

High Availability and Backup Strategies for the Lustre MDS Server



Karin Miers / GSI

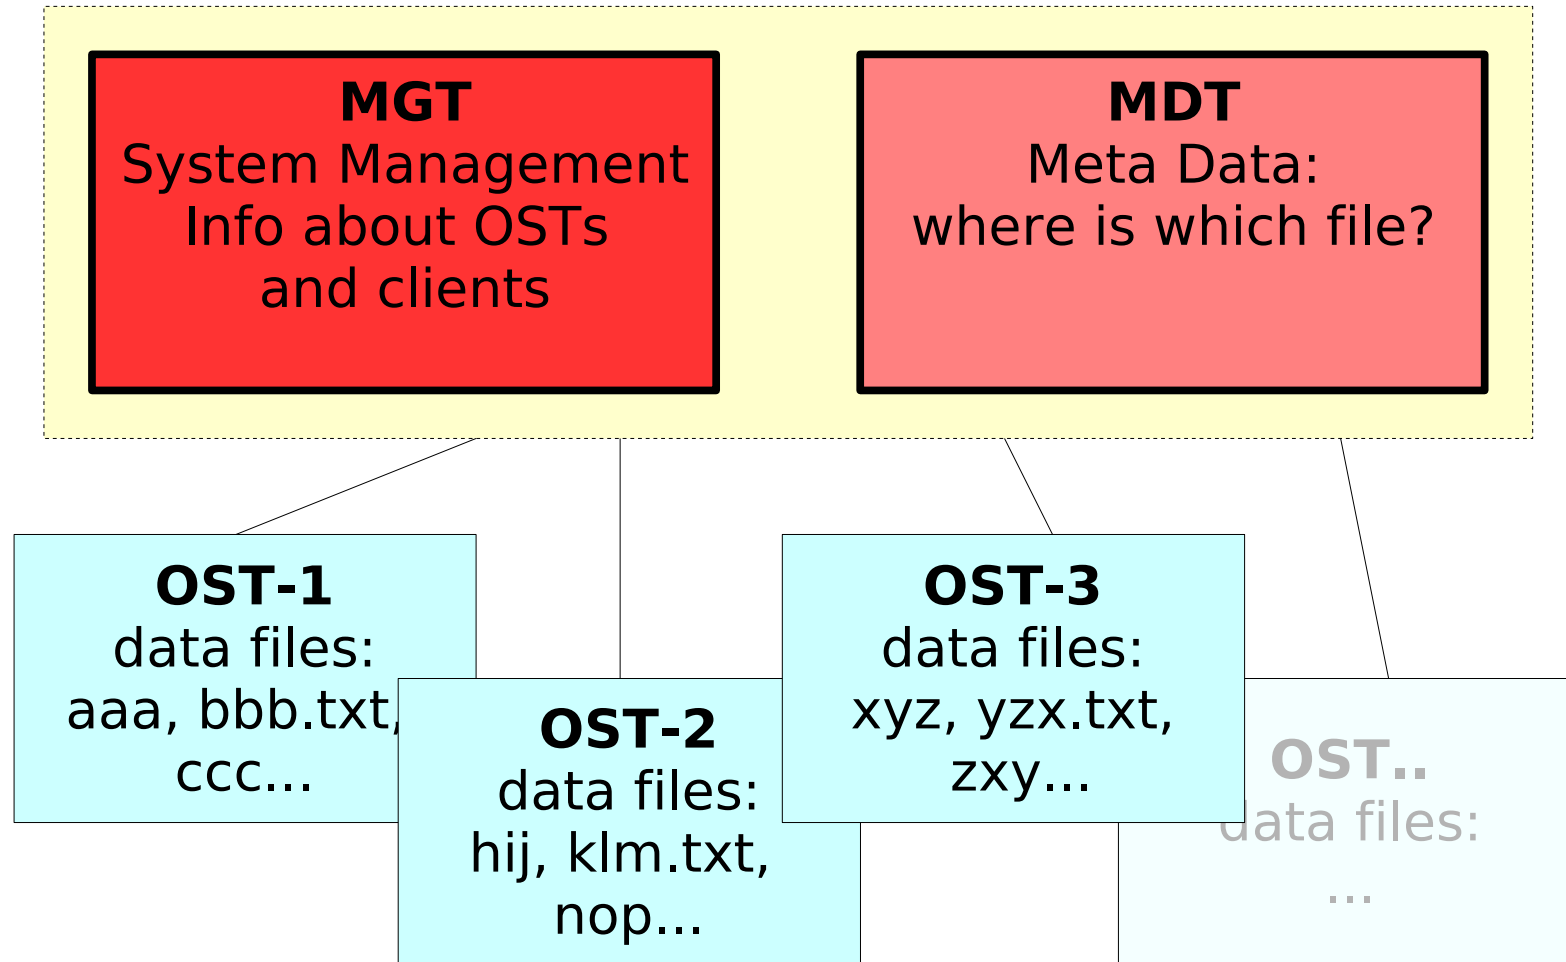
HA and Backup Methods for Lustre

Lustre:

- What is necessary for the production cluster?
or
- What will we do to make our new file system reliable?
 - high availability setup
 - backup of important parts (if high availability setup fails...)

Lustre Components

main parts of a Lustre file system:



In Case of Failure...

What happens if...

... an OST breaks down?

- all data on this OST are not longer available – Lustre continues

... the MDT breaks down?

- all data become inaccessible and probably lost forever...

... the MGT breaks down?

- no data loss, but Lustre becomes inoperable

HA Design

(...based on budget restrictions...)

OSTs are set up single, without backup...

... means data loss is accepted

Same situation as it is now for experiment data

MDT and MGT (=MGS) are set up in a cluster

2 nodes, master / slave which can take over

MDT is written to the backup

otherwise in case of failure ALL data could be lost

MGT is not written to the backup

...no need – can be set up new very fast

Cluster Tools (Software)

Software tools are Open Source (GPL or similar)

Main components:

Heartbeat V2 (2.1.3-5, debian package)

for cluster connection, management and monitoring

DRBD

for redundant data partition

Linux Heartbeat Package

Heartbeat-2...

- ... controls and checks the communication between both (or more) nodes connected by ethernet and / or serial line
- ...checks connectivity to local network
- ... monitors the resources (are MDT/MGT mounted... ?)
- ... (not implemented) can fulfil complicated conditions

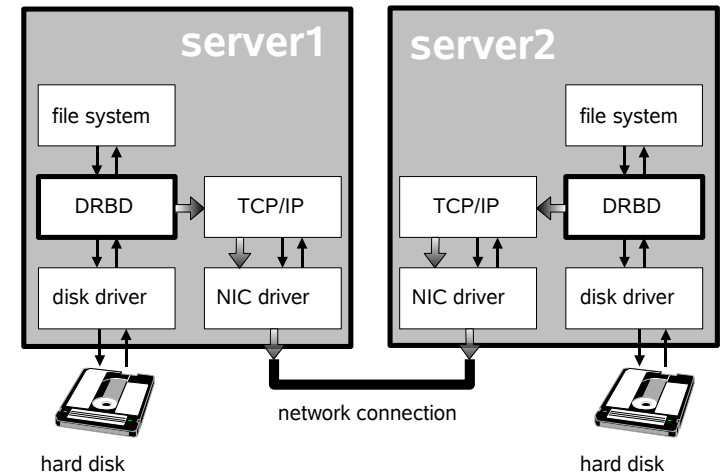
DRBD

Distributed Replicated Block Device

- in principle:

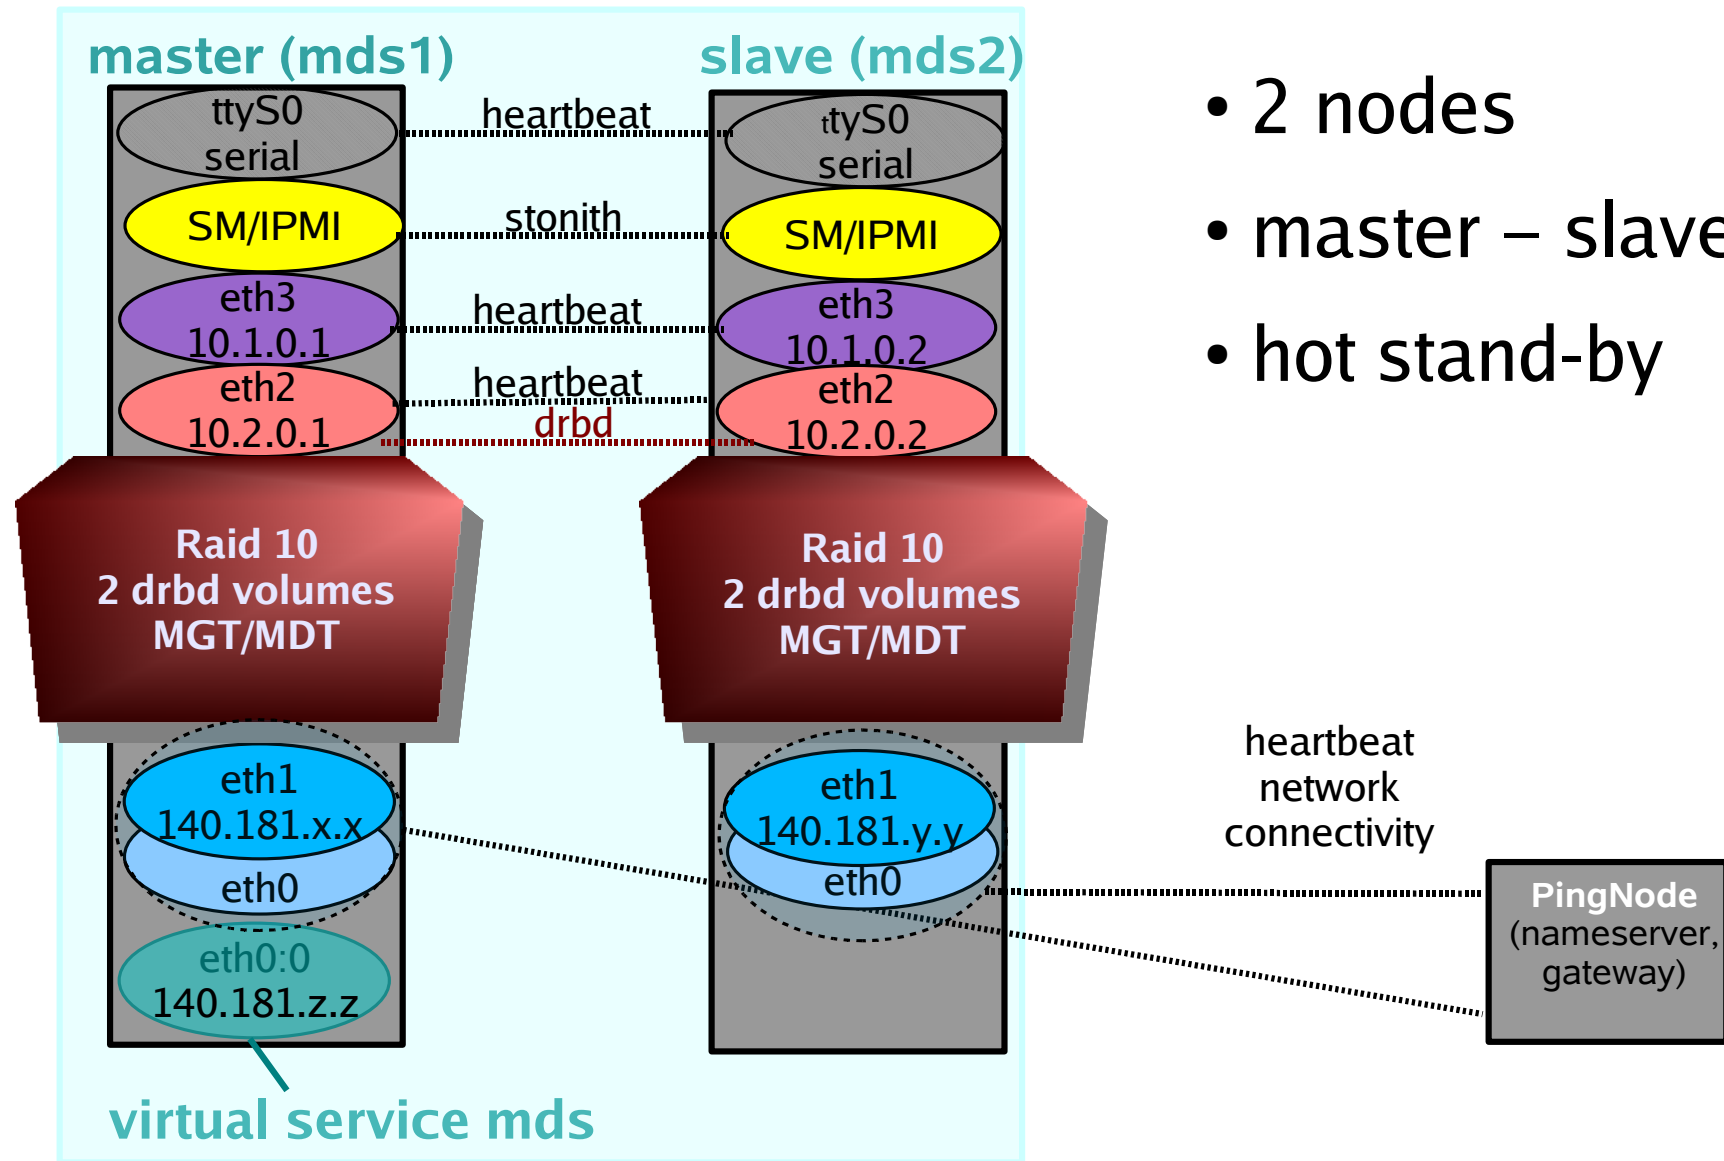
➔ RAID-1 over network

- data exist twice
- real time update on slave
- consistency guaranteed
- fast recovery after failover



- no load balancing
- overhead of drbd:
 - needs cpu power
 - write performance is reduced

MDS Cluster



- 2 nodes
- master – slave
- hot stand-by

lustre has to be told to use eth0:0 instead of eth0!

Failover

Failover tests:

master switched off -> slave takes over automatically, lustre is fully operable after a few minutes:

- heartbeat/drbd ~ 20-30s according to configuration
- lustre ~ few minutes (< 5 min)

MDT Backup Strategy


...just in case... if cluster fails...

Problem:

- permanent write processes on MDT
- backup on active MDT must fail
- no possibility to stop write access for backup duration

→ **DRBD ... there is a copy of the MDT!**



Backup Procedure



- „normal“ state:
- connected drbd
 - synced
 - HA

backup state:

- unconnected drbd
- snapshot MDT
- no HA



- after backup:
- connected drbd
 - syncing

Backup Steps in Detail

- *drbdadm disconnnet / detach* mdt
- mount mdt device as ext3 fs
- save extended attributes with *getffatr*
- make a tar archive of the directory and save it
- umount mdt device
- reconnect drbd
- time factor - ~25 s for 0.5 GB MDT space (lustre test system with appr. 200 000 files/800 GB) – but will depend mainly on size of MDT

Restore Procedure

(... worst case scenario – hopefully will never happen...)

destruction of MDT with „dd“...

- umount all OSTs
- umount mdt device
- format and tune mdt
- mount mdt with *-ldiskfs*
- restore tar archive and extended attributes
- umount mdt device
- activate mdt (*mount -t lustre*)

lustre recovers soon (appr. 5 min, time needed to restore sessions), no files lost since backup!

Backup Problems

No data loss with successful backup, but...

... appr. every third backup fails with error - inconsistent file system

drbd is a layer between hardware and file system and does not care or see the file system

2 possibilities:

- deactivate MDT shortly before DRBD is disconnected
 - no idea how much disturbance this causes under heavy used lustre?
- file system check on slave copy of MDT seems to help and produces correct backups - always?

Open Questions and Improvements

HA:

- setup well established and used successfully for other services for years
 - improvement of monitoring scripts and integration in heartbeat-2

Backup / Restore:

- test of backup / restore procedure on heavily used lustre system (until now – test system)
- no HA during backup procedure – 3 nodes?
 - 7zip instead of tar...?

Questions?