



DSS

Evolving the CERN AFS Infrastructure

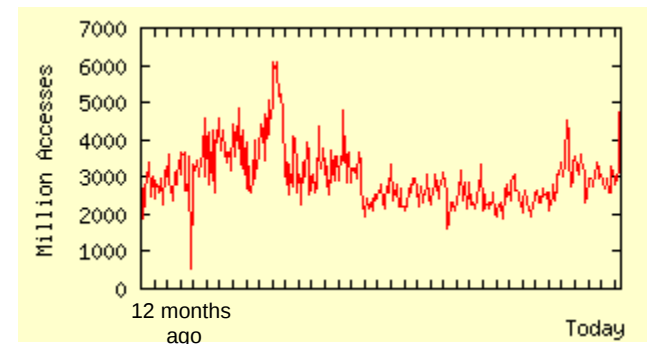
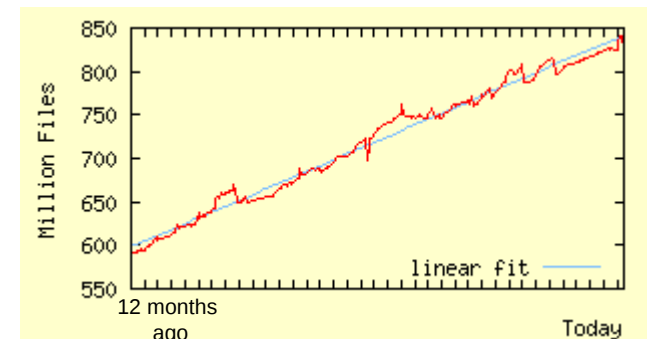
Arne Wiebalck

European AFS and Kerberos Workshop
DESY, 5th Oct 2011



Site Report: Numbers

- Service provides networked storage to CERN users
 - >30'000 home directories & ~300 project spaces
 - high availability, daily backup, security, access control, quotas, monitoring, operability, ...
- Service key data
 - ~55 file servers
 - ~850 million files (+250 m/yr)
 - ~55TB of data, ~100TB quota
 - ~65'000 volumes
- Service activity
 - 10'000 CERN clients
 - 5'000 off-site clients
 - 5'000 active users/week
 - ~3 billion accesses/day
 - ~300 million reads and writes/day

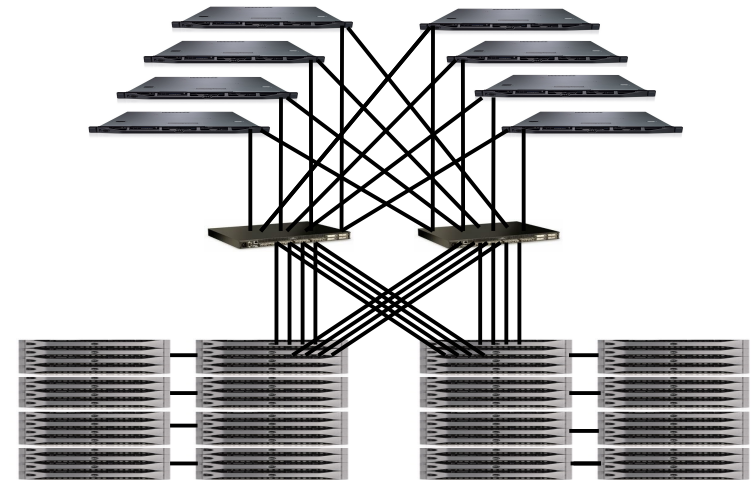




Site Report: Hardware

➤ Fibre channel fabrics

- some 45 servers
- ~80 FC disk trays
- 6 fabrics
- 300GB SAS disks
- home dirs, projects



➤ Standard disk servers

- some 10 servers
- directly attached disks
- s/w RAID over iSCSI
- scratch data





- Current production version is openafs-1.4.14 ...
 - both client and servers
 - since February 2011
 - no major issues

- ... plus a set of CERN-only patches
 - on-demand attach
 - handling of clients behind NATs
 - request throttling
 - remote network statistics
 - forced offline
 - client shutdown
 - ...

- openafs-1.6.0 is in the queue

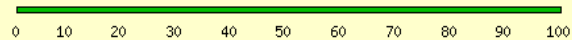


DSS

Site Report: Monitoring

CERN AFS Console ?

AFS Availability



AFS Available Performance



AFS Console's Promptness

Partition related information is from 2011-09-29/10:57:01
Volume related information is from 2011-09-29/10:37:01

AFS Alarms / Warnings

afs156 has a local response time for 64Kb of 160.15 ms (avg last hour)
[List hottest volumes](#)

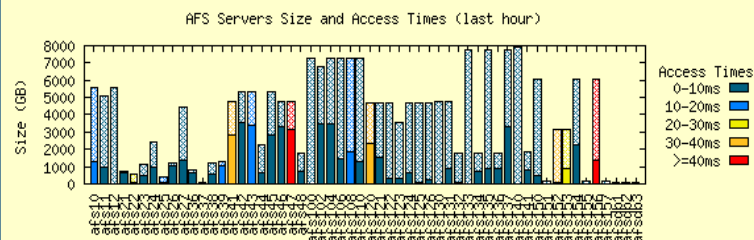
AFS Statistics

[Lemon](#)
[SLS](#)

Report Generator

[Volumes Report](#) [by server](#) [by project](#) [global](#)
[Servers Report](#)
[Partitions Report](#)
[Pool Report](#)

Servers Overview



Projects Overview

[atlas](#) [alice](#) [cms](#) [lhcb](#) [compass](#) [gd](#) [swicg](#) [user](#) [sw](#) [afs](#)

AFS Historic Monitoring Data

[Volume Statistics / Volume History](#)
[Partition Statistics / Partition History](#)

Service Incidents

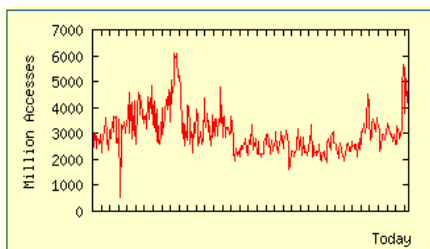
[Get Incidents Log](#)

AFS CERN Cell Logs

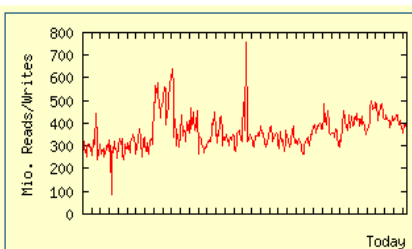
[Simple entry Log...](#)
[Complete Log...](#)

Get Log from date:

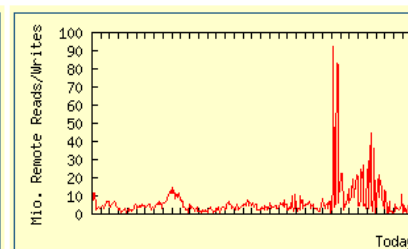
AFS Total Accesses (Last 365 days)



AFS Total Reads+Writes (Last 365 days)



AFS Remote Reads+Writes (Last 365 days)





- AFS Console: vos commands plus DB backend
 - sensors for volumes, partitions (and servers)
 - MySQL (new)
 - Web interface
 - data kept for 6+ months (averaging)
- Debugging
 - partition response times, hot volumes
 - also postmortem
- Service self-adjustment
 - initial volume placement
 - pool_monitor
 - readonly-monitor
- Trending
 - space, accesses



- Kerberos unification under MS AD
 - see John's talk tomorrow
- Introduction of ACL policy
 - information campaign
 - tools (for us and the users)
 - handling of special cases
 - ~20% of home dir changed (now <1%)
- Integration with MS Forefront Identity Manager
 - account creation/deletion
 - quotas
 - - passwords



- More space
 - 1GB home directories
 - some GB of work/scratch space
- More IOPs
 - access is always too slow
- Maintain service quality
 - availability
 - features (e.g. daily backup)
- Service simplification
 - backup, requests handling, space types, ...
- Service clarification
 - quotas, retention policies, incident handling, ...



- Ignore, do nothing or continue to grow slowly
 - users will go elsewhere
- Provide something else
 - What would that “something else” be?
 - “AFS sails on quietly and efficiently ...”
- Simply increase what we have done so far
 - technically possible, but expensive

Task: Increase the scalability of the AFS service



DSS Setup Review

- + Homogenous
 - simplifies operations
- + Reliable
 - hardware, dualpath
 - interventions do work
- + Sufficient performance
 - high-end, small SAS disks
- Non-standard hardware at CERN
 - procurement & metadata, testing & burn-in, tools (console)
- Complicated
 - initial setup
 - stress situations
- Price/GB
 - up to an order of magnitude more expensive
 - prevents easy scaling





Approach: SAS-based Storage Units

➤ Hardware Setup

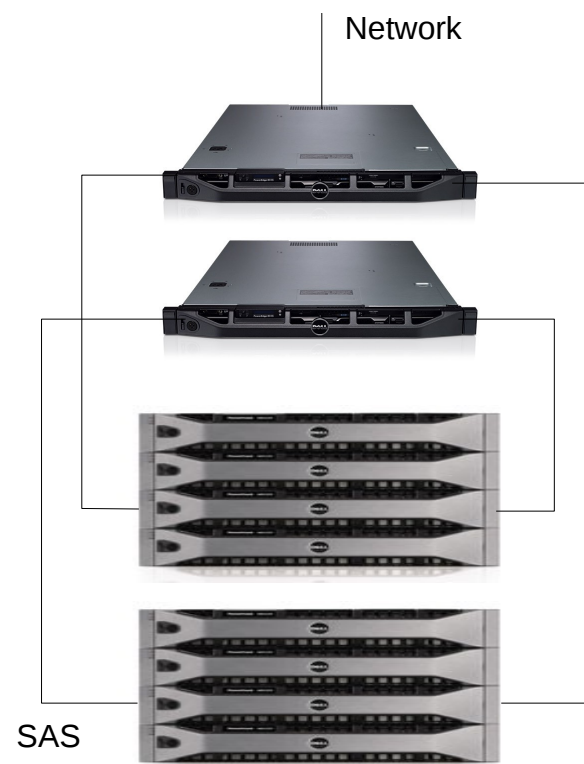
- 2 servers and 2 trays form a “unit”
- all disks visible on both servers
- 16x 2TB NLSAS, 4x 256 GB SSDs

➤ Reliability

- JBODs (no h/w raid controllers)
- s/w RAID across arrays
- “volume take-over”

➤ Performance

- make use of SSDs to compensate larger disks: FACEBOOK's flashcache



Storage Unit (SU)



- General purpose block cache for the Linux kernel

- developed by facebook for MySQL
- integrated into device mapper
- supports writeback/writethrough



```
# dmsetup status
...
vicepcd: 0 585937224 flashcache-wt stats:
      reads(396941574), writes(153485540)
      cache hits(293703891), cache hit percent (73)
      ...
      disk reads(103238560), ...
      ...
```

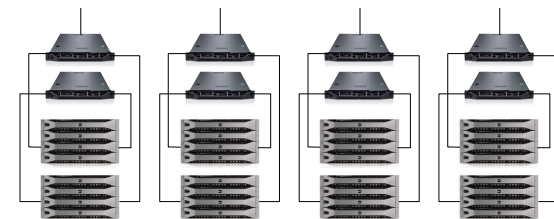
- In production since early this year
- Can give significant speed-up

<https://github.com/facebook/flashcache>



DSS SU Advantages

- **Lower Price/GB**
 - even with 5% SSDs we gain a factor 4
- **Data availability**
 - no SPOF (same as before)
- **Simpler setup**
 - limit the impact of finger trouble
 - ease things in stress situations
- **Performance**
 - SSD/flashcache shall compensate for larger disks
- **Streamlined hardware**
 - profit from well-established procedures (procurement, testing, burn-in, support)





DSS

What our users will get ...

- 10GB home directories (10x more)
- 100GB work spaces
- SSD read caching
- daily backups (shorter retention period)
- streamlined interfaces & procedures



Option 1: Both servers of the SU are active	Option 2: Only one server in an SU is active
+ both servers contribute to file serving memory, network, CPU	– one server “wasted” mitigated by 2 different servers?
+ impact by file server problem is less severe	– large servers backup, flexibility, cross-talk
– “joining” the data on one server is difficult: no sysid copy possible syncserv slow (1/sec, no bulk!) syncserv not (easily) possible when the other server is completely gone	+ “volume take-over” fast via sysid incidents/maintenance
– split after join?	+ switch back and forth possible
	+ understand limits add capacity if needed



- How good is the new hardware?
 - failure rate
 - incident procedures
- Can we backup these servers efficiently?
 - some 10 TB/server
 - depends on data change rate
 - backup setup under review
- How do users change their behavior?
 - more space, new use cases
 - separation of home dirs and work spaces
- Too few servers in the end?
 - ratio of data volume to #servers will grow
 - less flexibility, more contention?
 - isolation servers required?



- AFS cornerstone of data services at CERN
- User demands drive a change and the service needs to adapt
 - from FC fabrics to external SAS-based storage units
 - from small expensive to large mainstream disks
 - compensate performance penalty with SSDs
- New hardware arrives next week ...