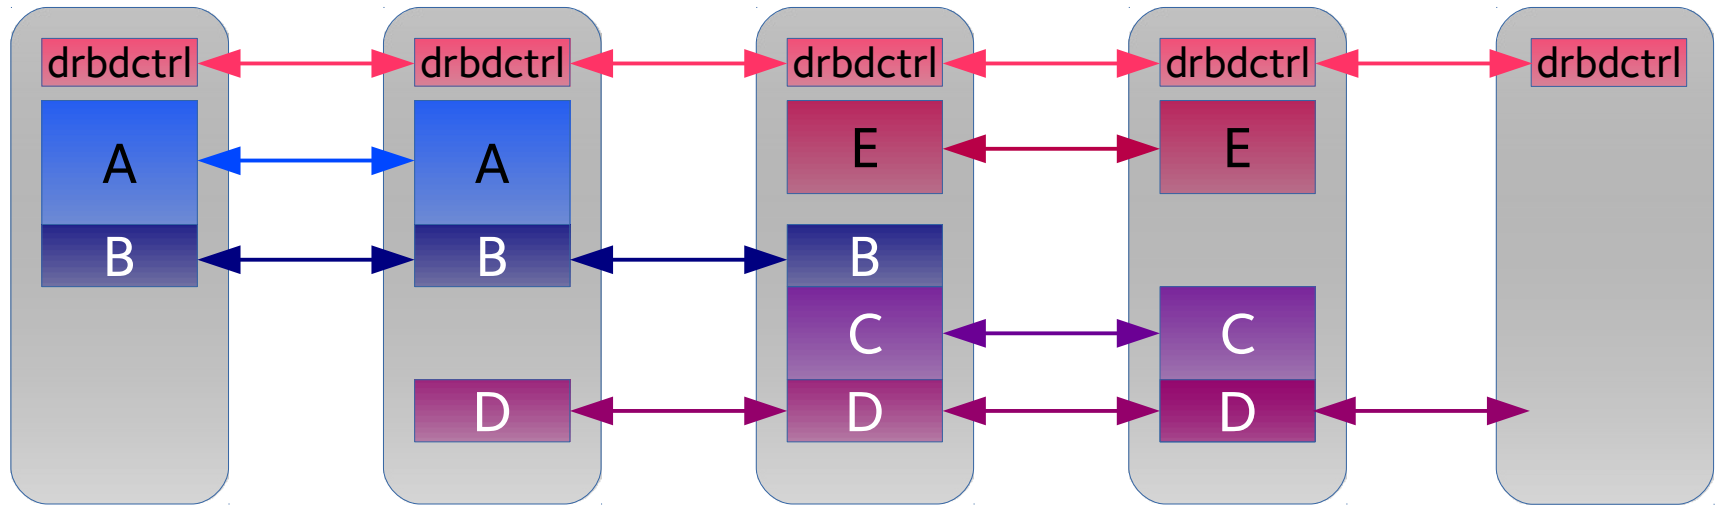
Adding a node





drbdmanage, volume management example


Rebalancing





drbdmanage, volume management example

Using the new space



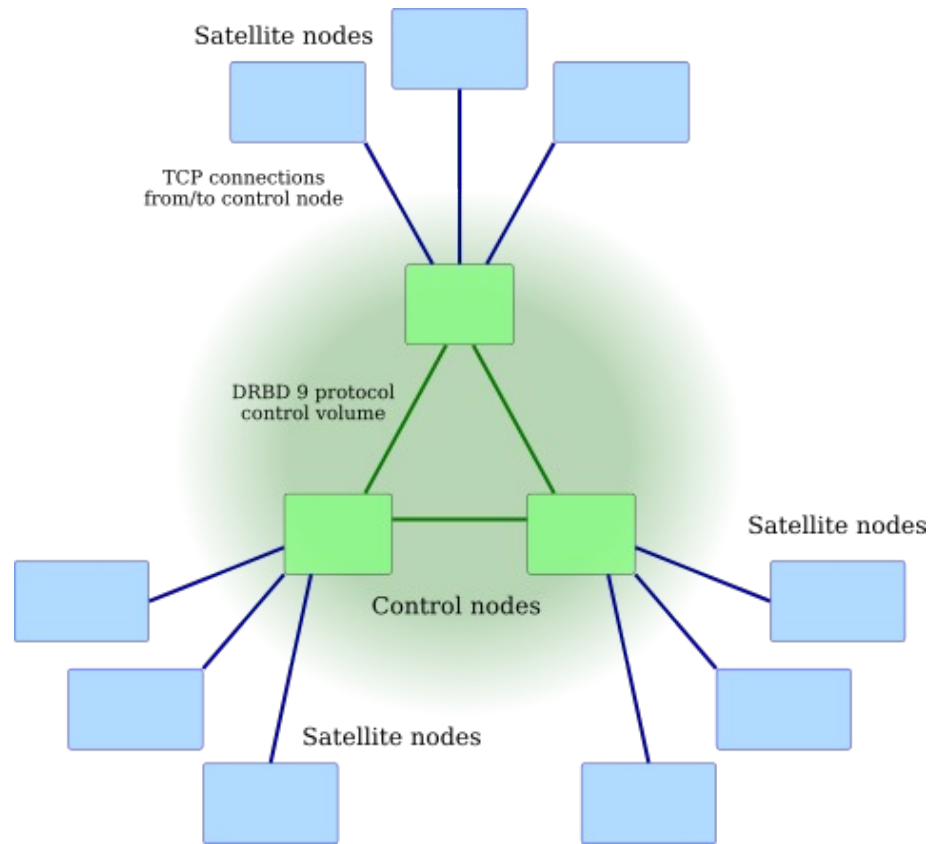


drbdmanage

- :provisioning solution for DRBD
- :implemented in python
- :manages LVs (LUNs) with name, size, replica count
- :manages snapshots
- :may base volumes on thinly provisioned LVM LVs
- :uses DRBD9 itself for its own internal database
 - up to 32 nodes with the complete database
- :scales to 1000s of nodes (with satellite nodes)



drbdmanage satellites






DRBD and the IaaS cloud

Drbdmanage is the glue to Cinder (OpenStack)



OpenStack



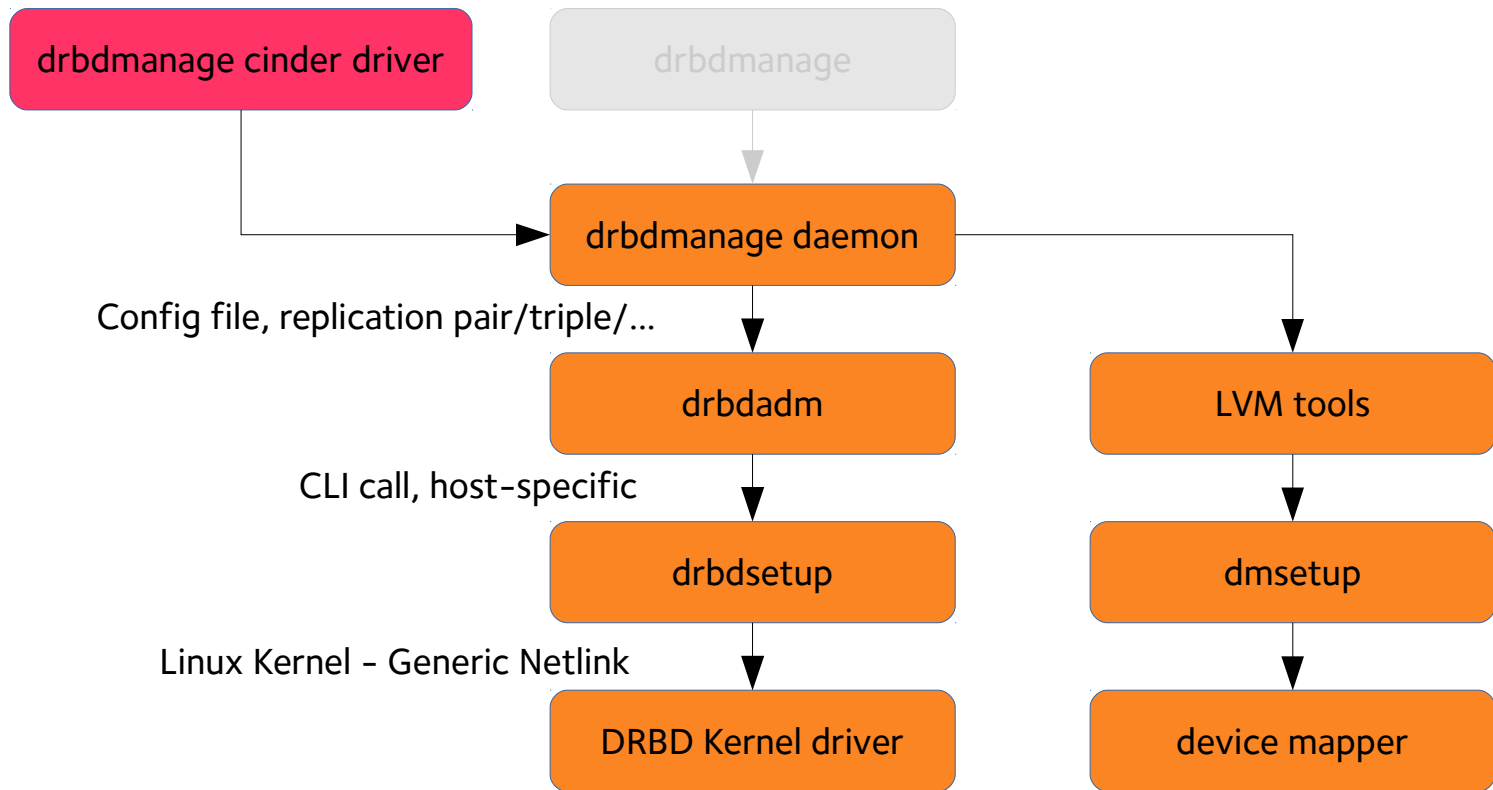
OPST0005

DRBD & drbdmanage dock on here as driver



drbdmanage cinder driver

on cinder node!



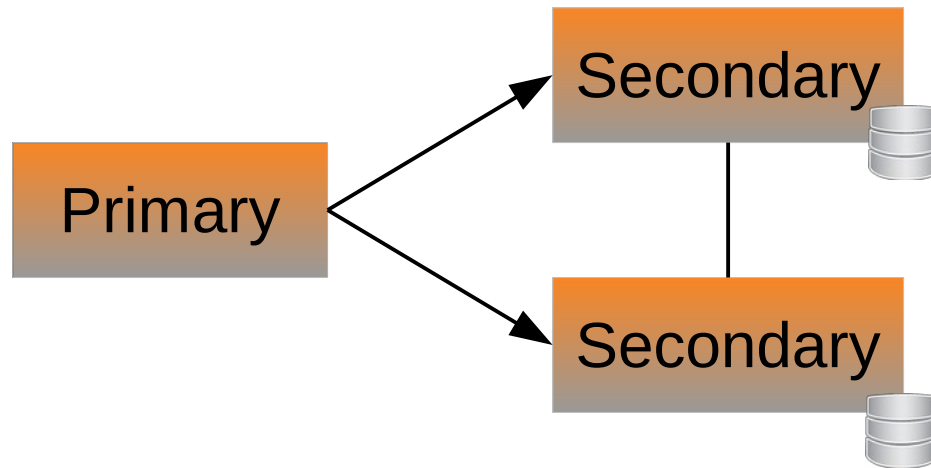


drbdmanage nova driver

- :the nova node...
- :... has a local replica of the volume
 - use it
- :... does not have a local replica: drbd client
 - a permanently diskless node, that is primary and connects to one or more secondaries that actually hold the data



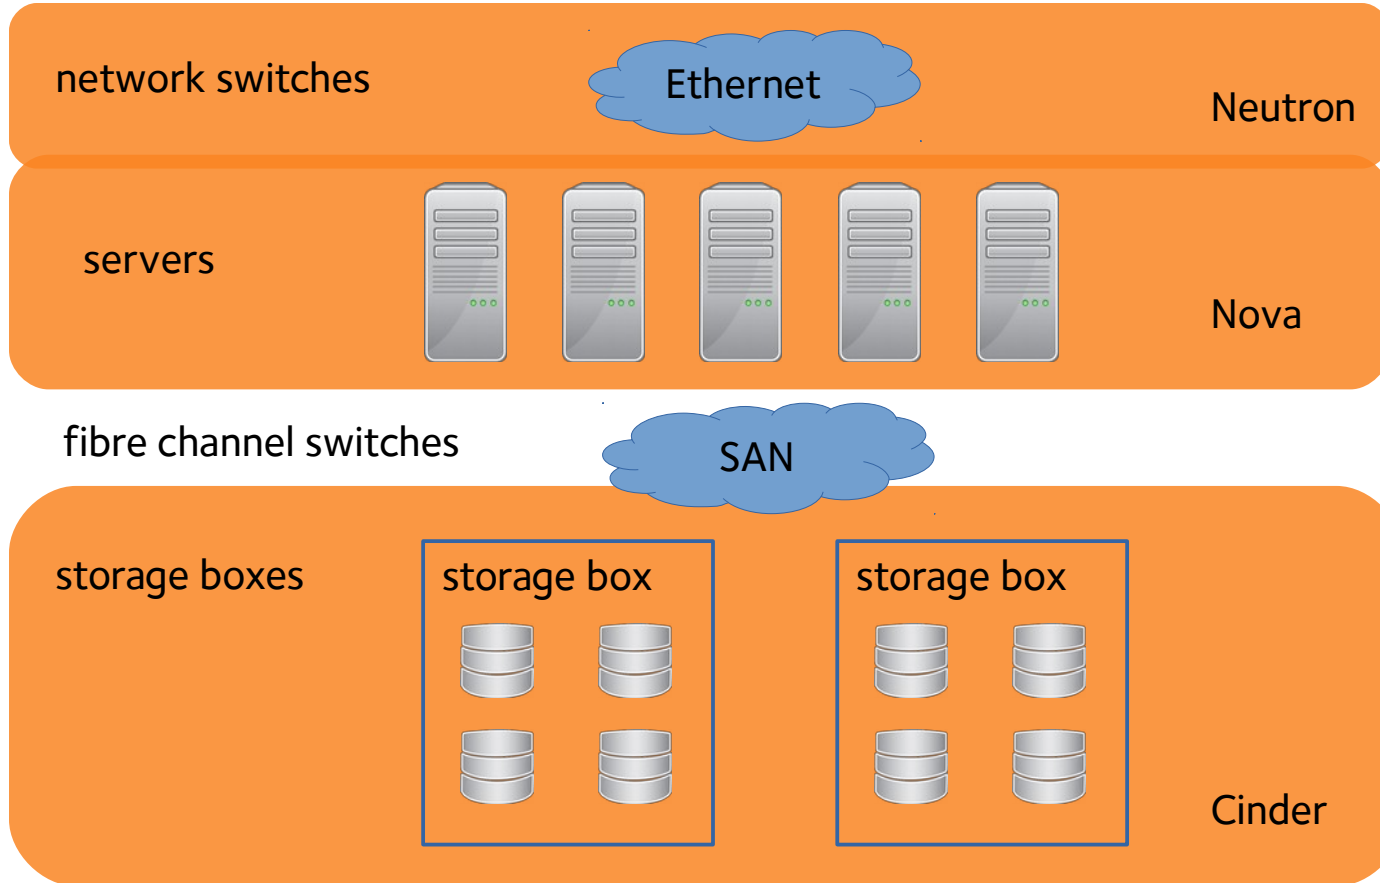
drbd client



: Ideally hint nova where to place VMs

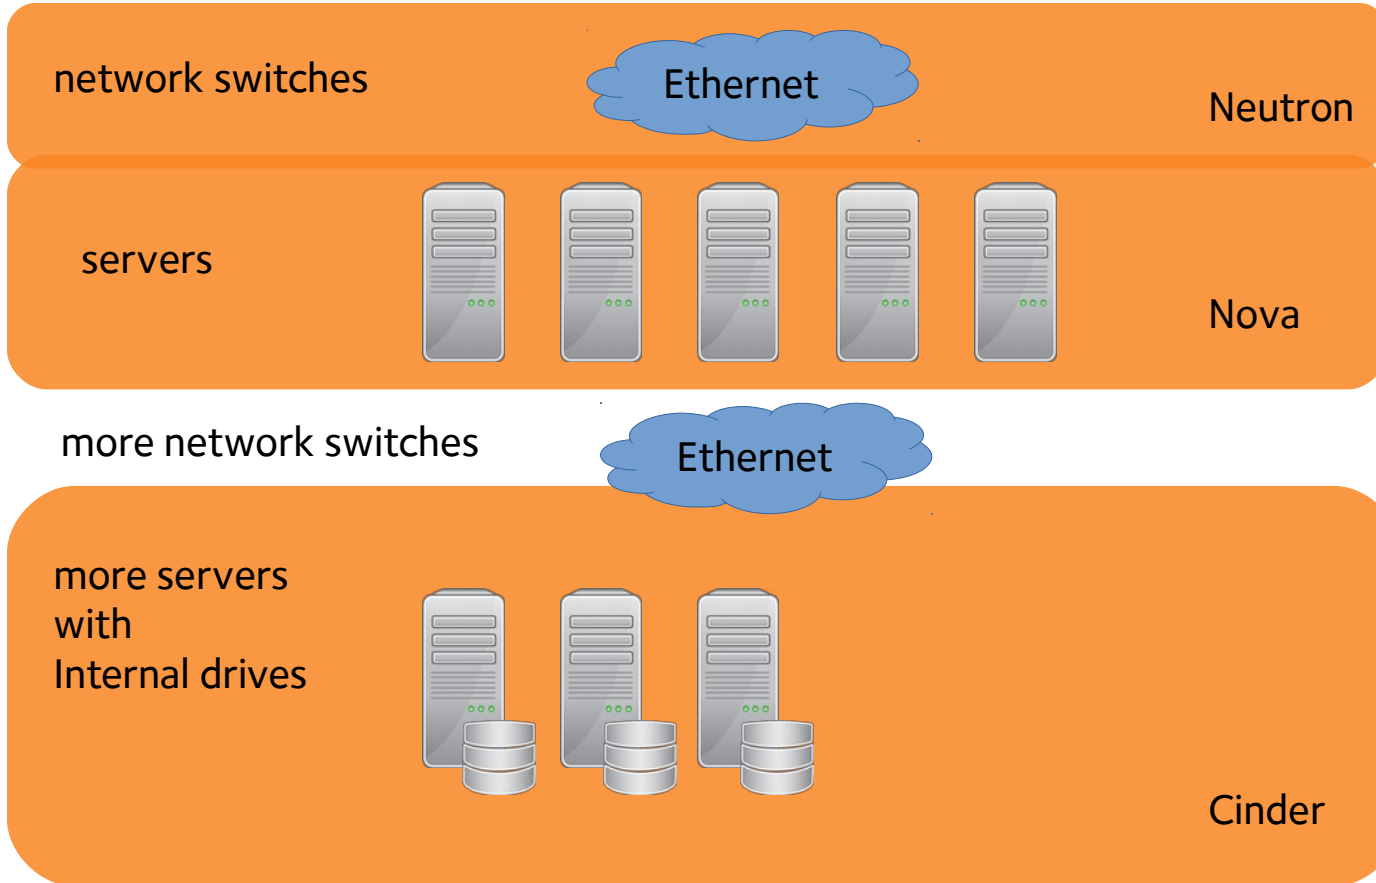


Big storage Vendor's idea of OpenStack






LINBIT's idea of OpenStack





LINBIT's idea of OpenStack



Low latency storage access possible by aligning nova and cinder allocations



Technology road map – achievements in 2015 and next 12 months

2015

OpenStack driver



Enables OpenStack users to base their clouds on DRBD9/DRBDmanage

Released - Tokyo Open Stack Summit - October 16

DRBD on Power



Give users of Power8 machines access to LINBIT's products

Released – with IBM in Germany – November 12

RDMA / InfiniBand



Multi-path support: Aggregates bandwidth of configured paths; increases replication link availability

Released – November 2015

2016

OpenNebula driver



Enables OpenNebula users to base clouds on DRBD9/DRBDmanage

Testing stage

MS Windows Support



DRBD for MS Windows

Expected Q2

PeerDirect RDMA



PeerDirect allows a write request to be sent from an InfiniBand HCA directly to an NVMe device

Planning stage



DRBDmanage roadmap for 2016

DRBDmanage 2016

Back end storage

Back-end storage driver besides LVM: zVols (Canonical plans ZoL for Ubuntu LTS 16.04)



DRBD-Proxy

In one instance per node and some named instances per site model



Network multi pathing

Support for DRBD's network multi-pathing and RDMA-transport



OpenAttic

Administer DRBDmanage clusters via the OpenAttic GUI



Multi-tiered storage

Multiple VGs per node, eg. HDD and SSD; explicit use of one pool or both via bcache or dm-cache



More stack drivers


Additional stack drivers as demand grows: Apache CloudStack






OpenStack driver road map

Q4 2015	<u>Cinder driver</u>	Q1 2016	<u>Nova driver</u>	Q2 2016	<u>Align Cinder and Nova</u>		
	<p>Benefit: Highly available storage volumes can be created via the Horizon dashboard and the cinder APIs. Snapshots supported. Replica count in cinder configuration visible as "pools"</p>	<p>Benefit: Replace iSCSI between nova node and cinder nodes with DRBD's native protocol. Improve storage write performance, enable read-balancing, faster and seamless failover.</p>	<p>Benefit: In a hyper-converged architecture, get read performance of local storage and lowest possible overhead for write accesses by aligning Cinder and Nova allocations</p>		<p>Released</p> <p>"Liberty" October 2015</p>	<p>Finished, waiting for release</p> <p>Release: "Mitaka" April 2016</p>	<p>Planning stage</p> <p>Target: "Newton" October 2016</p>



Storage access by DRBD protocol



Hyper-converged model



World's leading OS High Availability and Disaster Recovery Software

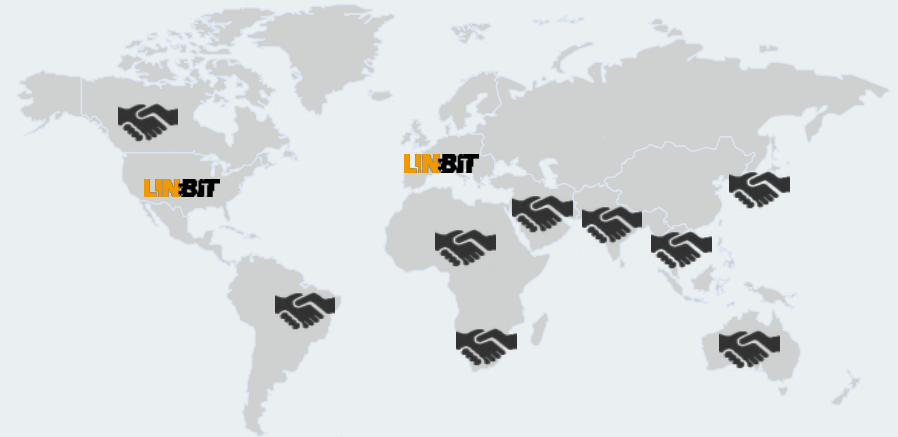


Best performing High Availability / Block device for Open Stack using common off the shelf hardware

20+ times faster than CEPH and GlusterFS*

Only replication technology exceling at both **synchronized short distance** and **asynchronized long distance**

- **Open Source DRBD** supported by proprietary LINBIT products / services
- **Hundreds of thousands** of DRBD downloads
- **OpenStack** comes with **DRBD Cinder driver**
- **100% founder owned**
- Offices in **Europe and US**



References



Partners



*synchronous, single-threaded 100% write workload



Get it! – What are you waiting for?

- : <http://drbd.linbit.com>
- : <http://oss.linbit.com>
- : <http://git.drbd.org>
- : get access to RPM repositories: office@linbit.com




Thank you for your attention

Questions?



LINBIT HA Solutions GmbH
DI Philipp Reisner
Vivenotgasse 48
1120 Wien

E-Mail: philipp.reisner@linbit.com
www.linbit.com



YOUR WAY TO HIGH AVAILABILITY