# UKI-NORTHGRID-SHEF-HEP: Storage system and performance in STEP09

Elena Korolkova
Matt Robinson
Paul Hodgson

# Sheffield LCG cluster storage

- SL 4.6
- DPM  1.6 11
- xfs
- No RAID controllers
- Software RAID: RAID5 and RAID6

# Sheffield LCG cluster: storage system

**SE0 DPM HeadNode+Disk Pool 13 TB**

2 X Quad-Core 2.0 GHz AMD Opteron processors 2347
Redundant Power Supply
8 GB RAM
20 X 1 TB Disks
RAID 5

**SE1 DPM Disk Pool 13 TB**

2XDual -Core 2.0 GHz AMD Opteron processors 2212
Redundant Power Supply
4 GB RAM
20 X 1 TB Disks
RAID 5

**SE2 DPM Disk Pool 13 TB**

2XDual-Core 2.0 GHz AMD Opteron processors 2212
Redundant Power Supply
4 GB RAM
20 X 1 TB Disks
RAID 6

# Sheffield LCG cluster: Network

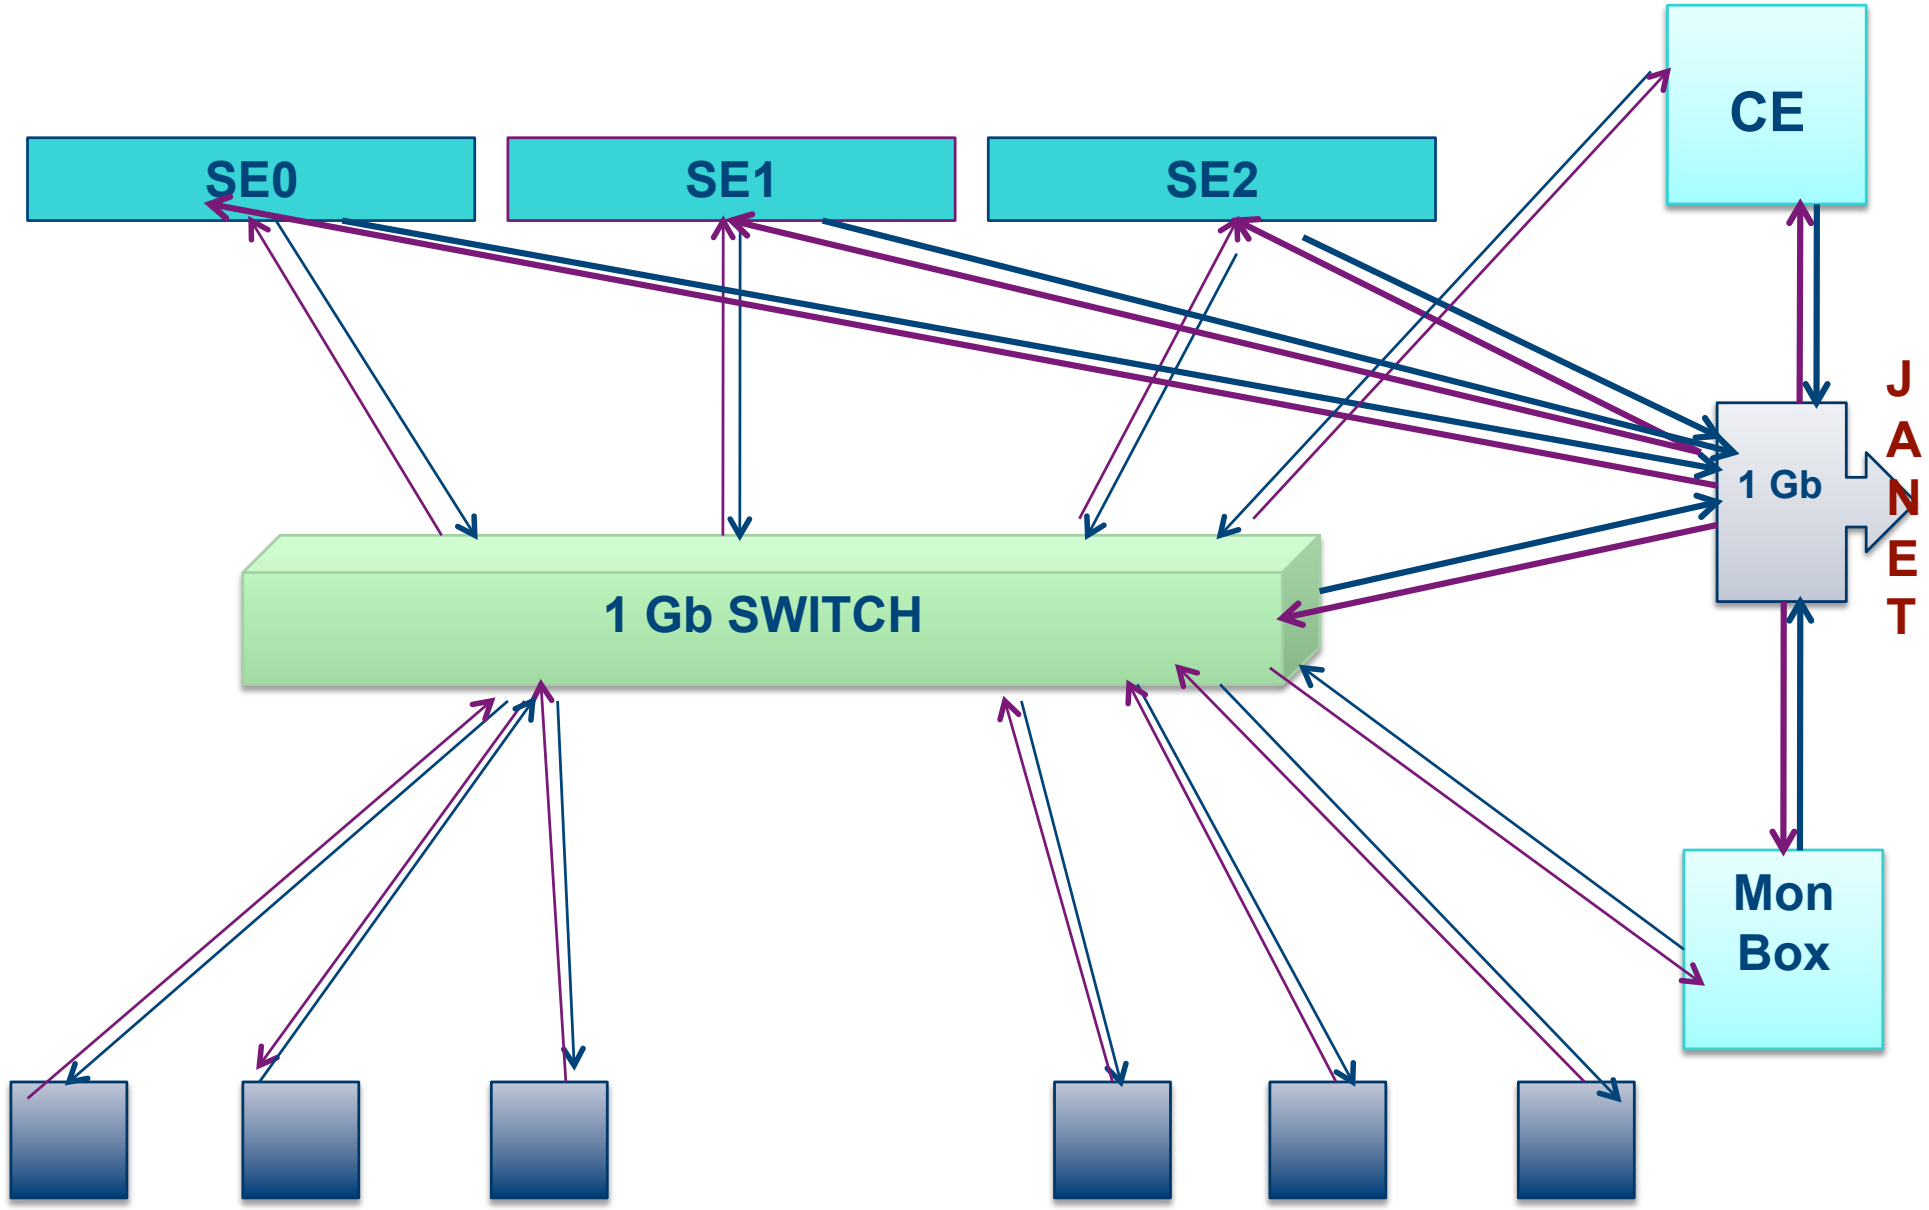


100 WNs (200 CPUs)

# Sheffield storage: future work

- **DPM 1.6 11**
  - will be upgraded to version 1.7 in July
  - need to drain the pool and keep space tokens
- **Put 2 TB disks instead of 1 TB into disk pools**

# STEP09 : Preparation

- **UKI-NORTHGRID-SHEF_HEP was subscribed for :**
  - Monte carlo production
  - data distribution
  - Hammercloud user analysis jobs (pilots and WMS)
- **expecting to run 200 jobs on 200 CPUs**
  - 100 production jobs
  - 50 analysis jobs submitted via WMS
  - 50 analysis jobs submitted via PANDA
- **expecting to accept 10% of UK AOD and DPD data**
  - 16 TB were reserved for ATLASDATADISK
  - 12 TB for ATLASMCDISK.

# STEP09 : Preparation

- **The batch system configuration (for 200 CPUs)**
  - non-LHC VOs was limited to 5 running jobs
  - LHCb was limited to 25
  - MAXJOB = 50 for WMS analysis jobs
  - MAXJOB = 50 for PANDA analysis jobs
  - MAXJOB = 25 for for individual atlas user (without ROLE)
  - no limit was set for production job

# STEP09 : Preparation

- **Hammercloud tests in December 2008 – April 2009 with small AOD**
  - high job completion rate (> 95%)
  - jobs efficiency (87%).
  - readahead buffer 128 MB to increase job efficiency
- **Pre-STEP09 Hammercloud test with large merged AOD 26.05 2009**
  - showed that our storage running out of memory
  - job completion rate and job efficiency dropped to 1%.
  - RFIO read ahead buffer to 32 MB.

**Storage load and memory during preSTEP09 Hammercloud test**

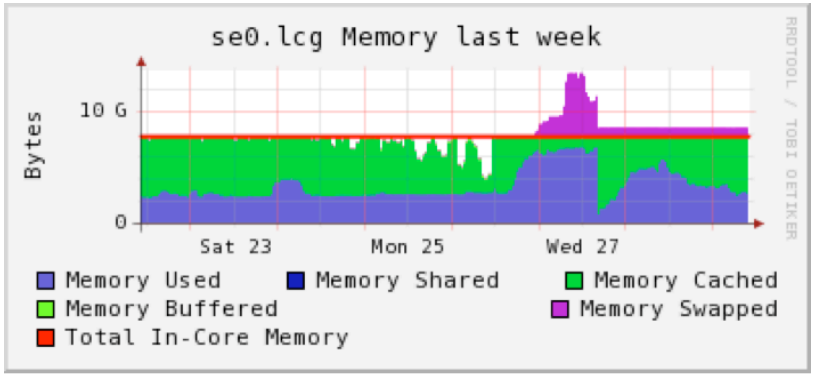


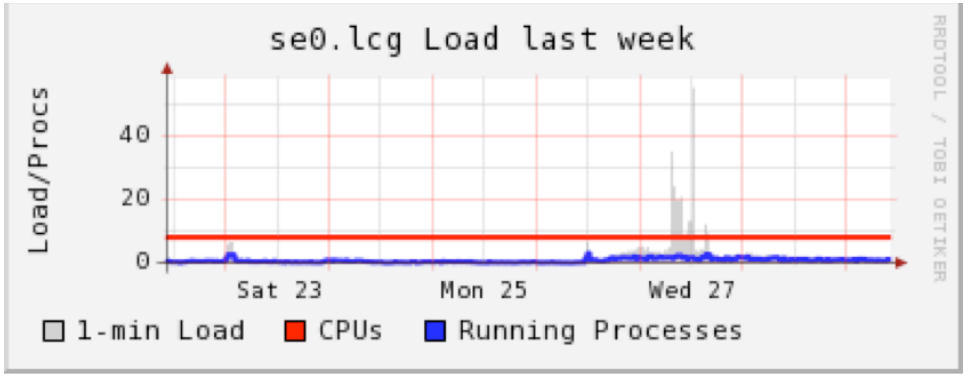

# STEP09 : Production

- **4900 production jobs were done In Sheffield**
  - job success rate is 84.2%  (86.6% if exclude jobs killed by Panda server)
  - average success rate for production jobs in 2009 (I- V) 94.2%.
  - 90.1% in the absence of analysis jobs submitted by WMS service
- **jobs failures were caused by network load:**
  - timeouts for contacting BDII in RAL, and LFC in RAL
  - we couldn't reproduce these errors (coping files from a WN to local SE)

| cluster | defined | assigned | waiting | activated | running | holding | transferring | success | failure | efficiency |
|---|---|---|---|---|---|---|---|---|---|---|
| UKI-NORTHGRID-SHEF-HEP-lcgce0-atlas-lcgpbs | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 4099 | 769 | 84.2% |

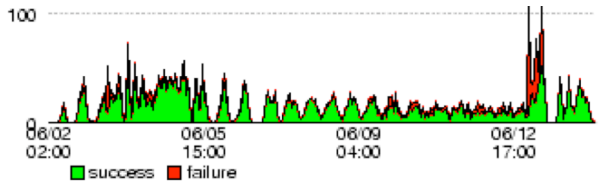


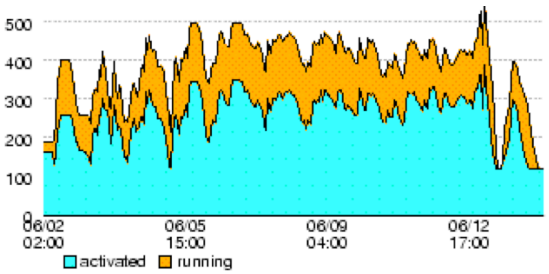


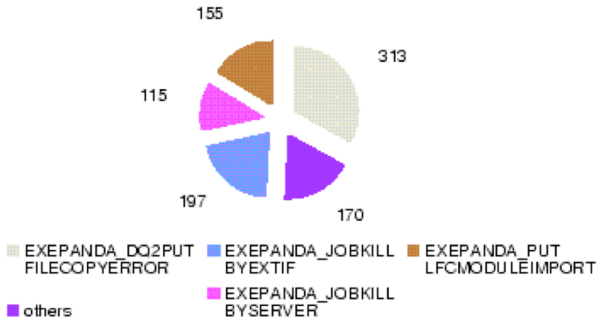


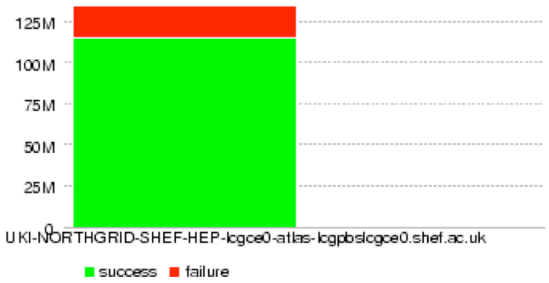

# STEP09 : Data Distribution

- **SHEF was subscribed a 10% share of ATLAS data during STEP09**
  - 6157 files were transferred in 1173 datasets
  - average throughput of 11 MB/s
  - efficiency of these transfers was 98 %.
- **No problem has been noticed for this part of the exercise**

Activity Summary ('2009-06-01 00:00' to '2009-06-14 13:50')

| | | Transfers | | Registrations | | Errors | | Services | |
|---|---|---|---|---|---|---|---|---|---|
| Cloud | Efficiency | Throughput | Successes | Datasets | Files | Transfer | Registration | Services | Grid |
| **UKI-NORTHGRID-SHEF-HEP_DATADISK** | **98%** | **11 MB/s** | **6157** | **1173** | **6063** | **155** | **0** | **0** | **ok** |
| **UKI-NORTHGRID-SHEF-HEP**_MCDISK | 98% | 1 MB/s | 1111 | 59 | 1119 | 27 | 0 | 0 | ok |

# STEP09 : Data Distribution

**TIER2S by sites**

last 24 hours

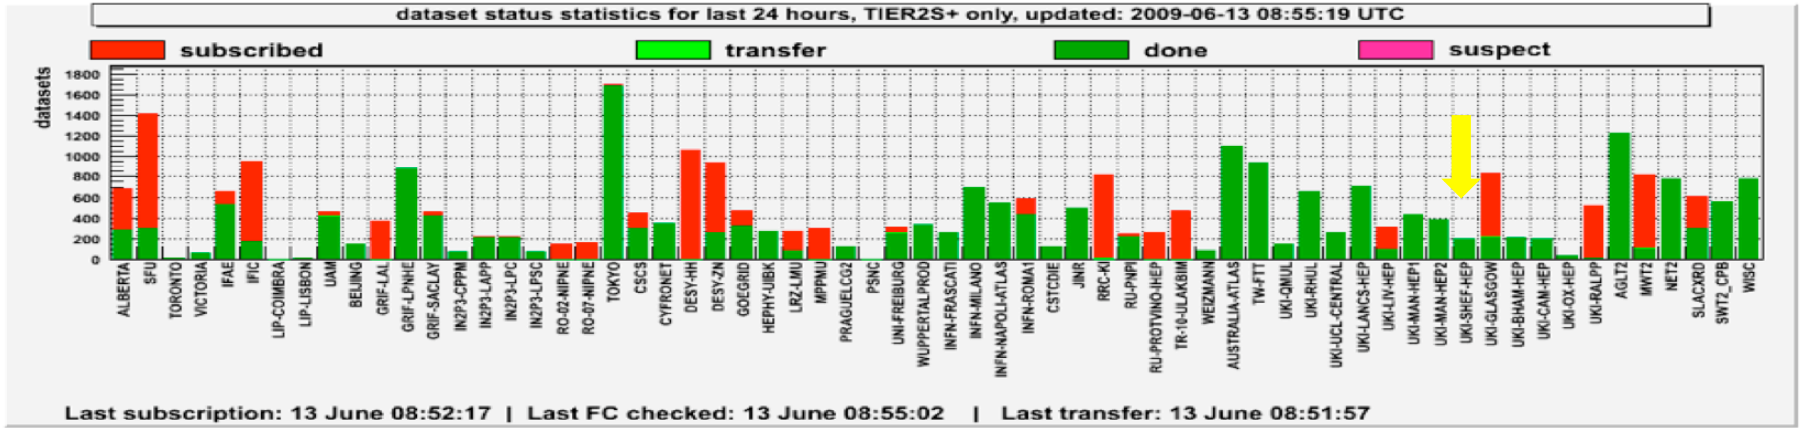


all test period

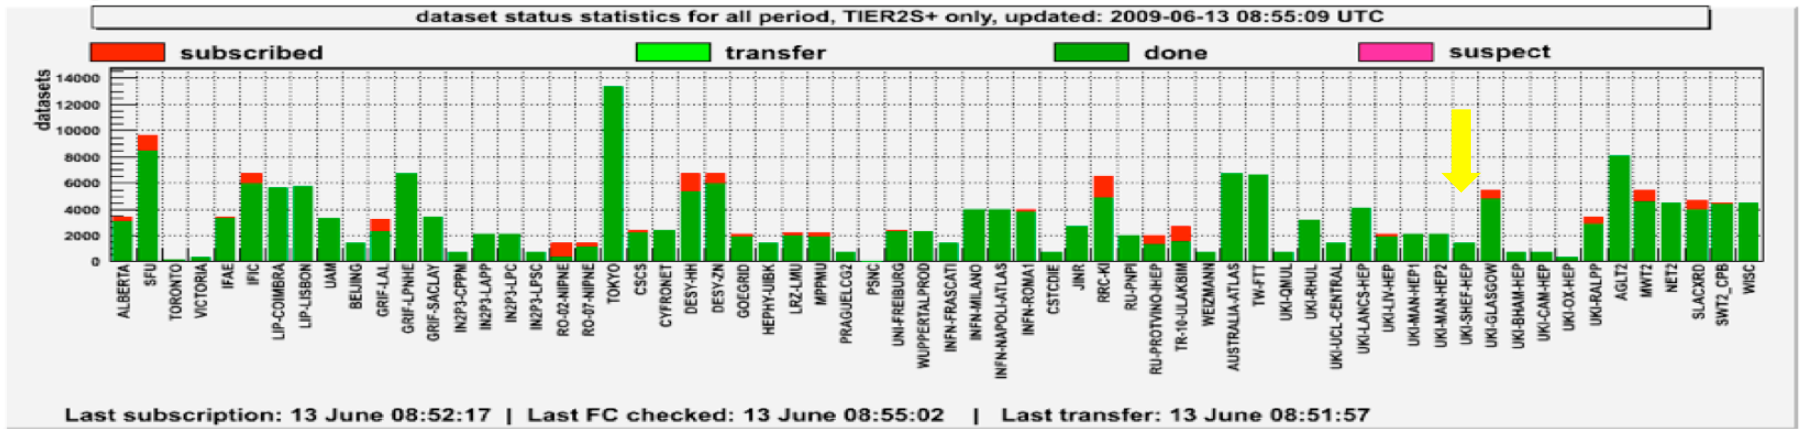


**FTS data transfer data backlog during STEP09**

# STEP09: Analysis jobs submitted via PANDA

- 6217 pilots analysis jobs were running at Sheffield
- job success rate was 82.1%
- job efficiency 50.8% (18$^{th}$ place in Winners: CPU/Wallclock)
- job success rate was higher at the end of week 1
- jobs failed mainly due to network problem:
  timeout errors while contacting LFC and BDII at RAL.

# STEP09: Analysis jobs submitted via WMS

- Sheffield was bombarded by Hammercloud jobs submitted by the WMS service in the morning of day 1
- 50 of these jobs (the maximum number allowed by our batch system) started to run on the cluster during one hour
- 1-2% efficiency
- > 200 jobs were queuing
- WMS jobs were competing for storage access with analysis pilots and ganga jobs submitted by 3 individual atlas users.
- by the end of day2 40 jobs were killed
- 10 jobs were running with very low efficiency and the storage network was still saturated
- 5 WMS jobs was running on the next day
- MAXJOB=5 for WMS analysis
- In the middle of week 2:
    - MAXJOB=1 and
    - RFIO buffer to 0

# STEP09: Analysis jobs submitted via WMS

- **Most of jobs haven't had a chance to run at Sheffield (proxy expired)**

- **Finally :**

  - 10 WMS analysis jobs were running in Sheffield
  - 5 of them were successful
  - job success rate of 50%
  - CPU/wall time ratio 10%

# STEP09: Analysis jobs submitted by individual ATLAS users

- Individual atlas users submitted GANGA jobs
- most of these jobs were using RFIO
- created additional load on the storage and network
- efficiency of 2-5 %
- on day 4 to improve cluster performance MAXJOB for each atlas user was limited to 5

# STEP09: Cluster Performance

## 1 Gb/s LAN and 1 Gb/s WAN links

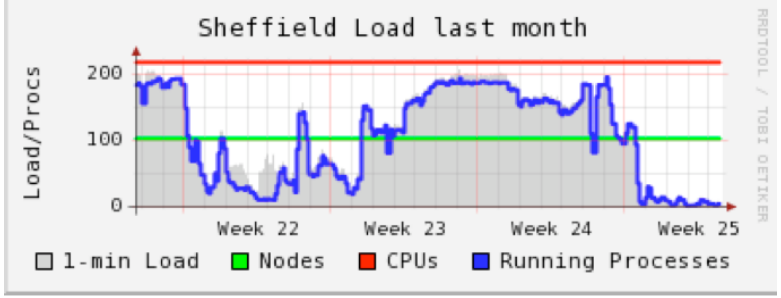


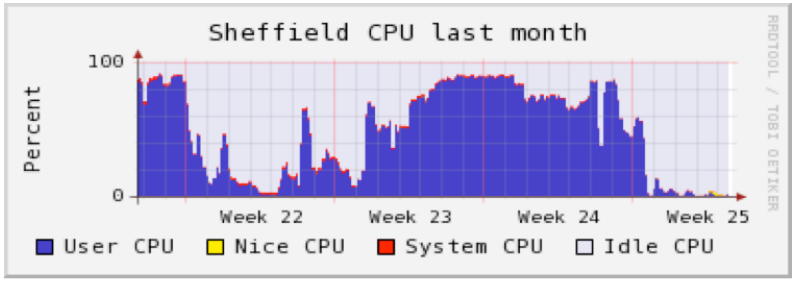


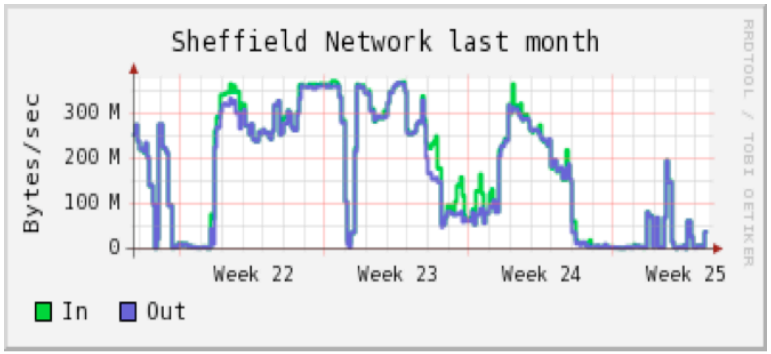


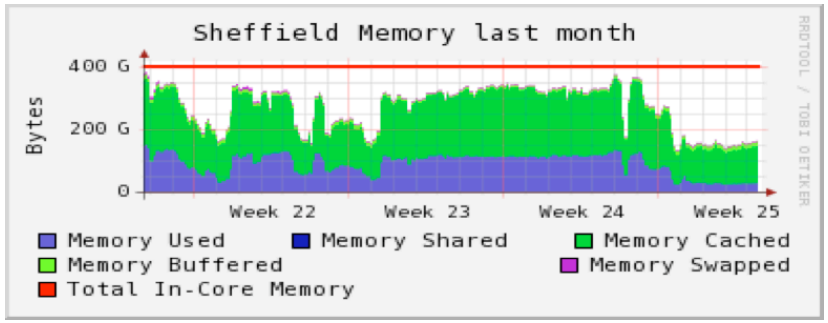


**Cluster usage during STEP09 (week 23 and 24)**

# STEP09: Cluster Performance

**The network bottleneck for storage on pre-STEP week (week 22) is probably related to the fact that we were using a 128 MB readahead buffer.**

**Server and Storage usage during STEP09 (week 23 and 24)**

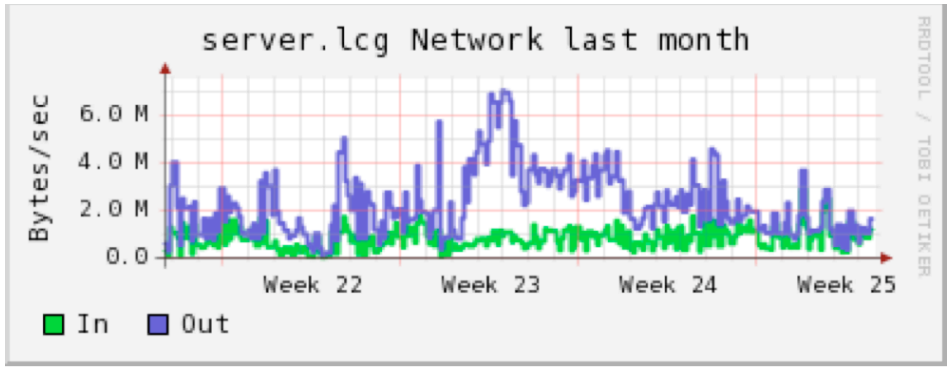


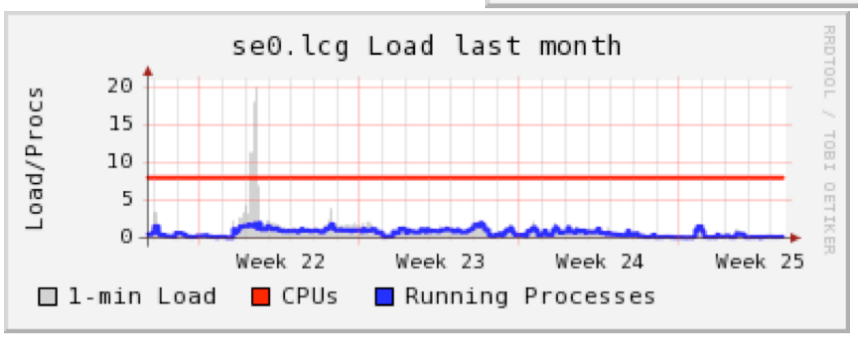


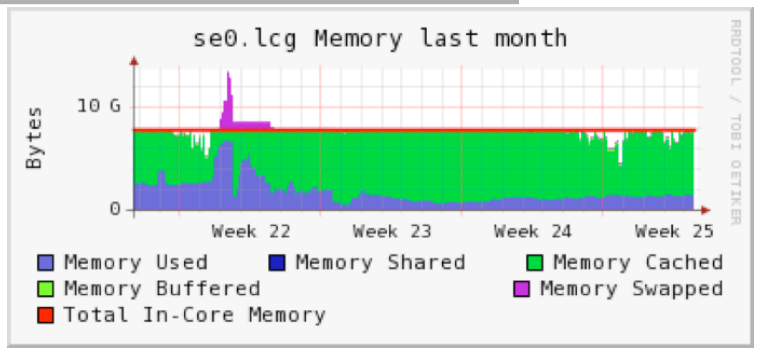


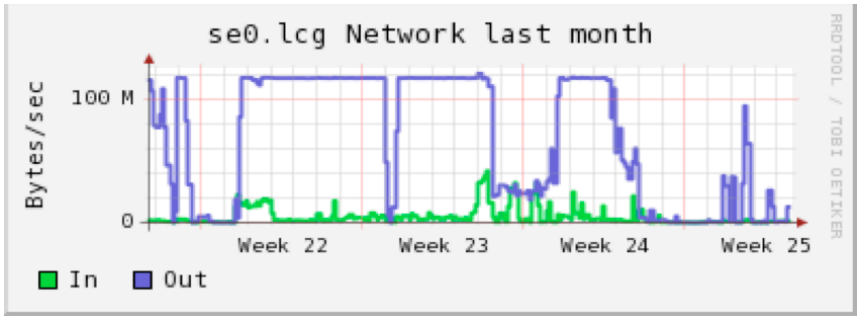


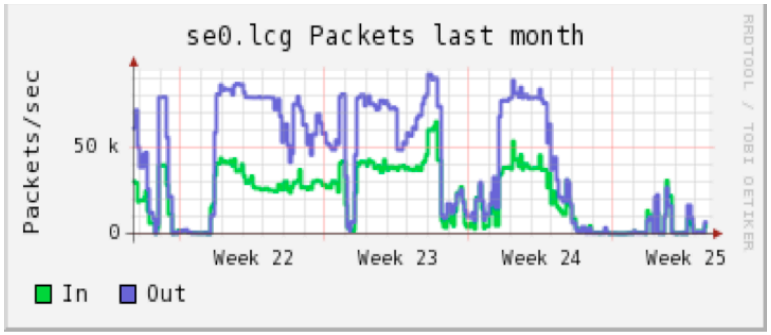

# STEP09: Conclusions

- Sheffield was successful in most STEP09 activities for Tier 2.

- succeeded in data distribution
  - no problems have been noticed
  - assigned of 12 TB of data ( 35 TB in total)
  - good share compared to larger sites
- successful in running Monte Carlo production
  - run 35% of the number of production jobs run in Glasgow
    200 old CPUs
- quite successful in running pilots analysis jobs
  - 1/3 less that the number of jobs in Liverpool
    ( 8% of jobs run in Glasgow).

- not successful in running WMS analysis jobs.

The number of these jobs submitted to our cluster was too large for us.

- need to understand why the production and pilots analysis job are failing with to the timeout errors when the cluster is heavily loaded. The percentage of such errors is not big but other sites didn't report this kind of problem.
-

# Sheffield LCG cluster

- Worker nodes
  - Dual 2.4 GHz AMD Opteron with 4 GB memory and 72 GB local disk
  - 100 WNs (200 CPU)
  - SL4 x86-64 gLite 3.1