

Storage status report for the praguelcg2 site

Alexandr Mikula

Institute of Physics
of the
Czech Academy of Sciences



Akademie věd
České republiky



Storage HW (1)

- DPM
 - Headnode with separate mysql node
 - 12 disk nodes (76-436 TiB capacity)
 - Total Capacity 2.98 PiB
- Each disknode has 10Gbit line
- 8 nodes are running consumer grade 4TB drives (Seagate 5.9k RPM NAS drives)
 - Two upgraded from 2TB RAID optimized drives
 - Rest bought with such drives

NAS drives


- 8 nodes are running consumer grade 4TB drives (Seagate 5.9k RPM NAS drives)
 - Two upgraded from 2TB RAID optimized drives
 - Rest bought with such drives
 - Fail rate is less than 20 drives per year out of ~1000 drives
 - Drive age varies from ~1 month to ~3 years
 - Comparable to fail rate of 2TB drives
 - Most failed in 2015 during AC shutdown (4 at once)

Storage SW

- Most of farm is running latest stable SL6.8
- DPM version 1.9.0 update was triggered by accident but proved stable and problem free for our instance
 - We still run legacy mode
 - DOME is on plan during next few months
 - One problem on DPM 1.8.11 with deletion over HTTP
- XrootD is on version 4.5.0

ZFS

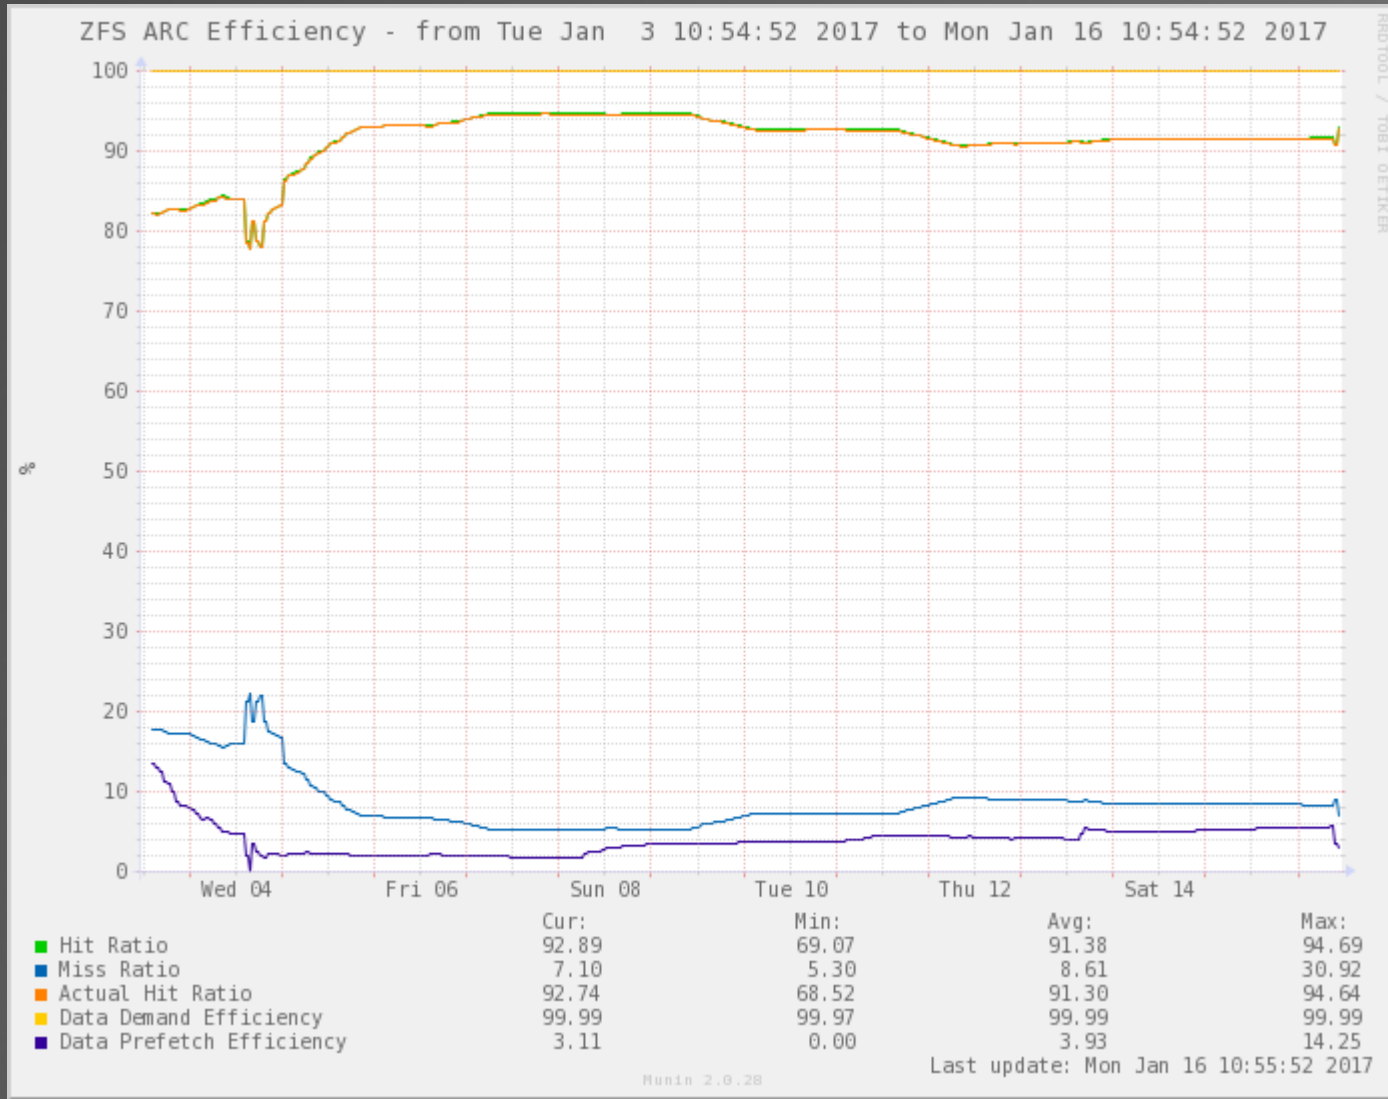
- Three of our DPM disknodes have storage fs on ZFS
- Used parameters
 - ashift=12
 - Enabled compression
 - Disabled suid, atime, devices
 - Roughly 95% quota (performance)
- Deduplication is not worth it



ZFS benefits

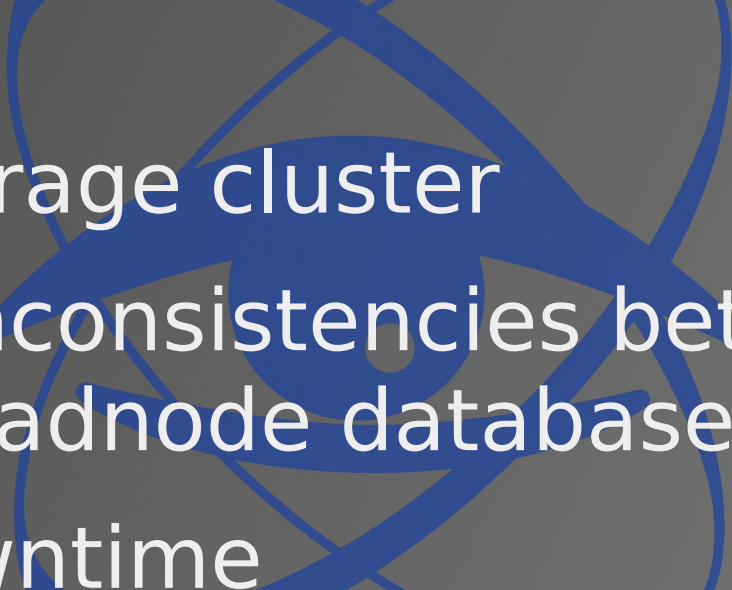
- Possible space savings with compression
 - 2% right now (approx 7% of space occupied)
- Data checksumming (consistency)
- Faster rebuilds (only on allocated data)
- Efficient cache (ARC)

ZFS benefits

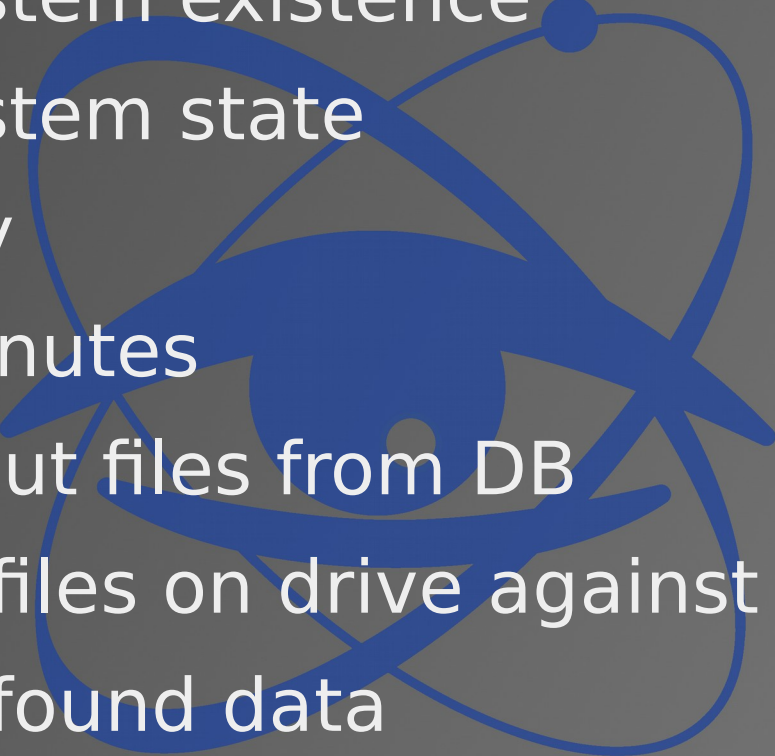


DPM cleaner

Dpm-cleaner is a:

- shell script
 - For DPM storage cluster
 - Searching inconsistencies between real data and headnode database
 - Without downtime
- 

DPM cleaner - How does it work?

1. Verify file system existence
 2. Verify file system state
 3. Set read-only
 4. Wait for 3 minutes
 5. Get data about files from DB
 6. Cross-check files on drive against DB data
 7. Remove not found data
 8. Set FS to state from stored in 2)
- 

Where to get it and contacts

- Github repository
 - <https://github.com/samuraiiii/dpm-cleaner>
- E-mail
 - mikula@fzu.cz
 - dpm-users-forum@cern.ch
 - Or again Github repo
 - PRs and bug reports are „welcome“ ;-)





Thank you for your attention