

Progress of resolving output data loss

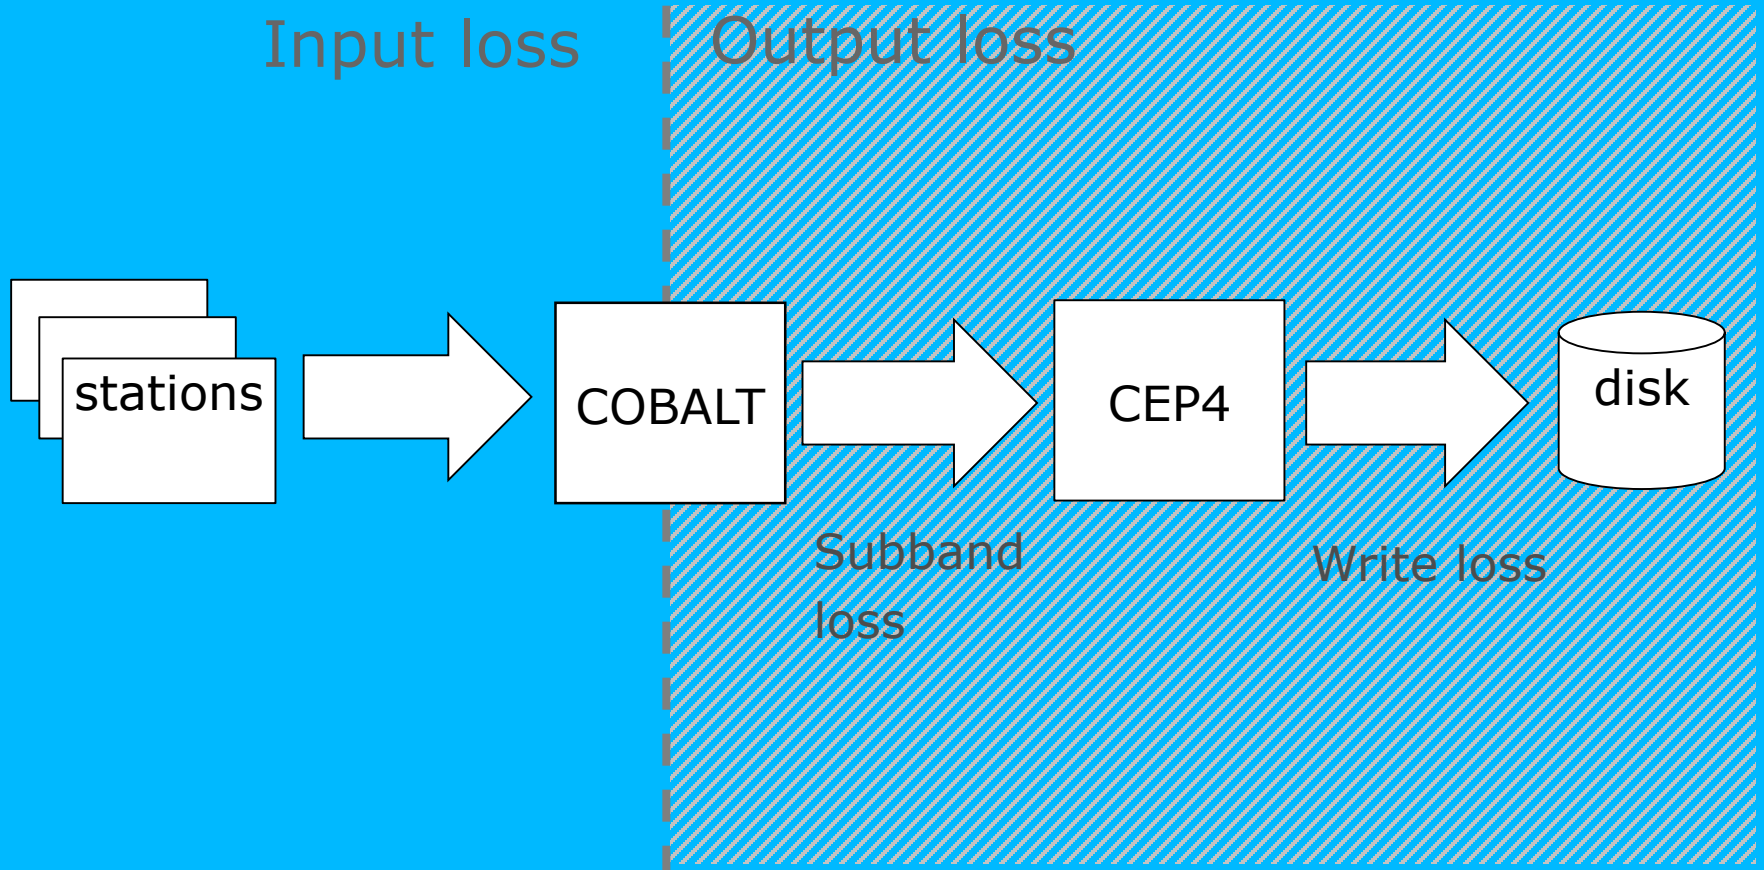
LSM 19-07-2017
Reinoud Bokhorst

About myself



- Reinoud Bokhorst, MsC
- LOFAR System Administrator & DBA in SDOS
- Started at ASTRON at June 15, 2016
- Background in Physical Oceanography and Meteorology
- Experience in system maintenance and development of (weather forecasting) production systems and products.

Data loss overview



Data loss symptoms

- Odd size data files (dropped blocks)
- Missing data files
- Zero-filled blocks in data files

L600057

- [Projects](#)

[Cobalt ERROR log](#)

Max file sizes (MB):

Correlated data : 1724
Beamformed data : 0

All data sets are there

Odd size data:

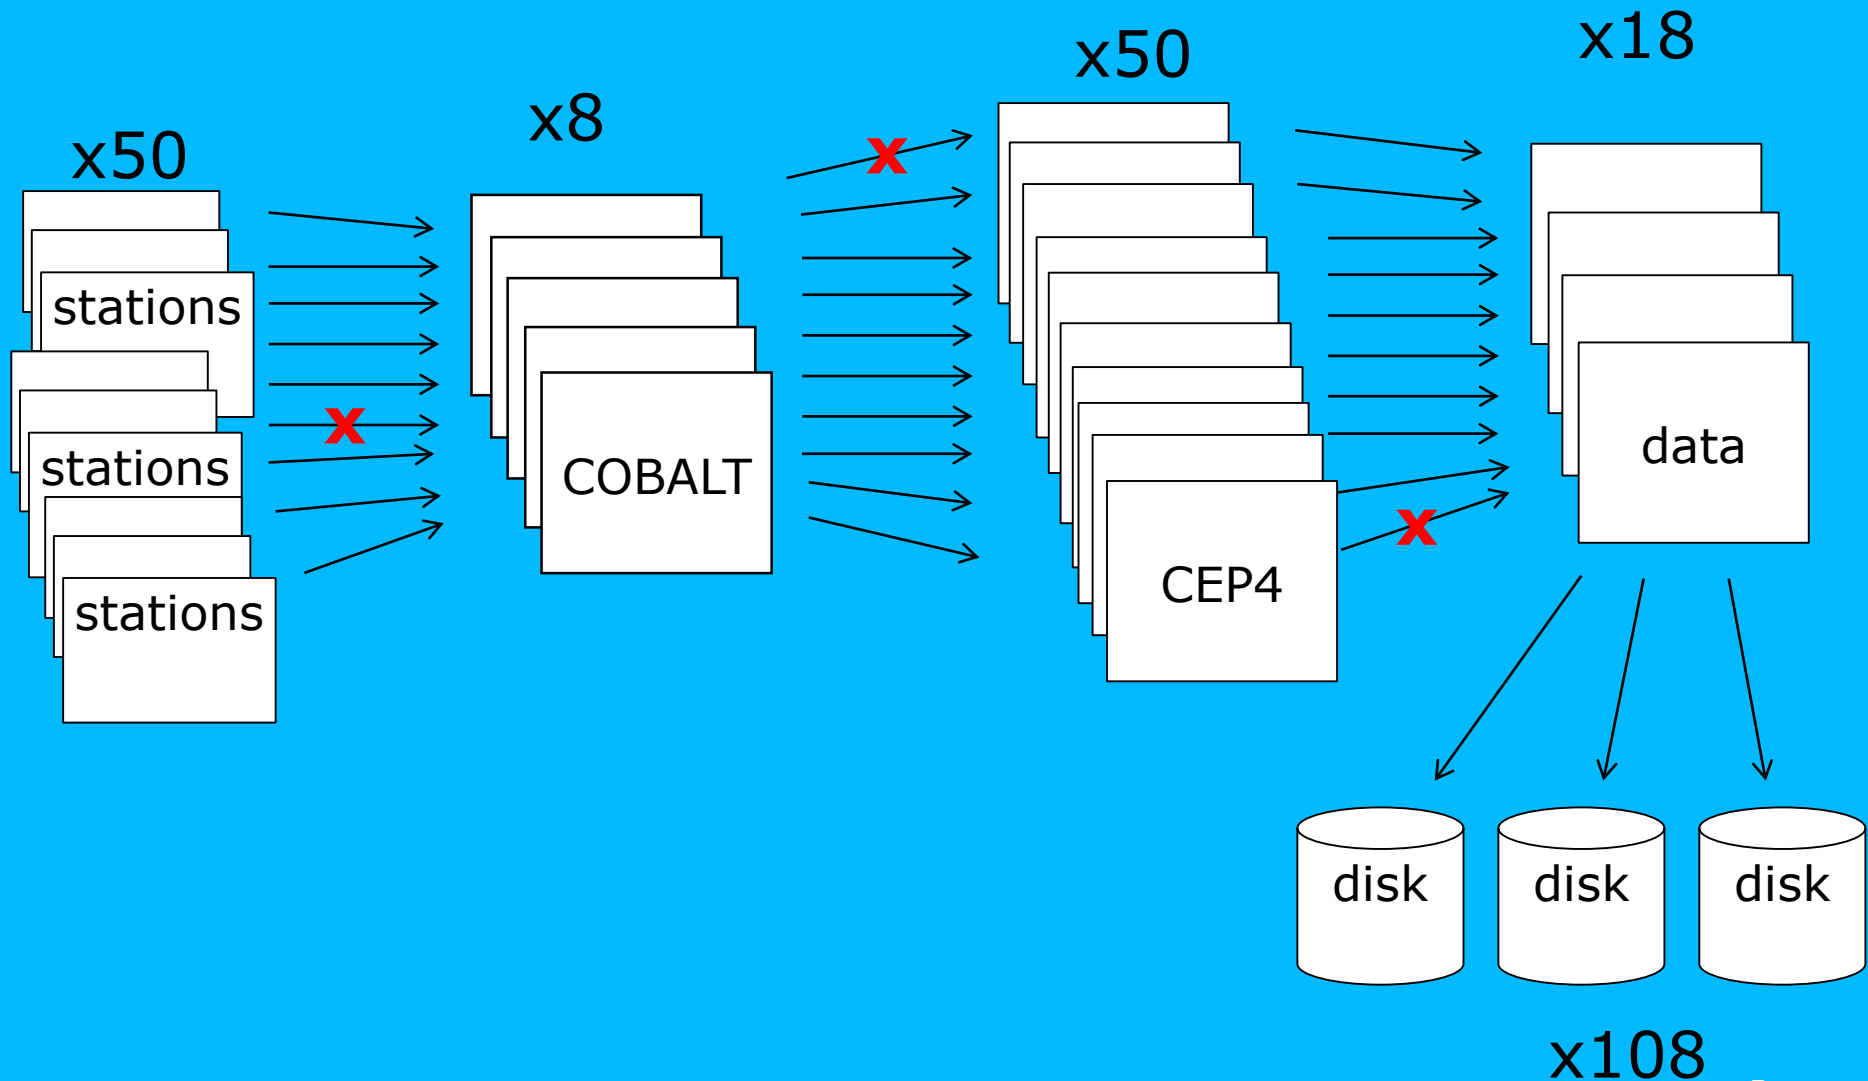
CEP4:L600057_SAP001_SB015_uv.MS : 1533 MB (89.0%)
CEP4:L600057_SAP001_SB060_uv.MS : 1539 MB (89.3%)
CEP4:L600057_SAP001_SB105_uv.MS : 1510 MB (87.6%)
CEP4:L600057_SAP001_SB119_uv.MS : 1712 MB (99.3%)
CEP4:L600057_SAP001_SB150_uv.MS : 1505 MB (87.3%)
CEP4:L600057_SAP001_SB195_uv.MS : 1580 MB (91.6%)
CEP4:L600057_SAP001_SB240_uv.MS : 1585 MB (92.0%)
CEP4:L600057_SAP002_SB254_uv.MS : 1695 MB (98.3%)
CEP4:L600057_SAP002_SB285_uv.MS : 1539 MB (89.3%)
CEP4:L600057_SAP002_SB299_uv.MS : 1706 MB (99.0%)
CEP4:L600057_SAP002_SB330_uv.MS : 1603 MB (93.0%)
CEP4:L600057_SAP002_SB344_uv.MS : 1706 MB (99.0%)
CEP4:L600057_SAP002_SB375_uv.MS : 1539 MB (89.3%)
CEP4:L600057_SAP002_SB389_uv.MS : 1706 MB (99.0%)
CEP4:L600057_SAP002_SB420_uv.MS : 1493 MB (86.6%)
CEP4:L600057_SAP002_SB465_uv.MS : 1551 MB (90.0%)

Input loss report



CS002HBA0: 0.0133%
CS002HBA1: 0.0324%
CS006HBA0: 0.0303%
CS006HBA1: 0.0445%
CS011HBA0: 0.0193%
CS011HBA1: 0.0049%
CS021HBA0: 0.0009%

Data loss overview: scale of the system



1. Registration of data loss cases for pattern discovery
2. Log file analysis
3. Monitoring of system parameters (Ganglia, Zabbix)
4. General knowledge building
5. Performance testing
6. Escalation to external resource persons

Iterative process..

Identified data loss causes



- Data loss typically occurred on CEP4 during high data-rate observations (beam formed), on nodes that were also loaded with pipeline processes => global file system? Network?
 - Some data loss cases typically happened during the initialization phase on CEP4, when large memory allocation is done
 - Some pipelines are 'heavy' on file operations, consuming Lustre resources
 - Global file system (Lustre) performance fell short under high load
 - Network connection COBALT to CEP4 slow
- ⇒ tuning required

- Tuning of Lustre was done.
18% read improvement, 0% write improvement.
Does not solve the problem alone
- Network connection COBALT-CEP4 is currently being fixed
- Alternative strategy has been tested:
=> Split CEP4 nodes into separate data writer nodes and pipeline processing nodes

Tests performed last week;

- 10 data writer nodes sufficient for high data-rate observations (beam formed)
- Still data loss for u/v data observations

Other obvious data loss causes

System instabilities:

- CEP4 nodes: operating system crashes (Sep - Oct 2016, **fixed**)
- COBALT operating system crashes (Feb - Apr 2017, **fixed**)
- CEP4 nodes: operating system crashes (May 2017, **fixed**)
- *Station network connection from Poland (June 2017, **fixed**)*
- COBALT operating system crashes (June 2017 - **now**)

Plans for the coming period



- Fix network connection between COBALT and CEP4
- Debug COBALT crashes
- Repeat dedicated data writer node test after network connection has been fixed.

1. Make COBALT software more robust for node crashes, observation can continue after reboot
2. Fix memory allocation problem on start-up of an observation (CEP4)
3. Improve general pipeline file handling, decrease number of file operations.