


# Overview of Storage Operations at CERN

*Or (better)*

# CERN IT file-based physics data storage operations

M.Lamanna, J.Iven; CERN IT DSS-FDO

CERN IT Department
CH-1211 Geneva 23
Switzerland

(other “CERN storage operations”, not presented here)

- experiment storage
- DFS,
- DBs

- Other IT-DSS activities:
  - AFS – see *Scaling the AFS service at CERN*
  - CASTOR tape+dev, EOS dev
  - Projects:
    - OpenStack, HADOOP cluster, single-replica disk cache, …
    - Huawei SingleStorage
      - “cloud” interface on custom (cheap) hardware

HUAWEI

CERN openlab

CERN



# Recap: CASTOR and EOS



## CASTOR2 = HSM

CASTOR
CERN Advanced STORage manager

- Successor of SHIFT (90s) and CASTOR-1 (99)
  - Related to DPM, LFC
  - DB-centric architecture
  - Feature-rich:
    - tape pools, disk pools, service classes, instances, file classes, file replication, scheduled transfers (etc)

## EOS

- Namespace “plugin” to xroot
  - In-memory: O(ms) latency
  - Redundant file copies
  - quotas

# CASTOR architecture (simplified)



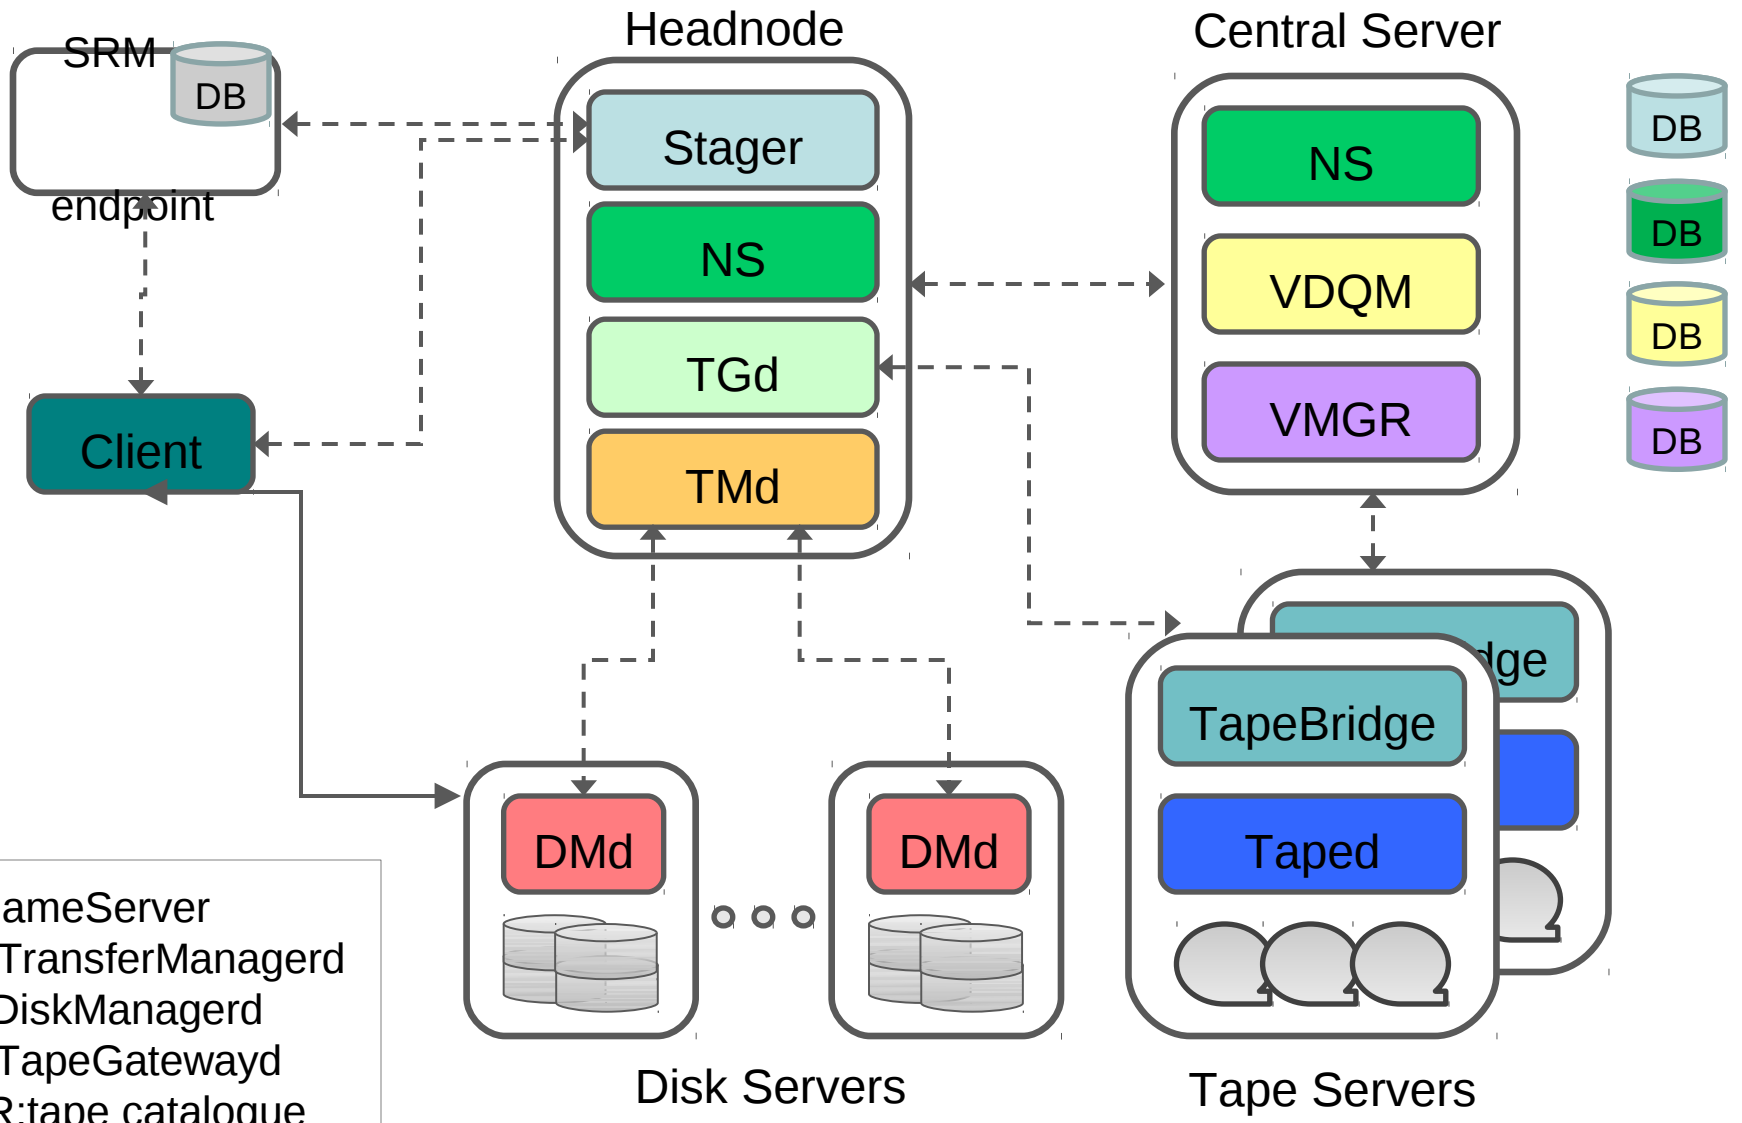


NS: NameServer
TMd: TransferManagerd
DMd:DiskManagerd
TGd: TapeGatewayd
VMGR:tape catalogue
VDQM: drive scheduler



(credit: L.Mascetti)

# EOS architecture



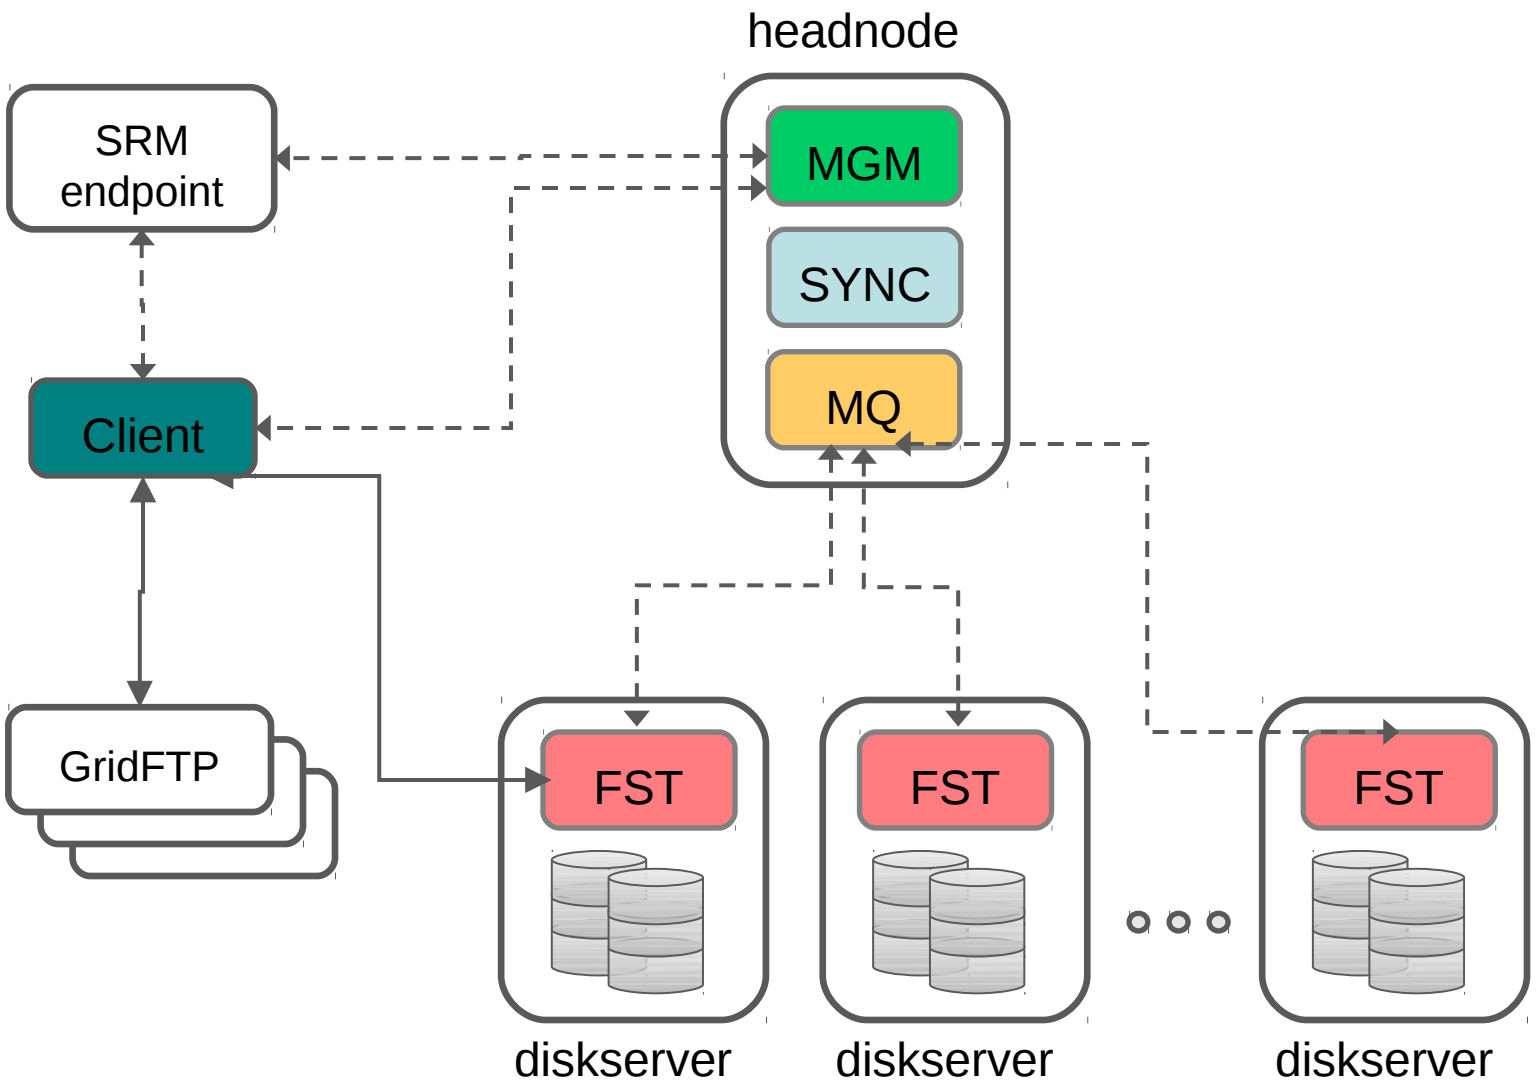




(credit: L.Mascetti)

# Why 2 - HSM vs managed



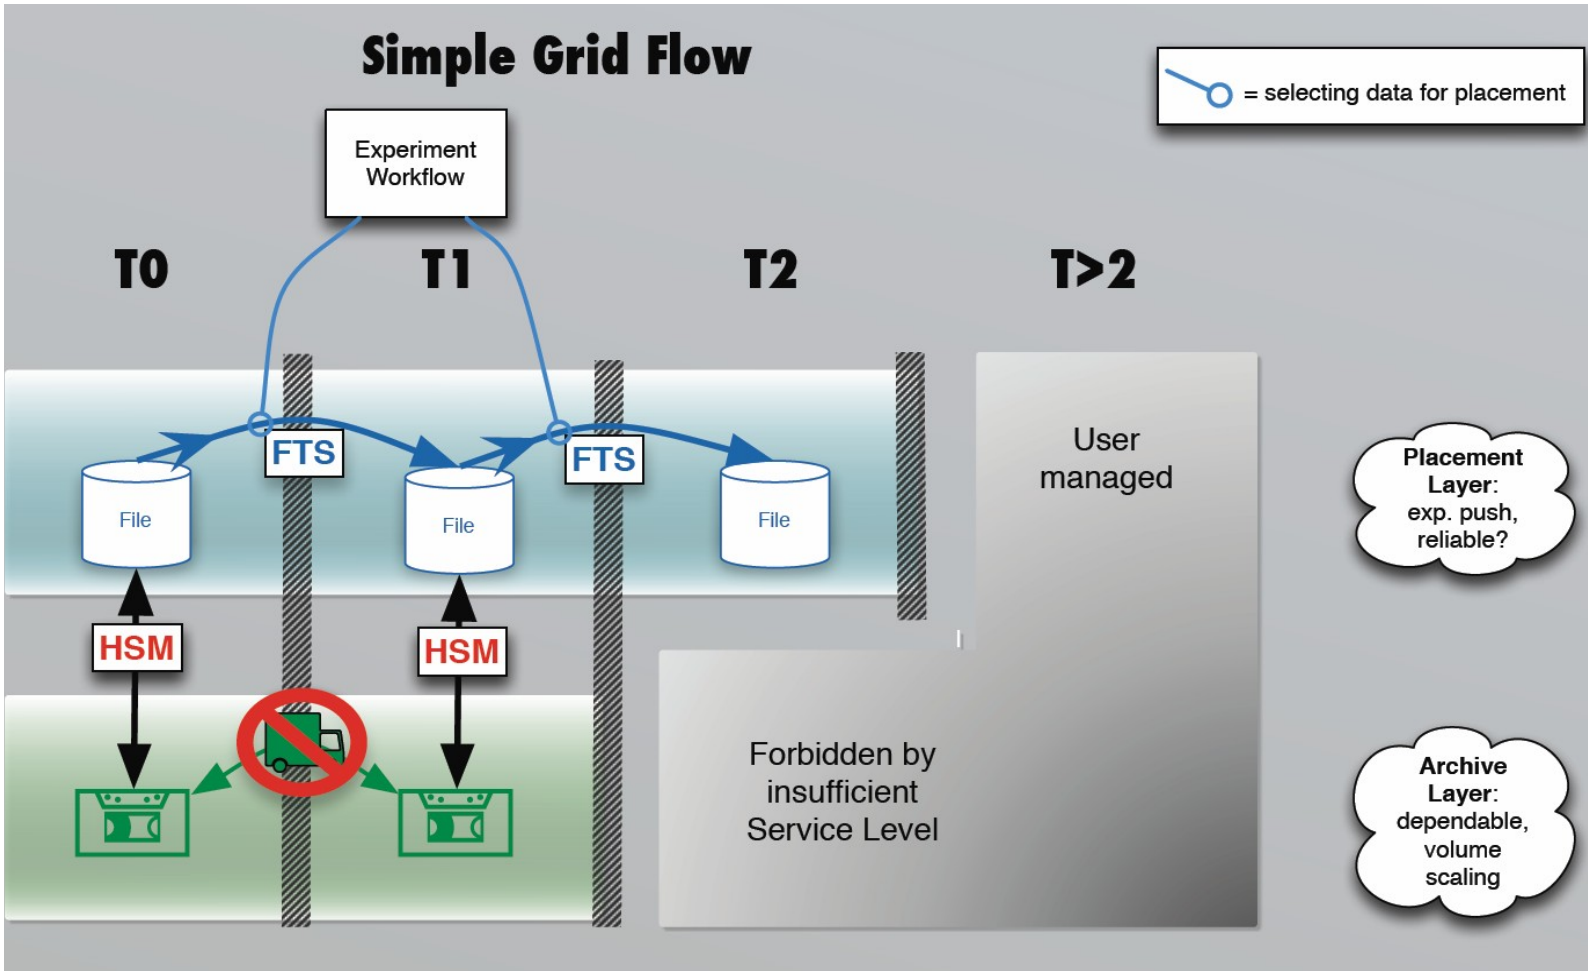


D.Duellmann from DM & SM TEG meeting 24.Feb.2012



(credit: L.Mascetti)

# Less HSM, more explicit transfers



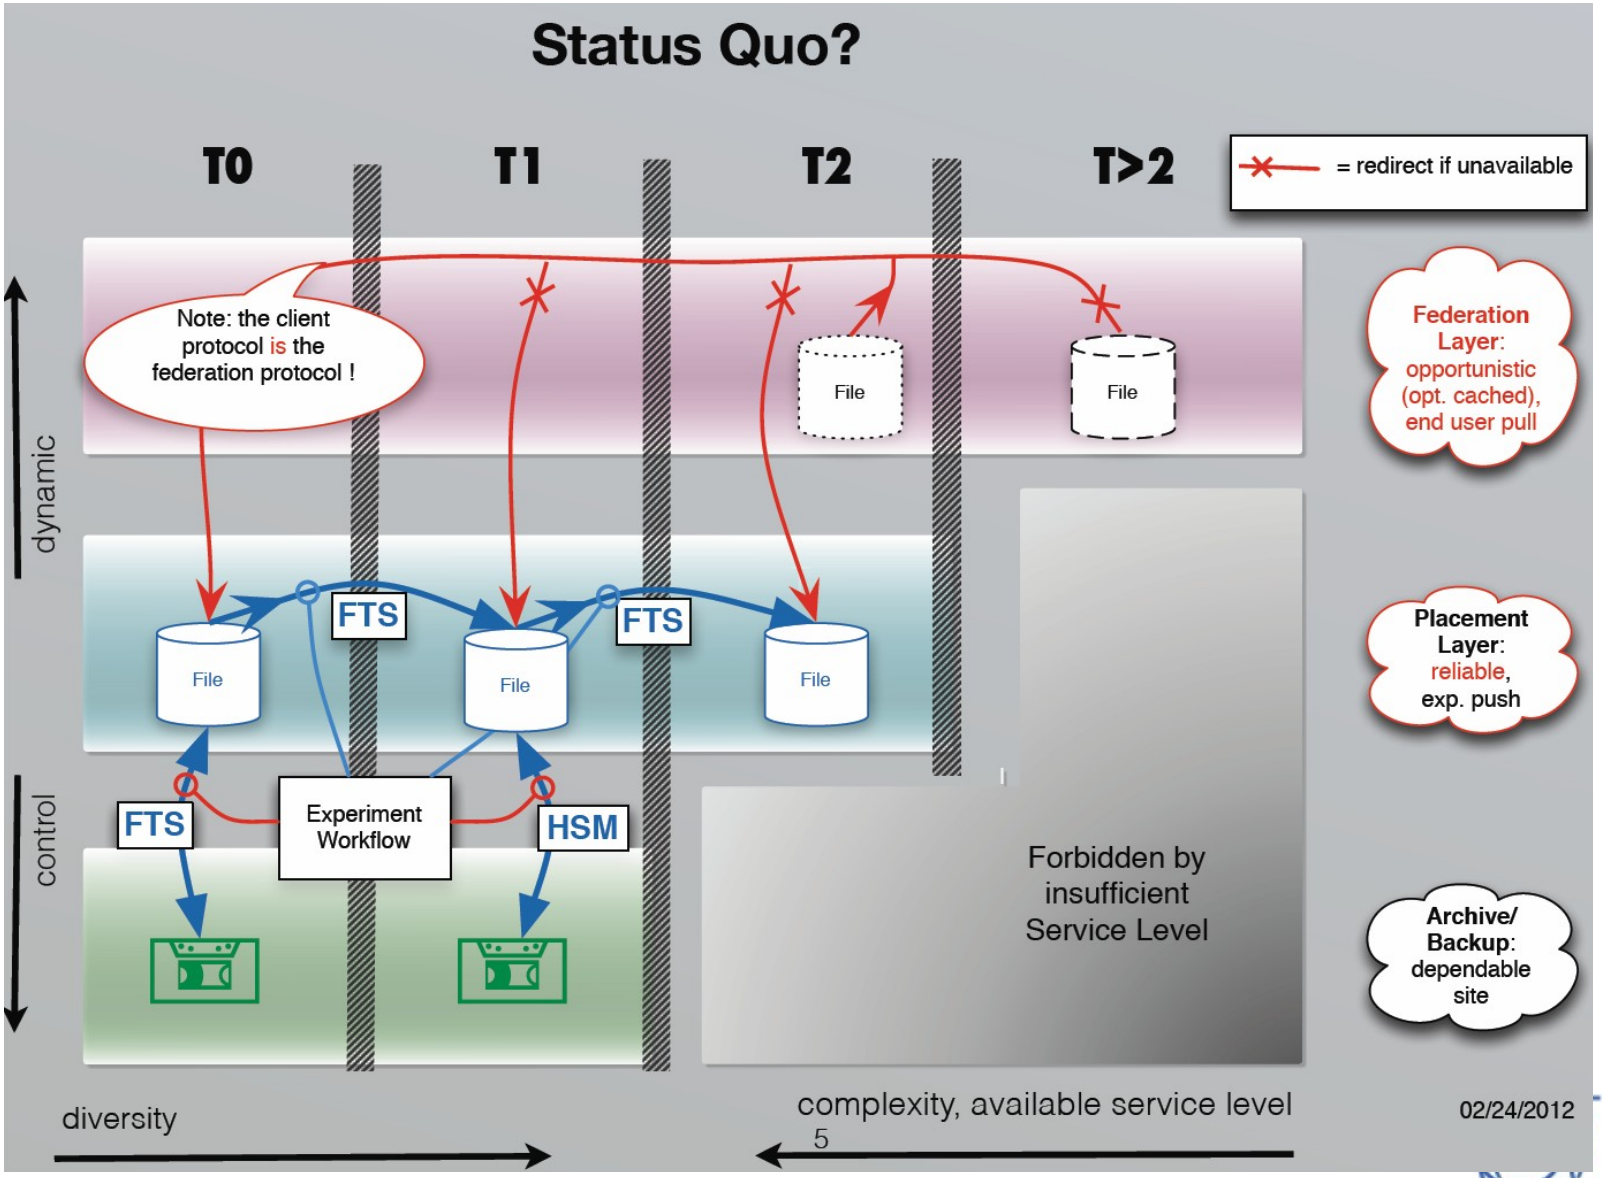


D.Duellmann from DM & SM TEG meeting 24.Feb.2012



# Why 2?



- CASTOR is geared towards Tier0 flow
  - CERN IT primary responsibility
  - DB-architecture allows easy horizontal scaling
    - Visible e.g. on SRM / high availability
  - Up to 2010: supporting all use cases
- New ingredients:
  - strong increase in “random” analysis stretches it
  - Our and experiment operational experience

*Our bet:* 2 targeted systems are easier than 1 do-it-all

  - New system cuts legacy cruft & patterns
  - Each use case has its system

# Overall numbers



- CASTOR (disk)
  - 5 instances (prod)
  - 1280 diskservers
  - 17PB (usable)
- EOS
  - 3 instances
  - 400 diskservers
  - 8.5PB (usable)

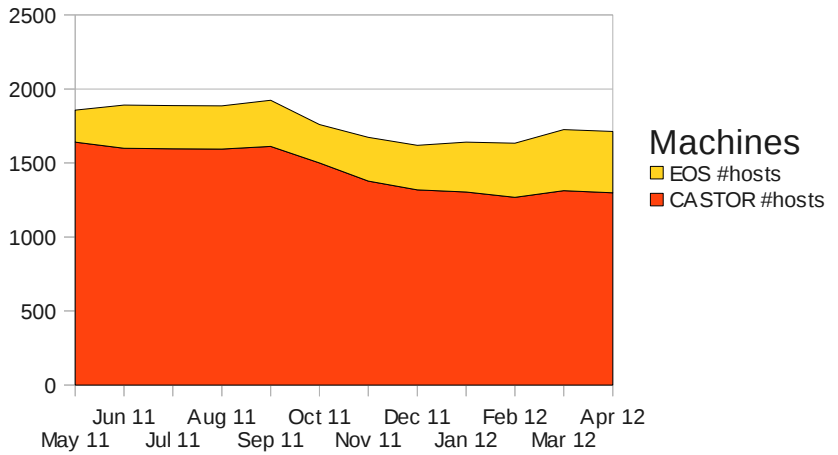


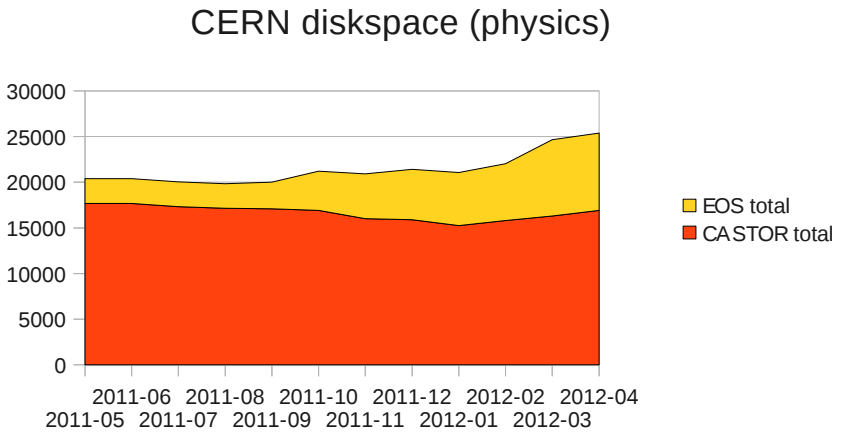

# experiment usage – ALICE



- CASTOR:
  - 1 tape-backed pool, 1 diskonly
  - Nice datarates during HI test (can saturate 4GB/s from pit, >12GB/s pool-internal)
- EOS (recent = small, 500TB)
  - expect to move CASTOR diskpool into EOS in 2012
- xroot-centric

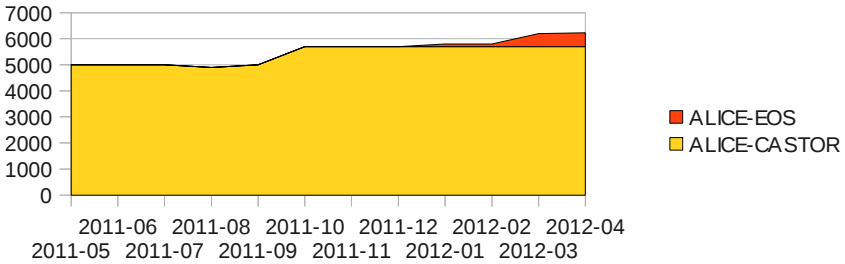


163 nodes: 3.3GB/s

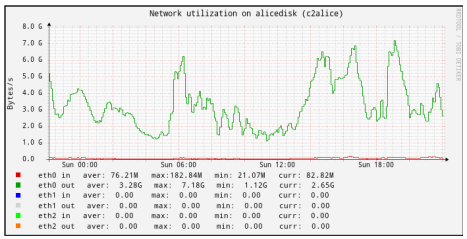


18 nodes: 2.3GB/s

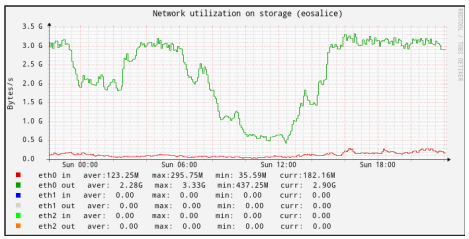

# experiment usage – ATLAS



- Restricted user tape access on CASTOR
- Consolidated 5 CASTOR pools into EOS
  - 3 more on the way out..
  - CASTOR CERNT3 instance gone/recycled
- RFIO- and SRM-less: use **xrdcp** (where possible)

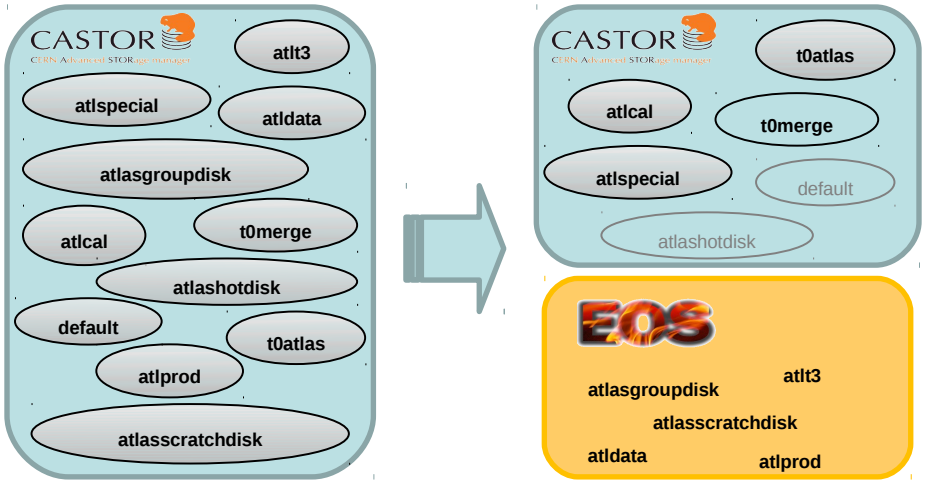


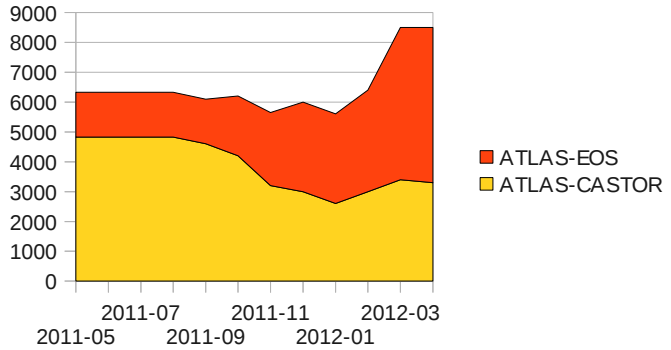


(credit: L.Mascetti)

# experiment usage – CMS



- 7 CASTOR pools moved into EOS (8 left)
- Extensive use of EOS quotas
  - user management done by CMS, own tools
- PhEDEx uses xrd3cp
- Looking at Xrootd federations

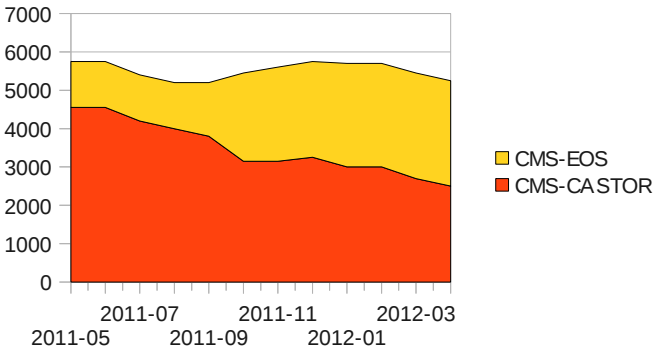

# experiment usage – LHCb



- Major CASTOR disk pool consolidation in 2011
  - only 5 pools left, major: *lhcbdisk* and *lhcbtape*
- Still using mostly SRM
- Do not want EOS (yet)

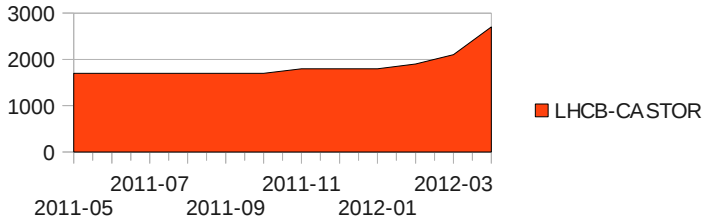

# experiment usage – non-LHC



- CASTORPUBLIC: 2.7PB, 16 pools 26TB..500TB
  - New: AMS 2 pools, ILC (heavy SRM usage)
  - COMPASS, NA48 could consolidate

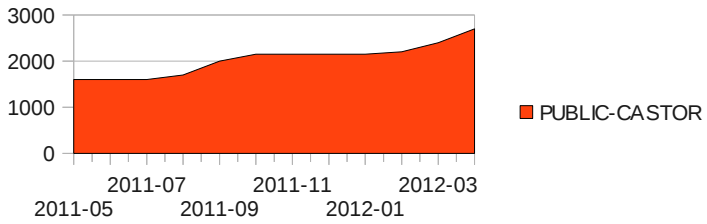


- EOSPUBLIC: under discussion
  - needs resource commitments from experiments
  - Several interested



# 2011 CASTOR changes



- LSF → Transfermanager:
  - LSF request scheduler: complex
    - but CASTOR did not exploit full capabilities, just “slots”
    - Scheduling delay >1sec; max rate ~20Hz; plugin “meltdown” if queue was big
  - Rewritten:
    - still slot-based, but greedy random allocation
    - >200Hz rate (throttled to 75Hz)
    - can cope with huge queues, stateless
- “tape gateway”: revamped tape migration
  - (and cleaned up tape pools)
  - Some teething trouble

DSS

# CASTOR Tape highlights

CERN IT Department

- 62PB on tape, 52K tapes, 9 libraries, 80 production drives (+20 legacy)
  - Beta-tested, validated and deployed IBM TS1140 (4TB) and Oracle T10000C (5TB) drives
- Boosted write tape speed writing by developing and deploying “buffered” tape marks (avoiding head repositioning) → factor 10x achieved in 1 year
- Introduced “traffic lights” and “bus lanes” for prioritising bulk read requests, reducing tape mounts by ~50%
- Investigating suitability of “commodity” equipment (aka LTO)
- Active verification of archive contents by re-reading tapes and comparing checksums
  - All newly filled tapes
  - “dusty” (not recently mounted) tapes

*Cf posters: 415 (S. Murray) and 247 (G. Cancio)*



CERN

# EOS vs CASTOR – operational comparison

CERN IT Department

- HW cost – similar (for disk layer)
  - Same hardware
    - CASTOR: RAID-1, EOS: JBOD (same replication)
    - Similar: 10% overhead (#nodes): headnode/SRM/dev
  - CASTOR: ~27CHF/TBmonth, EOS: ~17CHF/TBmonth (=newer HW)
- Setup effort:
  - EOS is more integrated with CERN Quattor infrastructure = new diskservers just pop up
  - CASTOR needs bespoke scripts (changing)
- SW updates – both RPM
  - CASTOR : +DB-side updates


# operational comparison (2)



- Daily operations:
  - EOS “draining” is more advanced
    - (Needs to! Disk “drains” are done by HW RAID on CASTOR)
    - Draining a whole machine is faster on EOS
    - But: both need manual action for leftovers
  - EOS has automatic space rebalancing (and more robust space allocation)
  - Debugging:
    - CASTOR has DLF : useful but fragile (DB jobs failing..); being rewritten
    - EOS has nothing (are considering SPLUNK, but daily volume is ~30GB/day = $$$). “grep” works OK, though.
    - EOS is noticeably less complex

→ EOS “feels” easier to operate.


# Support effort



- Shared (disk) support team
  - ~1.5 FTE EOS, ~3.3 FTE CASTOR
- Savannah (2011-01-01 .. now)
  - CASTOR:
    - 345 “support” (dev)
    - 298 “bug” (dev; incl RFE)
    - 364 “task” (operations)
  - EOS
    - 39 “support”
    - 226 “bug”
    - 116 “task”

GGUS alarms – unchanged:

| Error type | 2010 | 2011+Q1 2012 |
|---|---|---|
| Software bug+DB | 32% | 46% |
| Hardware failure | 19% | 21% |
| Config error | 17% | 11% |
| Human error | 8% | 11% |
| Overload | 8% | 11% |

Low absolute numbers: 28 in 2011

# Support: Hardware tickets



- CASTOR ~ EOS
  - Same basic HW sensors
  - EOS: combined ”replace disk” alarm
- (EOS spike is new alarms)
- CASTOR has Xmas break

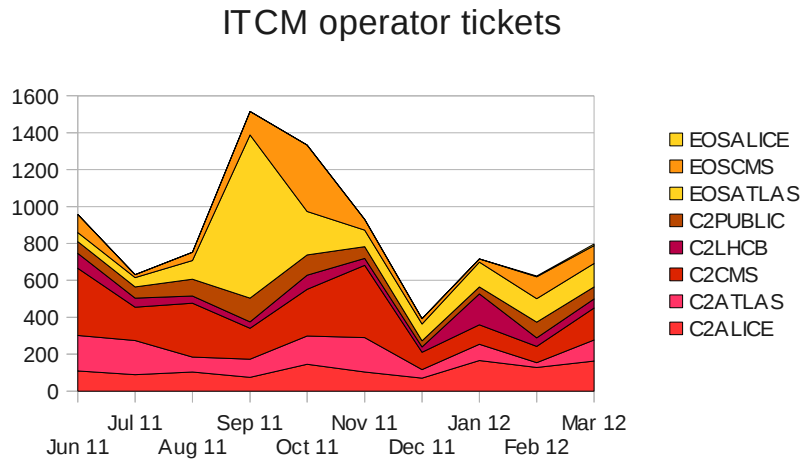


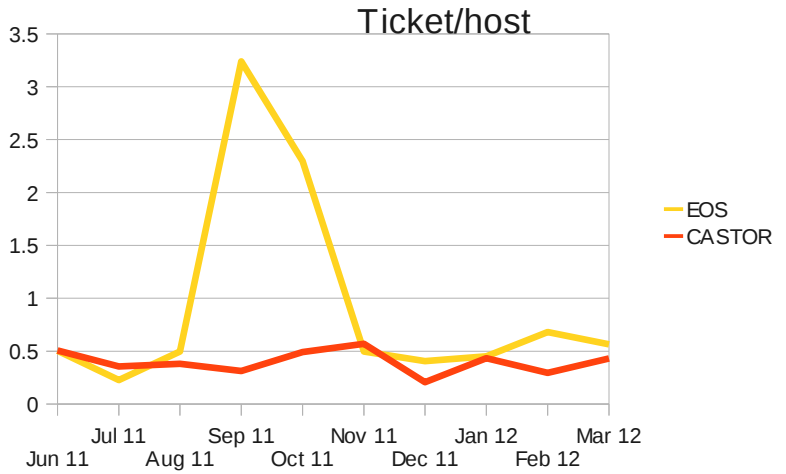

# Support: CASTOR tickets



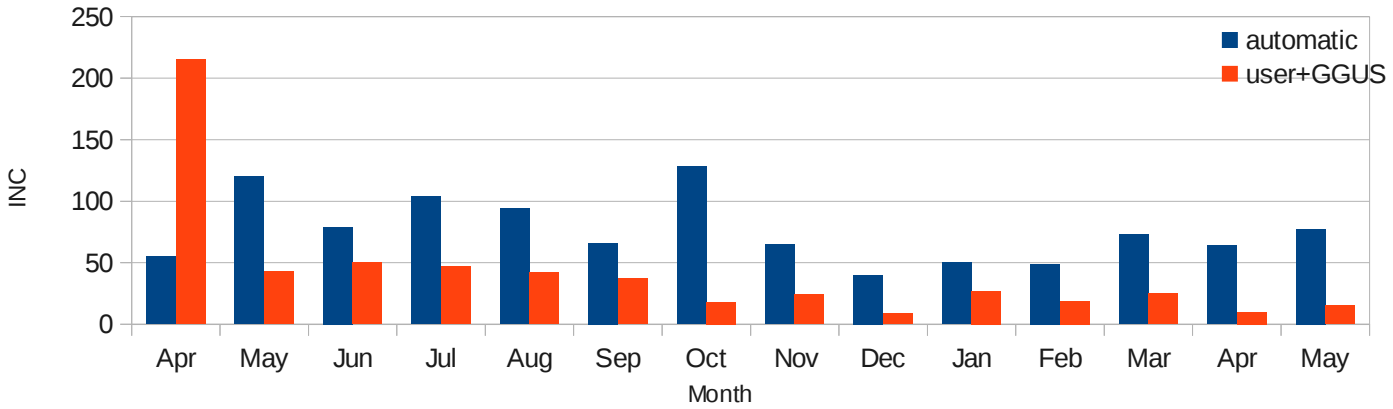


- (all CASTOR instances, all support lines, some assignment errors; spike in Apr 2011 is new ticket system)
- Internal tickets dominate
  - one underlying HW issue can create several tickets
  - Automatic tickets for tape-related problems
- Note: user support+GGUS decrease since Sep 2011

# Support (cont)



- EOS user support (still) largely handled by experiments – not on graph
- CASTOR GGUS has higher priority (T0 data)
- CASTOR (disk) + EOS support structure:

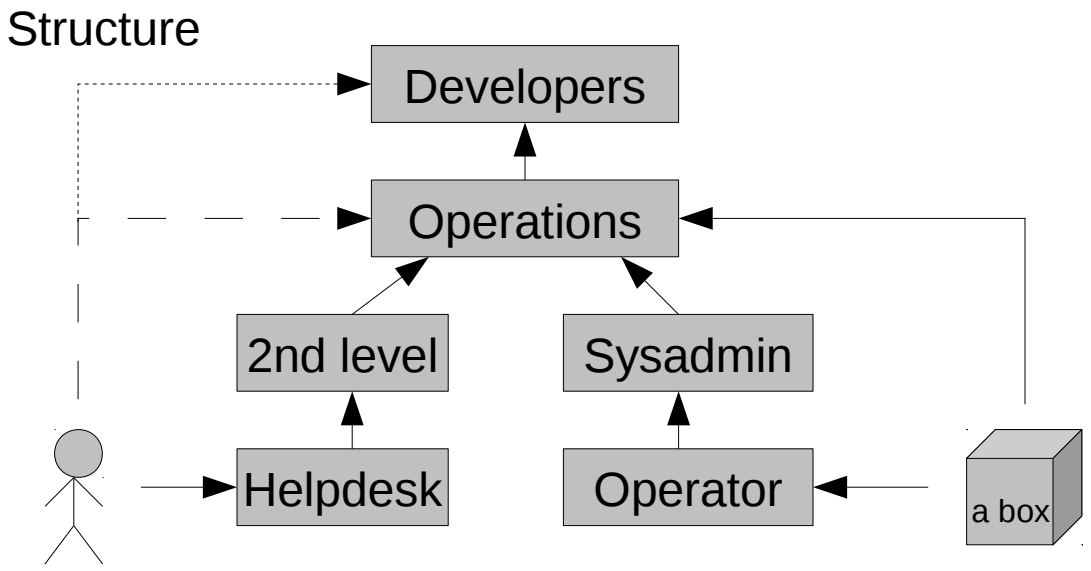


# Upcoming things



- CASTOR development
  - rewrite “tape recall” handling
  - namespace+stager protocols
  - RFIO → xrootd for internal transfers
  - Not: remove disk-only support
- EOS development – lots...
  - Faster consistency checks (FSCK)
  - HA for namespace – redo
  - Namespace storage alternatives
  - Replace Message queues
  - block based redundancy → < 2 replicas
- Xrootd federations (at least for EOS)


# Resources



CASTOR: http://cern.ch/castor

EOS: http://cern.ch/eos (service information)

http://eos.cern.ch (code – no support!)