



Status and Plans for the NIKHEF Tier 1

Davide Salomoni
for the NIKHEF CT/Grid Group

Presented at the LCG Workshop 2004,
23-24/3/2004, CERN

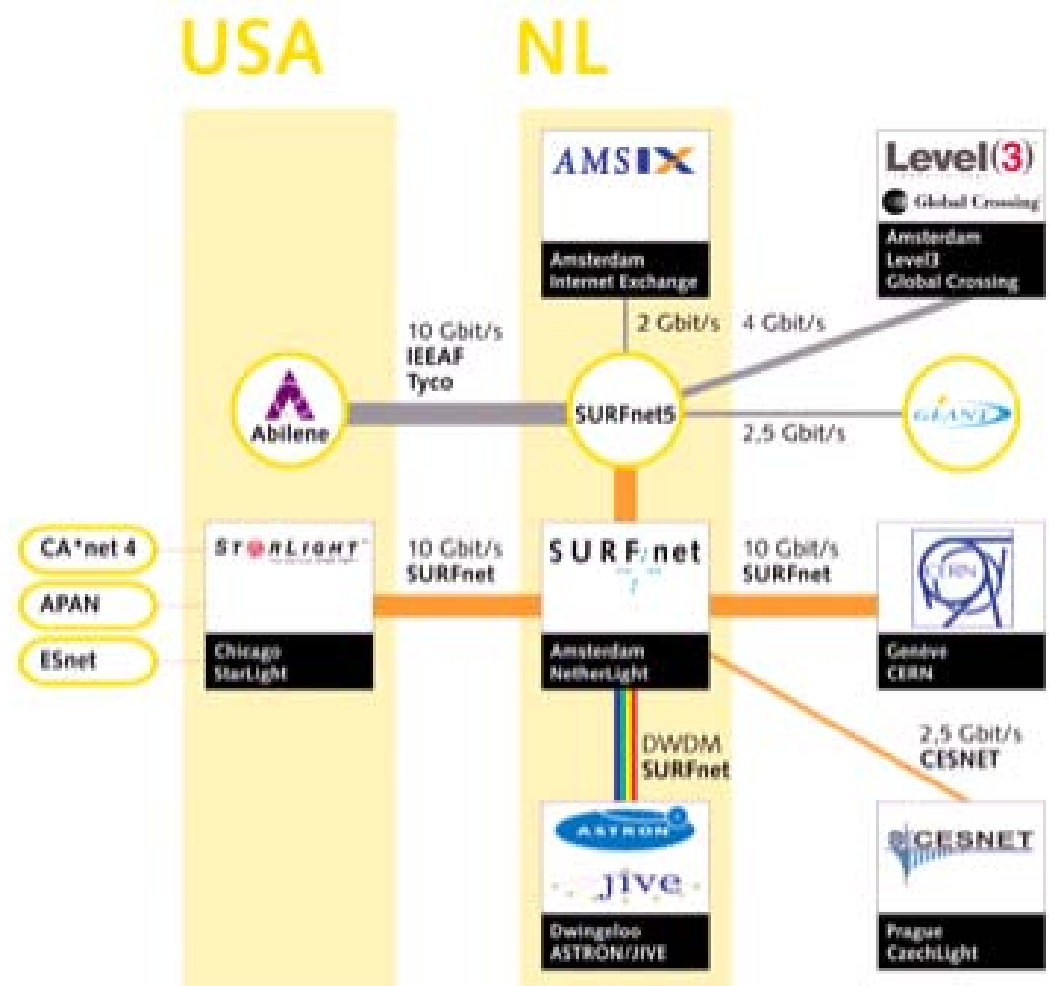


Agenda

- Infrastructure
- Funding and Support
- Main Activities
- Outlook

Infrastructure

- Network connectivity via SURFnet5
 - 10 Gbit/s to:
 - Abilene
 - STARLIGHT
 - CERN
 - GEANT
 - 4 Gbit/s to:
 - Level3 (GBX)
 - 2.5 Gbit/s to:
 - CESNET
 - 2 Gbit/s to:
 - AMSIX

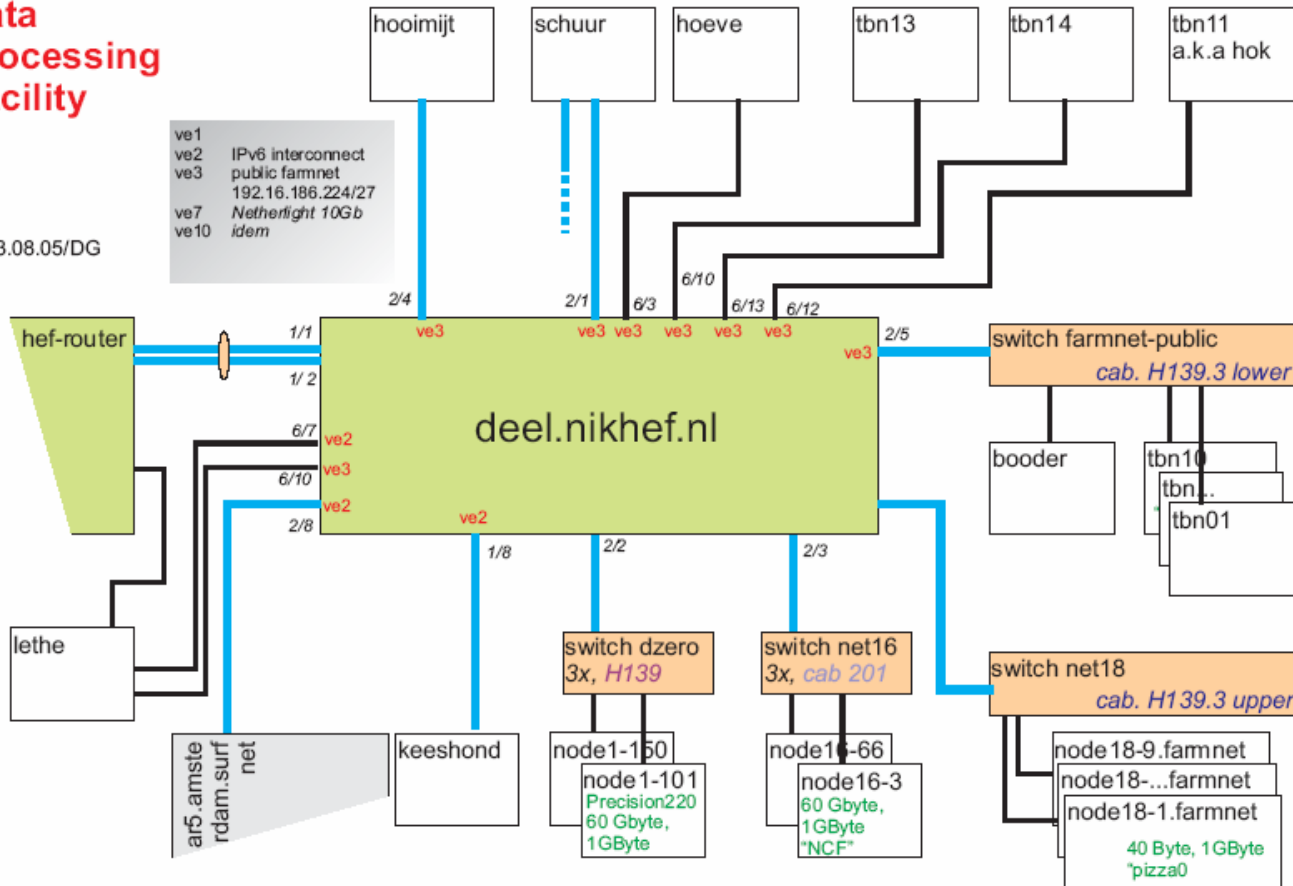




Internal NIKHEF Grid Network

NIKHEF
Data
Processing
Facility

2003.08.05/DG





NIKHEF Production and Test Clusters

- LCG2
 - ~ 250 CPUs currently published
 - The LCG1 infrastructure – LCFG, RB, UI, SE, CE – is still there, but with no WNs, and will be dismantled soon
- EDG Application Testbed
 - ~ 32 CPUs currently published
 - Will survive probably till April 1st
- EDG Development Testbed
 - Formally dead, but a few machines are still there
 - Will likely become the EGEE JRA1 testbed
- VL-E Certification Testbed
 - To test the latest VL-E software versions; will probably have around 10 WNs (plus CE,SE, etc)



NIKHEF Main LCG Cluster

- Actually composed of several pre-existing clusters:
 - Former D0 Farm: ~50 dual-PIII 866 MHz
 - NCF Cluster: 66 dual-AMD 2 GHz
 - Halloween Cluster: 27 dual-Xeon 2.8 GHz
- In total: ~300 CPUs, ~500 GHz → 200 kSI2k, ~6TB disk
- On versions:
 - LCG1 installed on Dec 1, 2003
 - LCG2 installed on Jan 29, 2004
 - WNs migrated to LCG2 on Mar 5, 2004

SARA

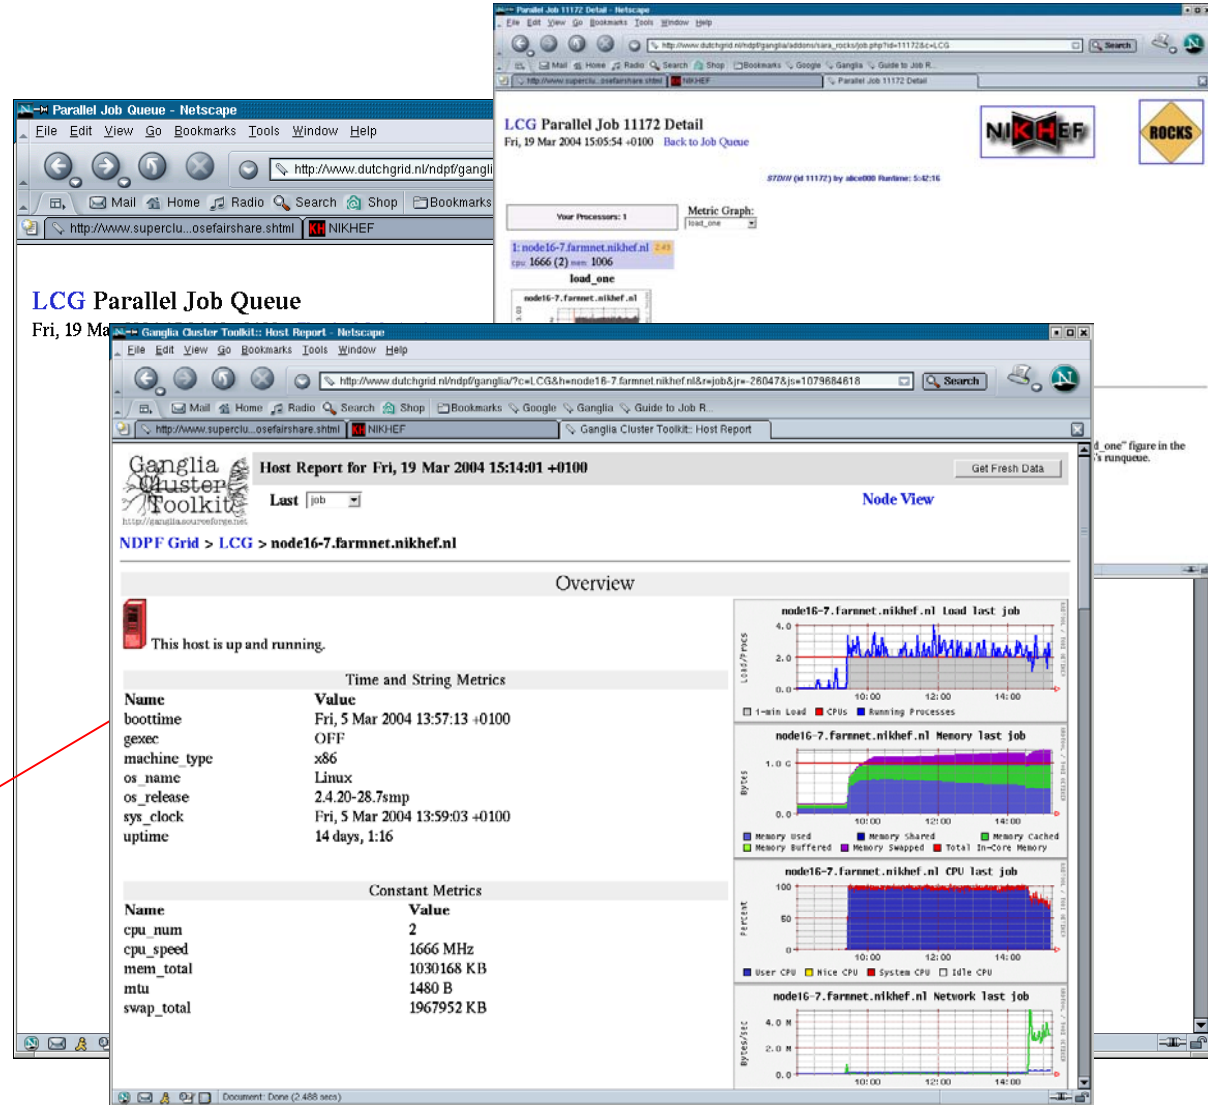


- Teras: SGI Origin 3800 w/ 1024 CPUs
 - Runs a version of GDMP ported to Irix
- Aster: SGI Altix 3700 w/ 486 CPUs
Itanium64 @1.3 GHz
- 400 TB tape storage seen as disk via DMF
- Brand new cluster in Almere (North of Amsterdam) w/ 72 CPU, 1 TB disk
- Running EDG software, will migrate to LCG Real Soon Now (before April 1st)

Network Monitoring

- Ganglia-based
- With PBS queue monitoring through a Ganglia add-on developed at SARA

Yes! We eventually managed to have Alice jobs running at NIKHEF 😊





Funding




- Mainly via NWO (Dutch Organization for Scientific Research) through the NCF (National Computer Facility)
- For the years 2004-2006, 2.8 M€ will be funded for Grid infrastructures in NL
- This is for projects that include LCG, but also non-HEP activities like LOFAR (Low Frequency Array), DEISA, EGEE, VL-E (Virtual Laboratory for e-Science) and others
- This brings us to an extremely important point...

The fabric must stay generic and multi-disciplinary
Hence (some of) our interest in security, scalability

Grid support at the NL Tier 1

- SARA: 5 dedicated people by April 1st
- NIKHEF: CT/GRID group – 12+ people



- Email: grid.support@nikhef.nl



Main (Current, LCG-Related) Activities

- LCG installations
 - Tide-up infrastructure and services
- Migration to TORQUE from OpenPBS
 - Always looking to share experiences with others
 - This includes e.g. fine-tuning of MAUI configurations
- Work on predictions of job start times on clusters
 - A very weak point of the current Grid architecture is the overly simple Glue schema and the effect this has on ETT and job submissions
- Security (e.g. JRA3 in EGEE)



Getting a bit off the point: MAUI

- A word on our MAUI configuration:
 - Fair Share: “specify usage targets to limit resource access or adjust priority based on historical resource usage” (from the MAUI admin manual)
 - Work needs to be done to dynamically adjust the configuration to changed patterns or fabric set-up

```

# Priority Weights
QUEUETIMEWEIGHT      1
XFACTORWEIGHT        10
RESWEIGHT             10

CREDWEIGHT           1
USERWEIGHT           1
GROUPWEIGHT          1

FSWEIGHT             10
FSUSERWEIGHT         5
FSGROUPWEIGHT        1

# FairShare
# use dedicated CPU ("wallclocktime used") metering
FSPOLICY              DEDICATEDDPS
FSDEPTH               192
FSINTERVAL            00:30:00
FSDECAY               0.99
FSCAP                 25

# fair shares and limits
GROUPCFG[DEFAULT]    FSTARGET=0      PRIORITY=100    MAXPROC=99999
USERCFG[DEFAULT]     FSTARGET=0      PRIORITY=100    MAXPROC=99999

GROUPCFG[users]      FSTARGET=20+    PRIORITY=500
GROUPCFG[tmpusr]     FSTARGET=1-     PRIORITY=10

GROUPCFG[dteam]      FSTARGET=5      PRIORITY=1000   MAXPROC=50
GROUPCFG[iteam]      FSTARGET=2      PRIORITY=500    MAXPROC=10
GROUPCFG[wpsix]      FSTARGET=3      PRIORITY=500    MAXPROC=10
GROUPCFG[atlas]      FSTARGET=20     PRIORITY=500    MAXPROC=100
GROUPCFG[dzero]      FSTARGET=20     PRIORITY=500    MAXPROC=100
GROUPCFG[alice]      FSTARGET=20     PRIORITY=500    MAXPROC=150
GROUPCFG[lhcb]       FSTARGET=20     PRIORITY=500    MAXPROC=100
GROUPCFG[cms]        FSTARGET=0      PRIORITY=0       MAXPROC=10
GROUPCFG[biome]      FSTARGET=5      PRIORITY=500    MAXPROC=2
GROUPCFG[eo]         FSTARGET=5      PRIORITY=500    MAXPROC=32
GROUPCFG[tutor]      FSTARGET=5      PRIORITY=500    MAXPROC=0


USERCFG[svens]       FSTARGET=5      PRIORITY=0       MAXPROC=32
USERCFG[sander]      FSTARGET=20+    PRIORITY=500    MAXPROC=150
USERCFG[s64]         FSTARGET=5      PRIORITY=0       MAXPROC=32
USERCFG[willem]      FSTARGET=20     PRIORITY=200    MAXPROC=180
USERCFG[stanb]       FSTARGET=20+    PRIORITY=500    MAXPROC=150
USERCFG[versto]      FSTARGET=2-     PRIORITY=0       MAXPROC=120
USERCFG[bosk]        FSTARGET=0      PRIORITY=0
USERCFG[dauidg]      FSTARGET=0      PRIORITY=50000
USERCFG[templon]     FSTARGET=0      PRIORITY=50000

```

Activities (cont.)

- We are interested in developments/ideas about a new software installation system
 - Remember: we want the fabric to remain generic and scalable
 - We like the idea of an on-demand cached filesystem to leave software installations directly to the user (no root or SGM privileges)
 - Don't involve grid core services if at all possible
 - Perhaps use ideas from slashgrid, zero-install, etc.

Let's build a "Thermal Grid"!
(Federico Carminati &
Jeff "Thermodynamic" Templon)

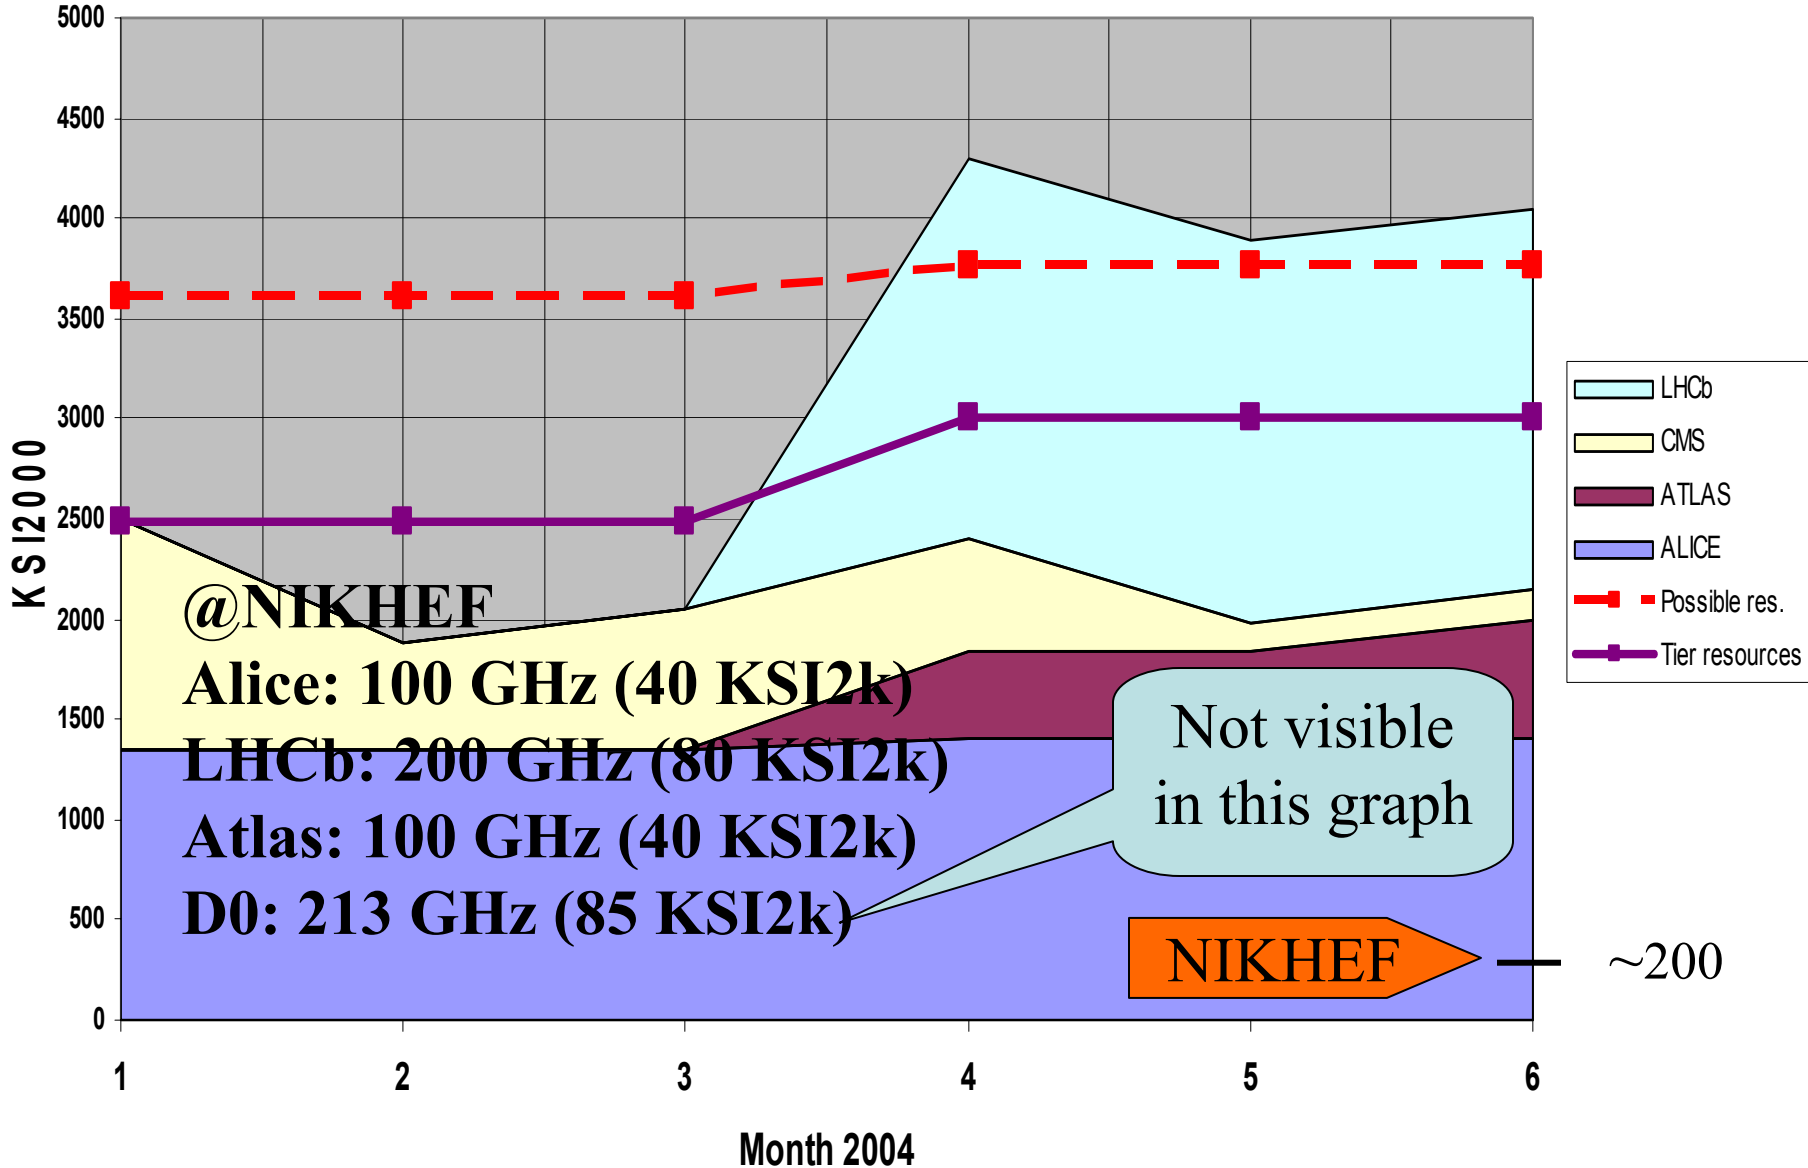


Activities (cont.)

- Simplify installation of LCG software
 - Keen interest in Quattor
 - We successfully tested manual installation of LCG UIs using apt (orig. from Debian) on RPM systems
 - We maintain an experimental apt-rpm repository on www.dutchgrid.nl/mirror/apt with scripts to install an LCG UI
 - Using apt on an RPM system greatly simplifies the resolution of sw dependencies: apt automatically fetches all required packages from the network
 - You are welcome to try this out (and give us feedback)



Resources and Requests





Outlook

- Alice, ATLAS, LHCb and D0 can almost get what they requested
- This year we probably won't significantly expand the fabric any further
- We need our partners in order to get more resources: SARA, NCF, NWO
- Eventually, most CPUs and all data archiving will be at SARA
- NIKHEF will provide central services and expertise