

Symptoms

- Hardware problems: ~10 disc failures on 3-4 machines since November
- HDFS warnings: two blocks with the same block ID - spreading
- A crash (null pointer exception) in the block manager
- Excessive I/O run times - up to 8 minutes for a "write"
- Hbase crashes on nodes put into "drain"
- balancer not working
- draining nodes causes client failures

Treatment

- 4 modern disc server nodes borrowed
- suspect nodes set to drain
- debugged and by-passed hbase errors
- produced list of files lost
- monitor consolidation of remaining data (luckily over 99.99%)

Status

- Cluster is running, throttled, but usable
- broken/wrotten/corrupt hdfs files should be deleted
- hbase is running, needs administrative action
- rebuild ongoing, estimate is 1-2 days
- failure analysis ongoing