NIKHEF Test Bed Status

David Groep

davidg@nikhef.nl

NIKHEF: Current Farms and Network

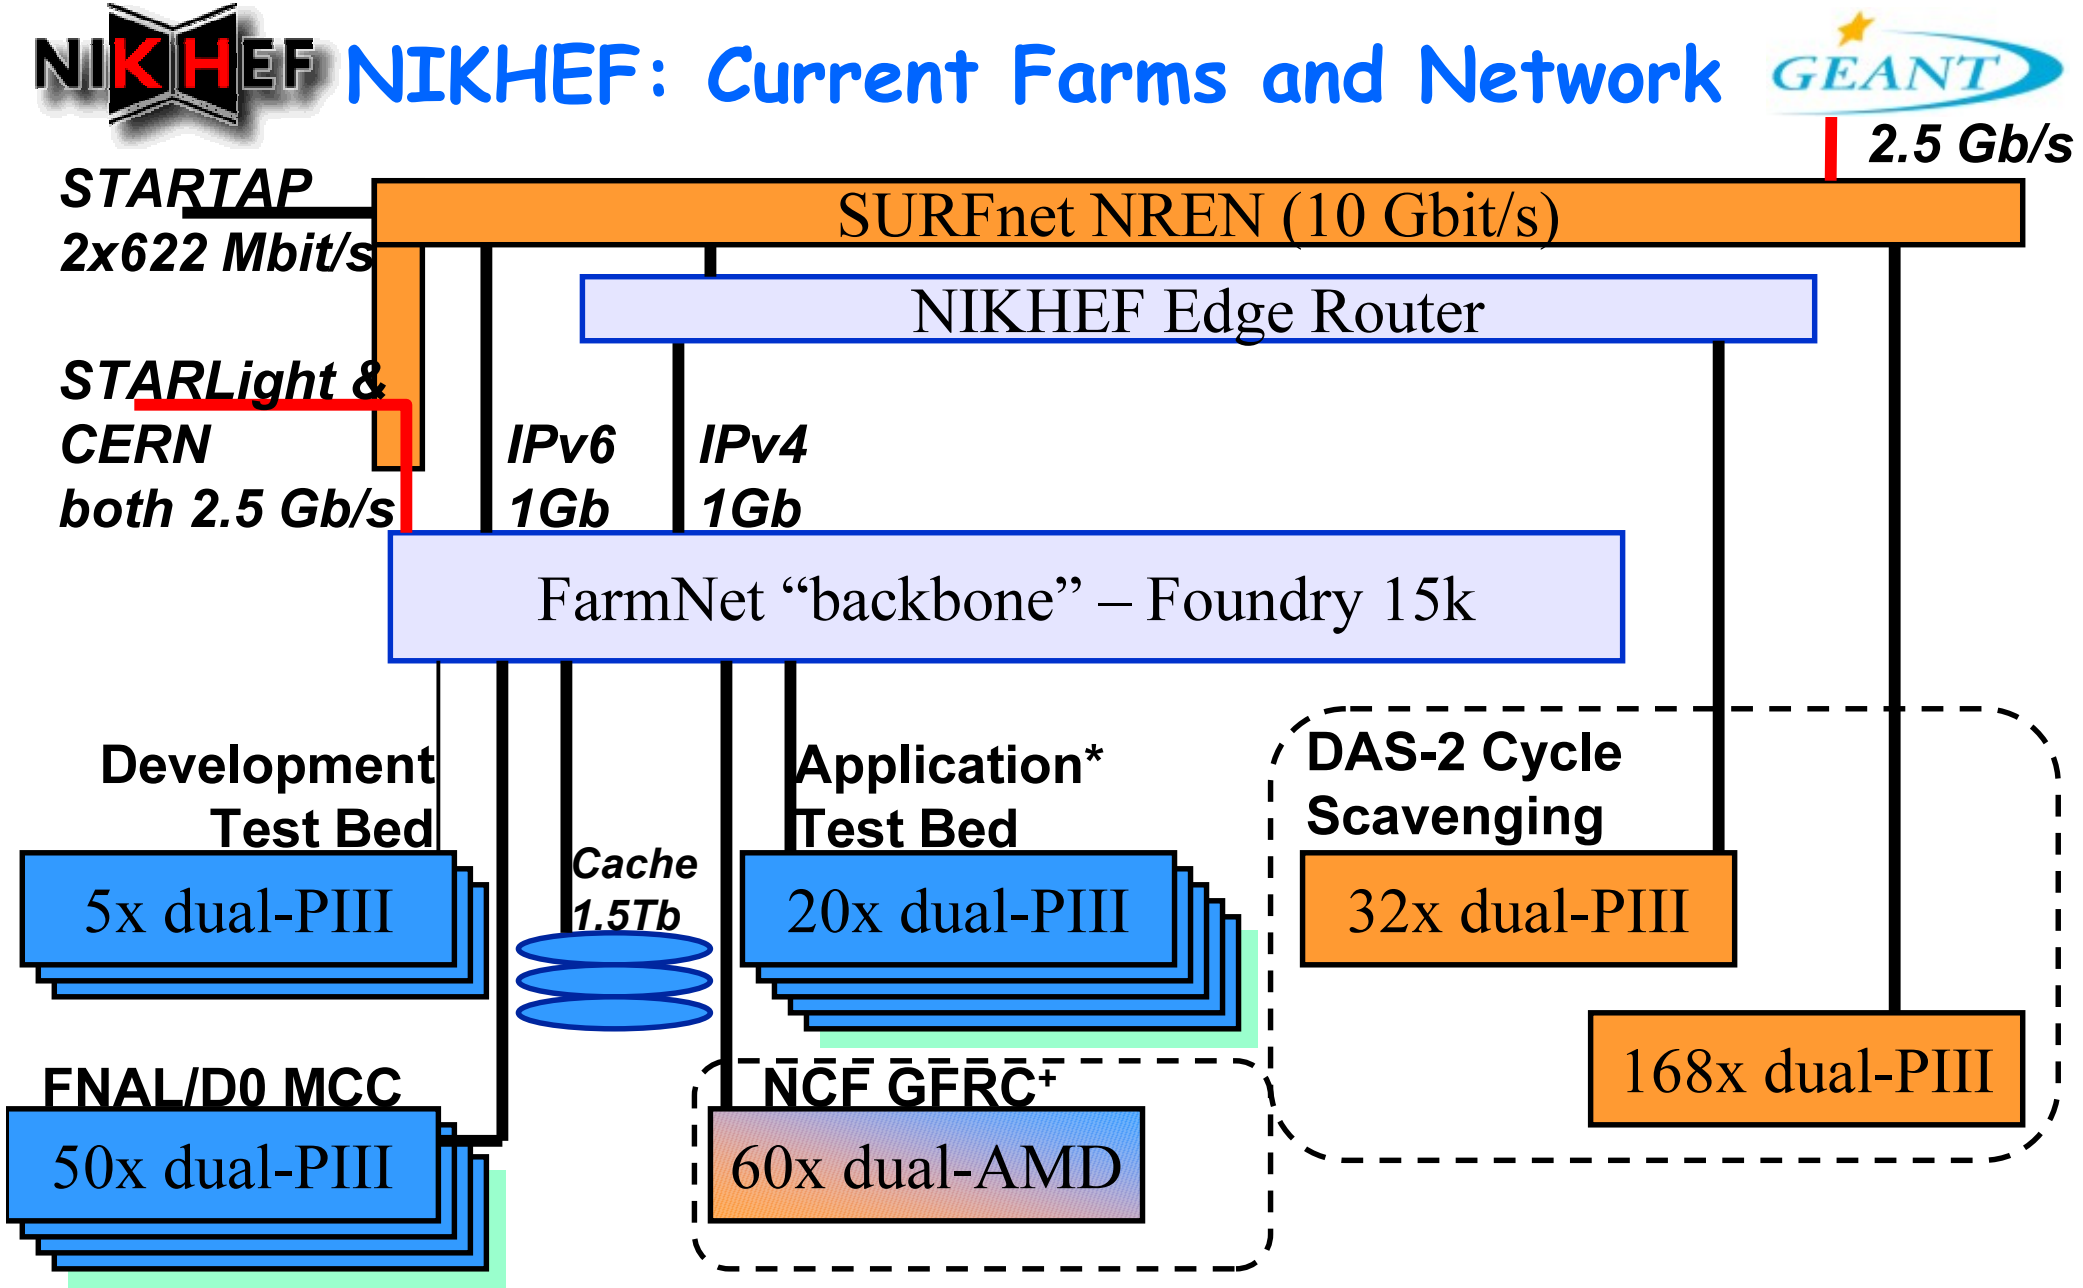

Testbed: current state

- Development test bed
  - EDG release 1.2.0
- Application test bed
  - EDG release 1.2.0
  - Support for all VOs (incl. EDGtutorial, CMS install pending)
- "External" services
  - Some of the VO membership and RC services (LHCb, Alice, EarthOb, BioMed, …)
- Resource sharing with Dzero MCC production

Test bed contacts

- For requests, please mail the generic address

  grid.support@nikhef.nl *

- Also for upgrade requests, RC problems, GDMP host quasi-VO, …

- We operate as a best-effort service

Test Bed Buildup stategy

"Why buy farms if you can get the cycles for free?"

- Get lots of cycles in "scavenging" mode from CS research clusters
- Attracts other support from CS faculties
- Get cycles from national supercomputer funding agencies

Downside:

- Many different clusters (but all run Globus and most EDG middleware)
- Middleware shall (and should) be truly multi-disciplinary!

SARA: Mass Storage

- NIKHEF "proper" does not do mass storage – only ~ 2 TByte cache
- SARA: 200 Tbyte StorageTek NearLine robot
- 2 Gbit/s interconnect to NIKHEF

- Front-end: "teras.sara.nl" 1024 processor MPP – SGI IRIX
- Ron Trompert ported GDMP to IRIX.
  Now running!

Challenges and Hints

- Farm installation using LCFG works fine
  - Re-install takes 15 minutes (largely due to application software)
  - Adapts well to many nodes with different functions (2xCE,2xSE,2xUI, external disk server, 2 acceptance-test nodes, 2 types WN, D0 nodes, …)
- Some remaining challenges
  - "edg-release" configuration files are hard to modify/optimize
  - RedHat 6.2 is really getting old!
  - Netbooting for system without FDD
  - Get all the application to work!

LCFG configuration

- Use EDG farm to also accommodate local user jobs

- disentangled hardware, system, authorization and app. Config
- using autofs to increase configurability (/home, GDMP areas)
- Installed many more RPMs (DØMCC, LHCb Gaudi) and home-grown LCFG objects (pbsexechost, autofs, hdparm, dirperm)
- Force RPM install trick (`+updaterpms.offline`)
- modified rdxprof to support multiple domains

- Shows flexibility of LCFG (with PAN it will be even nicer!)

```
[davidg@booder source]$ cat node18-10
/* node specific profile */

#include "inc/macros-cfg.h"                      /* Some useful macro */
#include "inc/nikhef-macros.h"

/* Host specific definitions */
#define HOSTNAME                node18-10
#define LOCALDOMAIN             farmnet.nikhef.nl
#define SITE_GATEWAYS           192.168.18.254
#define SITE_NETMASK            255.255.255.0
#define PBS_MASTER              ce02

/* Linux and boot configuration */
#include "inc/nikhef-site-config-app.h" /* Site specific definitions */
#include "edg-1.2/linuxdef-cfg.h"       /* Linux default resources */
#include "inc/nikhef-sysconfig-core.h"  /* LCFG client specific resources */
/* Hardware configuration */
#include "inc/nikhef-nodetype-pizza0.h" /* hardware config pizza0 nodes */
#include "inc/nikhef-disklayout-wn0.h"  /* disk layout for WN without home */
/* Software configuration */
#include "inc/nikhef-filesys-wn0.h"     /* non-local filesystems (autofs) */
#include "inc/nikhef-auth-config.h"     /* Core auth definitions (root) */
#include "inc/nikhef-users.h"           /* permanent local users */
#include "inc/nikhef-poolusers.h"       /* EDG DataGrid leased users */
#include "edg-1.2/WorkerNode-cfg.h"     /* WorkerNode specific resources */
                                        /* be sure to override rpmcfg */
#include "edg-1.2/pbs-cfg.h"              /* PBS specific config stuff */
#include "inc/nikhef-unmount-nfs.h"     /* get rid of fixed NFS mounts */
#include "inc/nikhef-pbsclient.h"       /* make this a client of the CE */
+update.rpmcfg                          rpmlist-net18
+updaterpms.offline                     upd020823171537-4165
```

RedHat 6.2 – modern-processor breakdown

- Recently acquired systems come with P4-XEON or AMD K7 "Athlon"
- Kernel on install disk (2.2.13) and in RH Updates (2.2.19) say "?????"

- Baseline: RedHat 6.2 is getting really old
- But a temporary solution
  can still be found (up to kernel 2.4.9):

  use new kernel (without dependencies) in existing system
- Requires you to build a new RPM
- You can even get the Intel 1Gig card to work (for installs -> Steve)
- See http://www.dutchgrid.nl/Admin/Nikhef/edg-testbed/

Installing systems without an FDD

- Most modern motherboards support PXE booting
- stock LCFG-install kernel works well with PXE

- "just" need a way to prevent an install loop
  - thttpd daemon with a perl script to "reset" dhcpd
  - called from modified dcsrc file
  - script will only reset dhcpd.conf when $REMOTE_ADDR matches

  - CNAF did something similar using temporary ssh keys

DutchGrid Platform

www.dutchgrid.nl

- DutchGrid:
  - Test bed coordination
  - PKI security
  - Support
- Participation by

  NIKHEF, KNMI, SARA

  DAS-2 (ASCI):
  TUDelft, Leiden, VU, UvA, Utrecht

  Telematics Institute

  FOM, NWO/NCF

  Min. EZ, ICES/KIS

  IBM, KPN, …

Our test bed in the Future

- We expect continuous growth

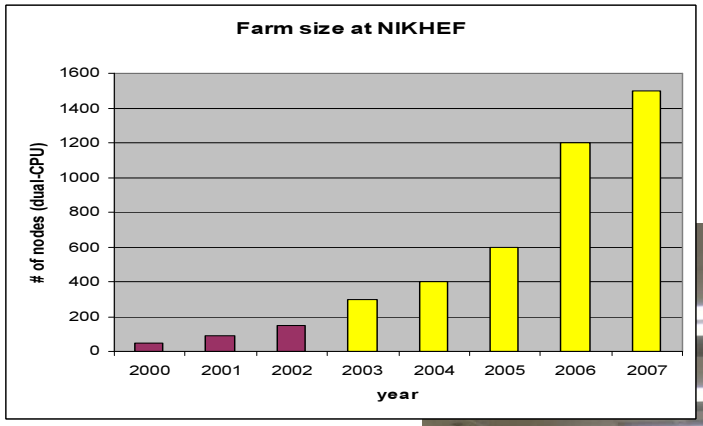

- Our Aims:
  - ~ 1600 CPUs by 2007
  - "infinite" storage @ SARA
  - 2.5 Gbit/s interconnects now
  - > 10 Gbit/s in 2003/2004?

- Our constraints:
  - The fabric must stay generic and multi-disciplinary