# Testing CernVM-FS scalability at RAL Tier1

Ian Collier
RAL Tier1 Fabric Team
ian.collier@stfc.ac.uk
WLCG GDB - September 2010

# Contents

Experiment software server issues at RAL

Required characteristics for software area

CernVM-FS - a possible solution?

Current tests at RAL

Outlook

# Experiment Software Server issues

RAL software servers are similar to the setup described at PIC, with NFS servers, nodes for running install jobs etc.

We see the same issues and problems, particularly load related performance degradation on the Atlas software server

Not in a position to buy a NetApp or BlueArc

– but even they appear not be immune

Upgrading the NFS server helped a bit

– but depending on job mix 2000-3000 jobs Atlas can still bring the server to its knees – and knock WNs offline

# Some Experiment Software Server characteristics

Do not require write access

Many duplicated files, even within releases, never mind between releases.

In Atlas' case many repeated accesses of same file during jobs

Currently maintained with local install jobs at sites

Caching read-only filesystem would be ideal

# CernVM-FS ……. perhaps?

HTTP & Fuse based (first implementation was based on GROW-FS)

Developed to uniformly deliver experiment software to CernVM appliances

- Not inherently anything to do with virtualisation
- May deliver even greater benefits to physical WNs

With local squid proxies it should scale easily – just add squids (although our initial tests suggest one proxy will be fine)

File based deduplication (a side effect of integrity checks) – no file gets transferred twice to a given cache – software areas contain many, many identical files.

Caches at CERN, at local squids, and on client WN

- Repeated access of software area during jobs (eg Atlas conditions files) all becomes local after first access

# CernVM-FS testing at RAL

Currently testing scalability at RAL

So far just 800 or so jobs

**RAL most interested in resolving file server issues**

If we can fix that then more jobs will succeed

Of course if it means jobs run faster too that would be nice

# CernVM-FS tests

- Now testing at RAL:
- So far just 800 jobs or so - the squid proxy barely misses a beat:

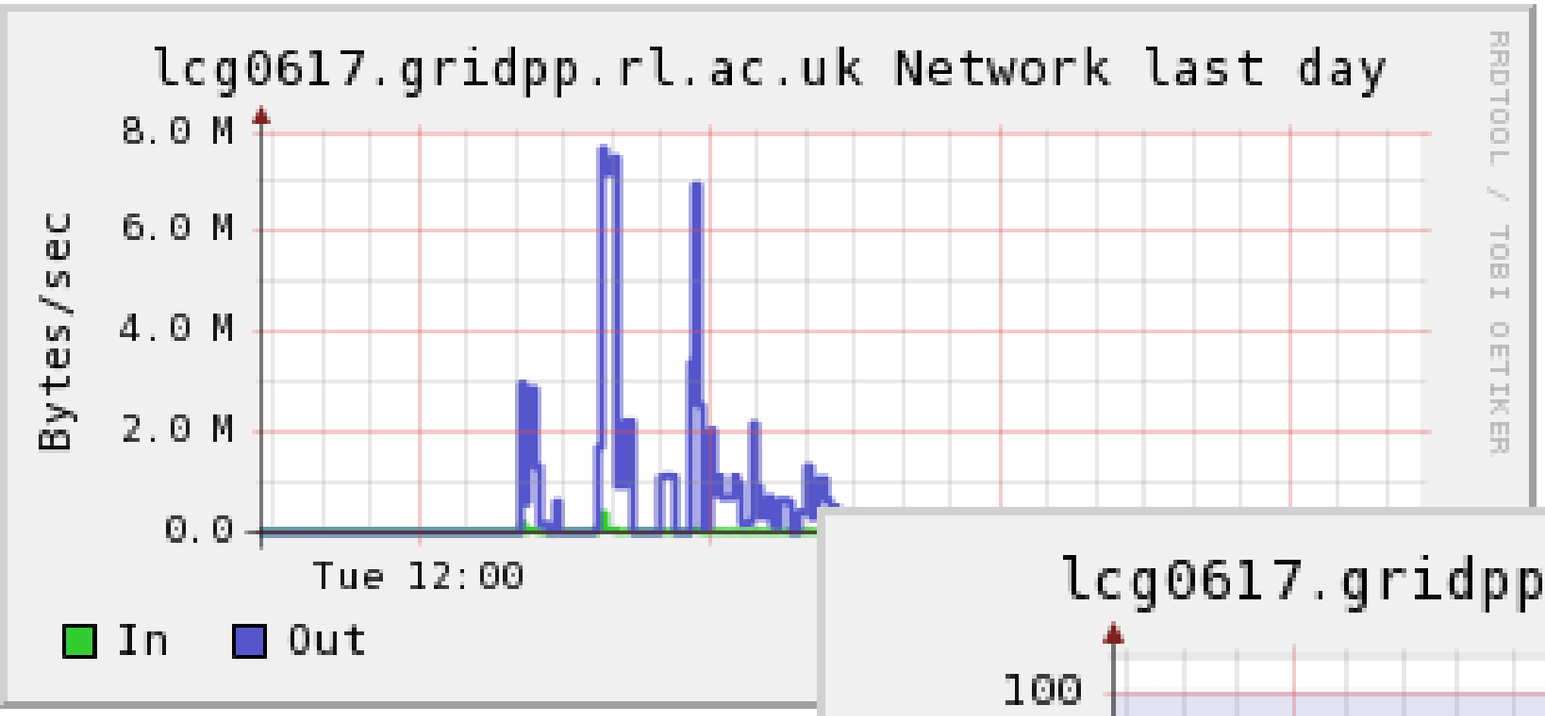


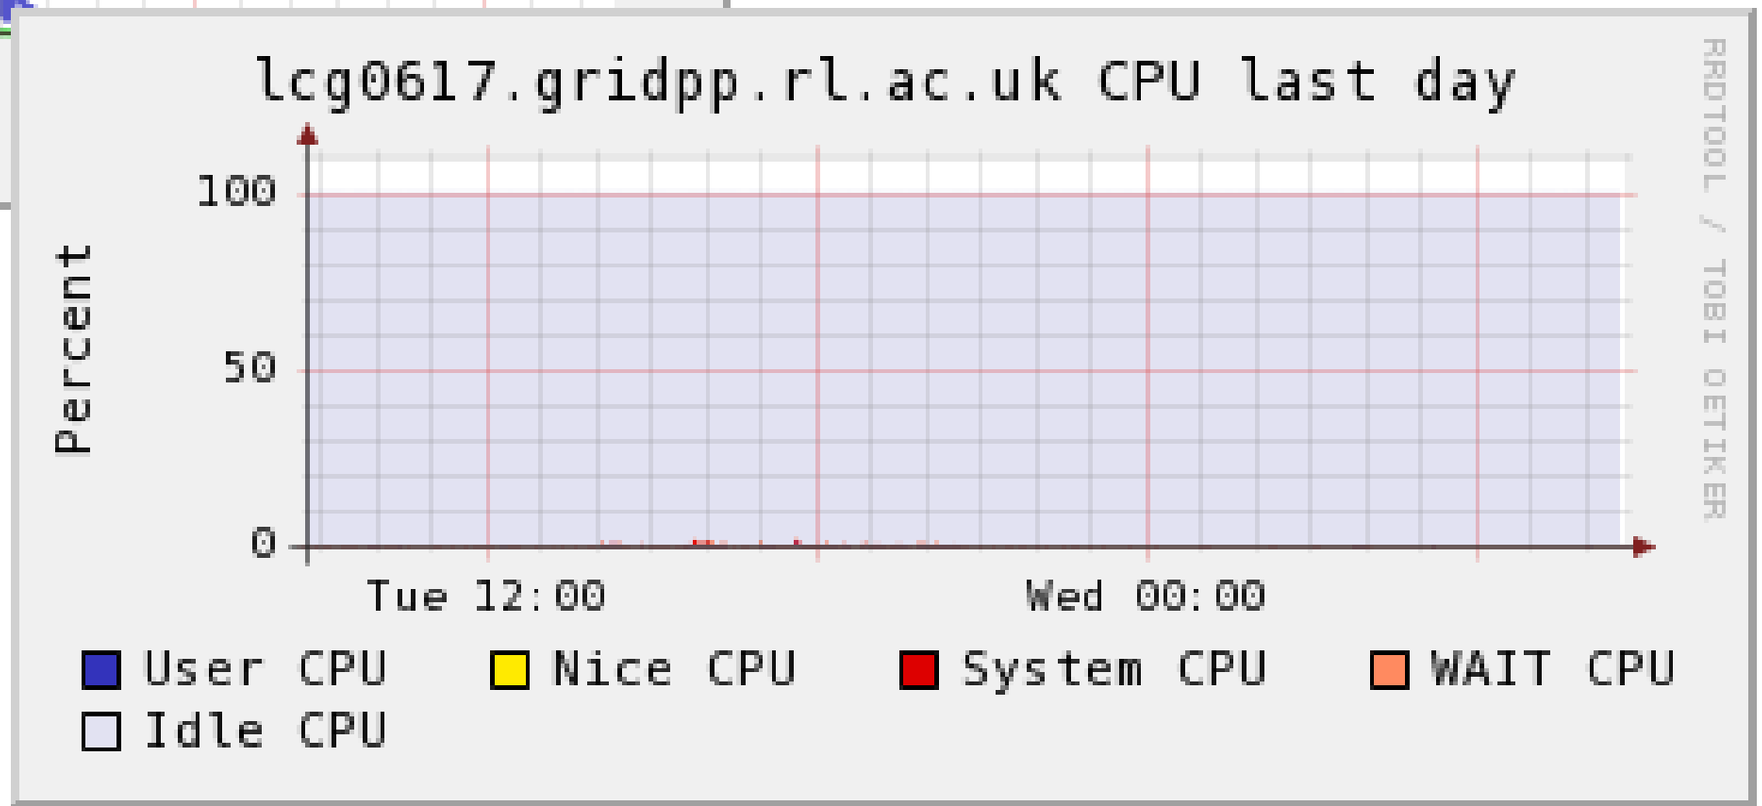

# CernVM-FS tests

By comparison the production Atlas Software server in the same week (not especially busy):

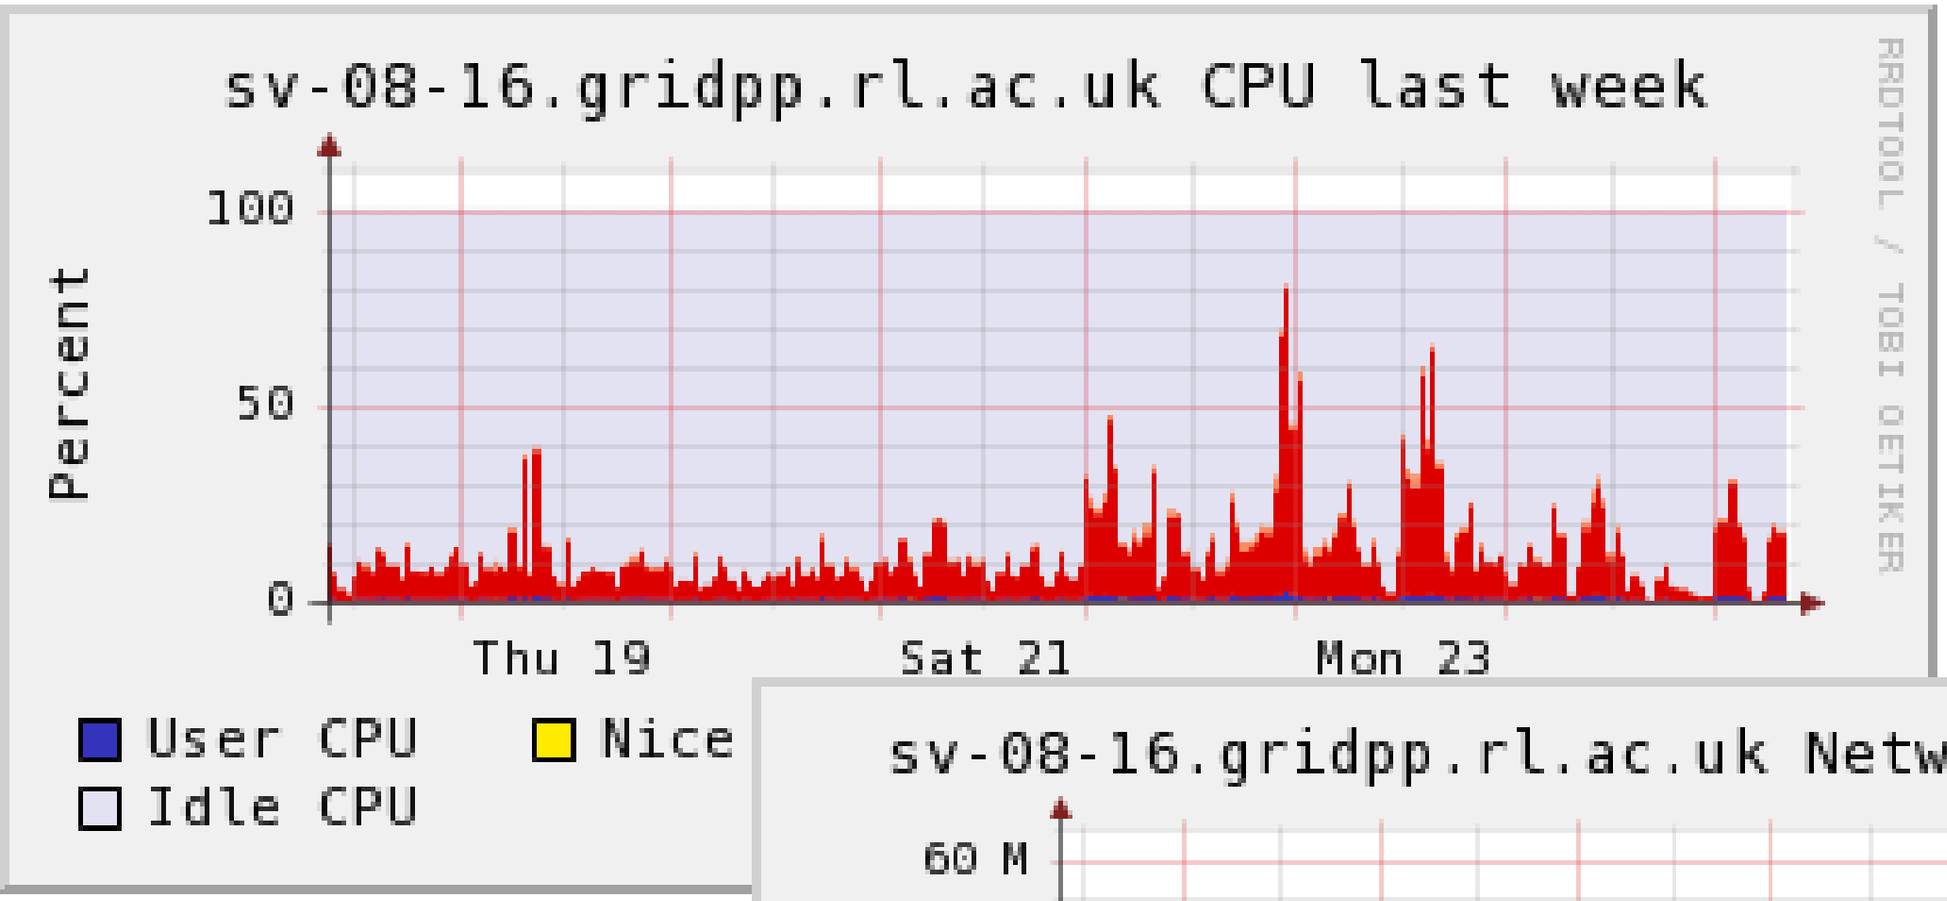


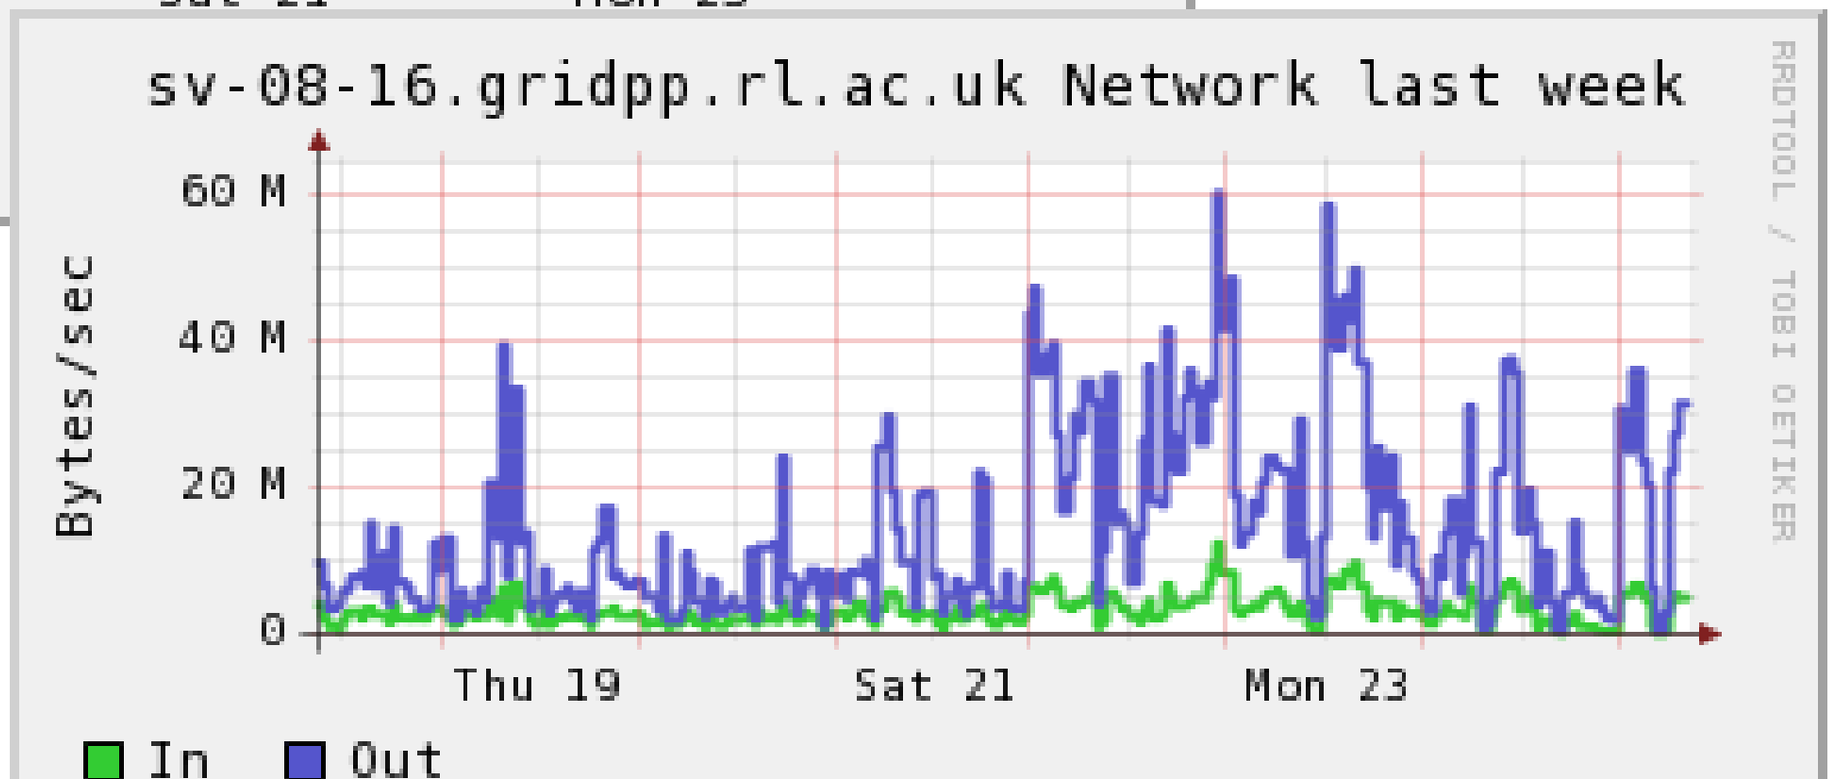

# CernVM-FS at RAL - next steps

Scale tests to thousands of jobs – planning now.

Compare performance with NFS server as check

Install latest version – supports multiple VOs

Work out and document use with grid jobs to allow wider tests

# CernVM-FS - outlook

Security audit of software still to be completed

Question about production support still to be answered

Would eliminate local install jobs

Potential for standard software distribution to all sites

(with any configuration areas for each site on server)

# So far it looks very promising.

Some more detail:
https://cernvm.cern.ch/project/trac/cernvm/export/1693/cvmfs-tr/cvmfstech.preview.pdf
and
http://indico.cern.ch/getFile.py/access?contribId=36&sessionId=3&resId=0&materialId=slides&confId=89681