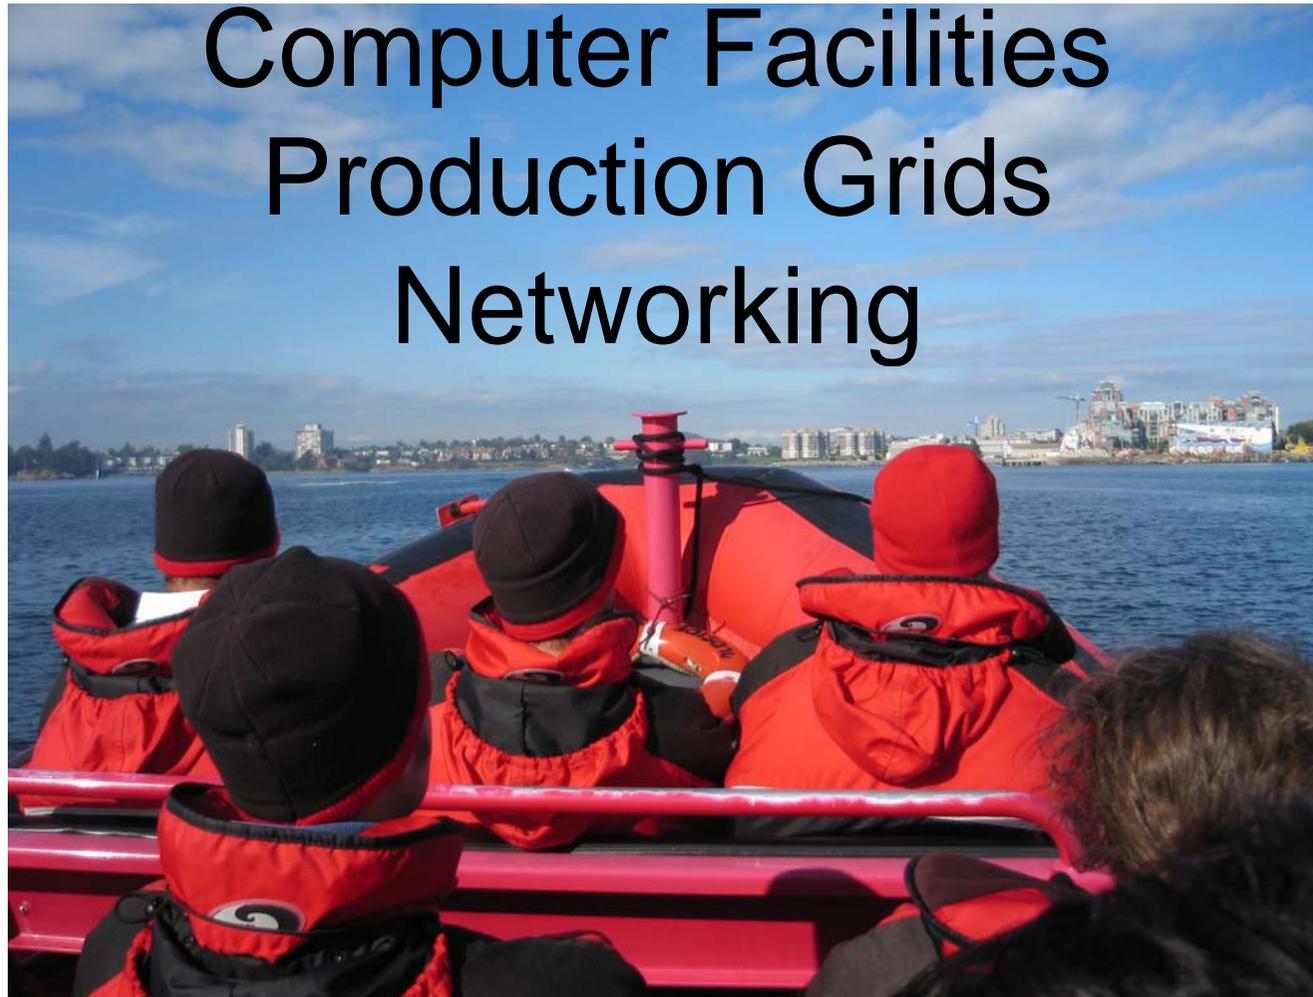


CHEP 2007



Computer Facilities
Production Grids
Networking

Session Summary

Kors Bos, NIKHEF, Amsterdam

Initial remarks

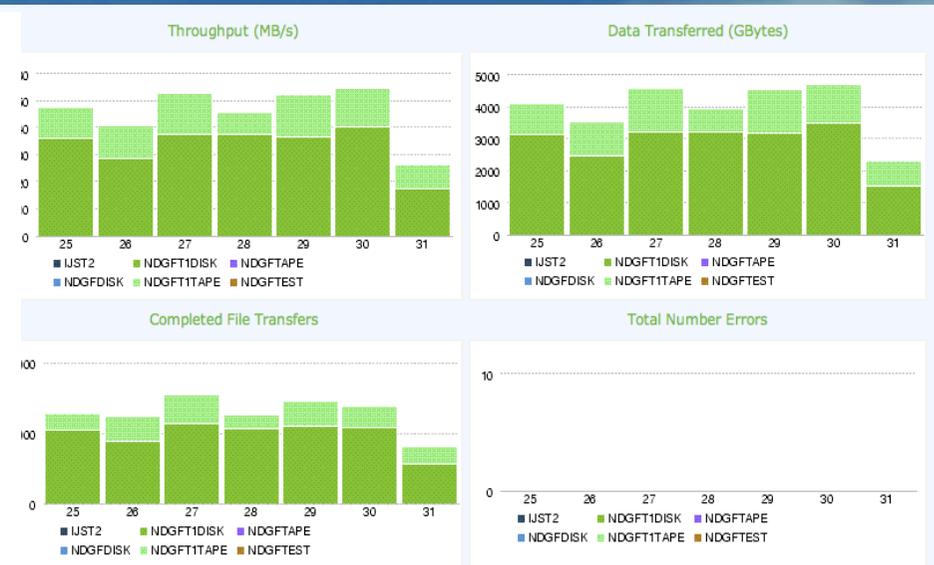
- 742 slides in 37 talks in 7 sessions
- Impossible to talk about them all
- Selection is my choice
- Apologize for not giving enough exposure
- All very interesting, have a look at the slides!

I. Grids

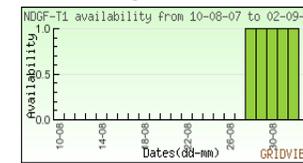
- WLCG Service Deployment (Jamie Shiers)
- Security Incident Management (Markus Schulz)
- OSG (Ruth Pordes)
- GridPP (Jeremy Coles)
- NDGF (Michael Grønager)



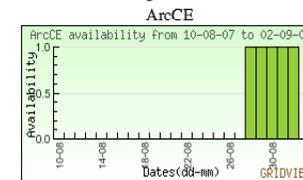
- We have build a distributed Tier-1
 - dCache – for storage
 - ARC for computing
- Interoperabel with:
 - ALICE
 - ATLAS
 - ARC monitoring and accounting
 - LCG monitoring and accounting
- It works !



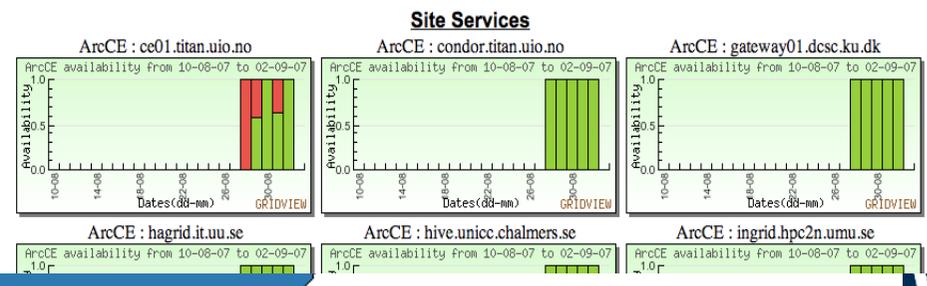
Overall Service Availability for site NDGF-T1 : Daily Report



Individual Service Availability for site NDGF-T1 : Daily Report



Service Instance Availability for site NDGF-T1 : Daily Report

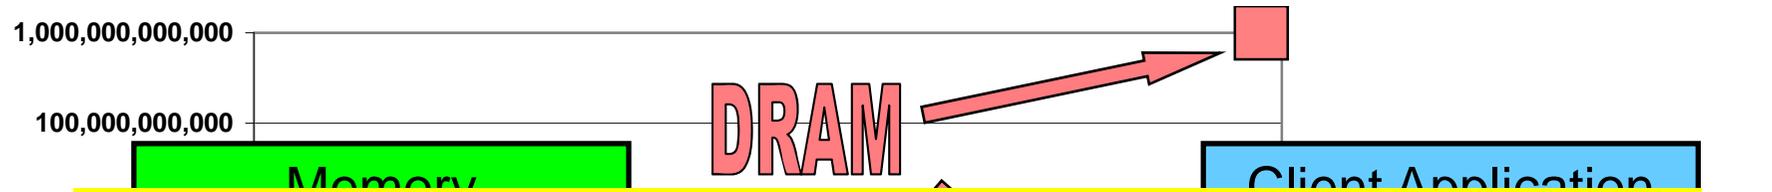


II. Storage Systems

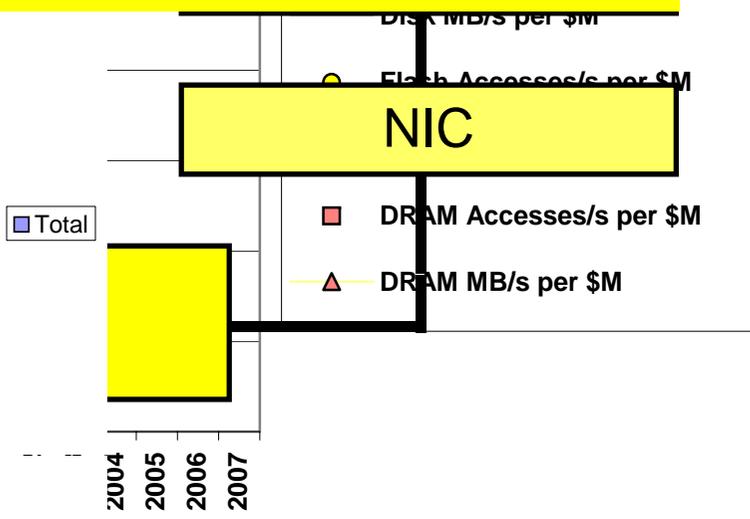
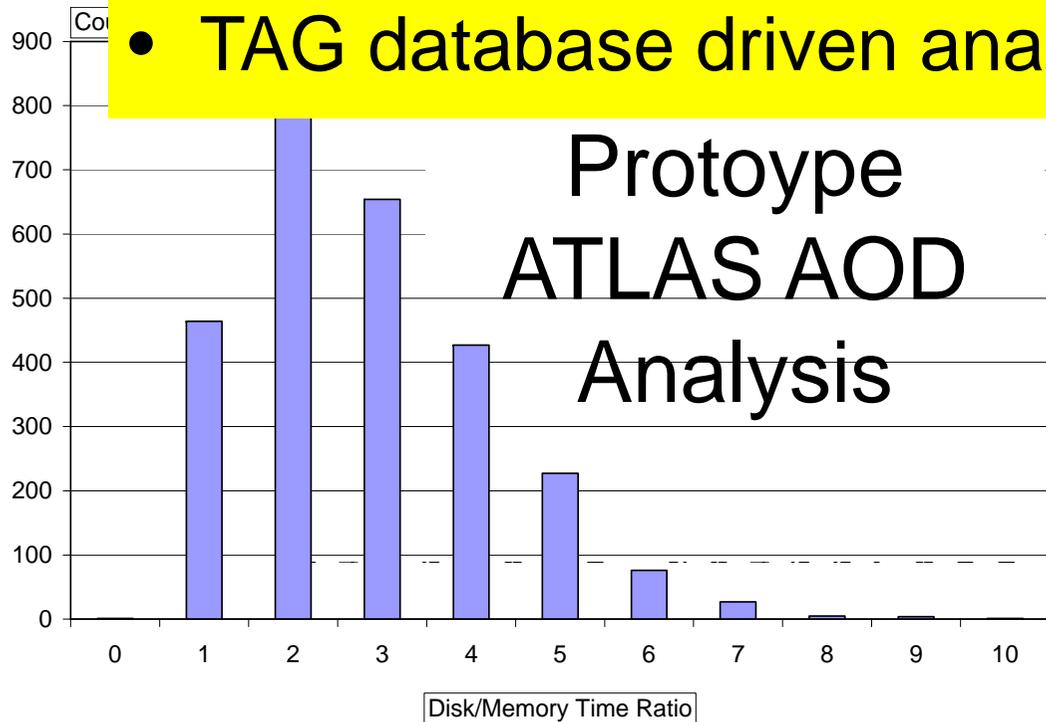
- PetaCache (Richard Mount) ←
- Castor2 (Giuseppe Lo Presti)
- gStore (Horst Göringer)
- UltraLight (Paul for Rick for Jan for Harvey)
- dCache vs xrootd (Alessandra Forti)



Price/Performance Evolution: My Experience



- Significant, but not revolutionary, benefits for high-load sequential data analysis – as expected.
- Revolutionary benefits expected for pointer-based data analysis – but not yet tested.
- TAG database driven analysis?



III. Tools and Services

- CMS Challenges (Ian Fisk)
- Quattor (Michel Jouvin)
- GGUS (Torsten)
- SAM (Antonio Retico)



eGEE SAM (simplified) architecture

All the four LHC experiments are running (or planning to run) custom tests using the production instance of SAM

Goal: sanity checks against selected grid and application services.

CMS, Alice, LHCb

running custom tests in production using two different submission approaches

Atlas

running standard tests in production using Atlas proxy. preparing to submit custom tests

The production SAM platform is supporting the four VOs

Only minor changes were needed to support Alice

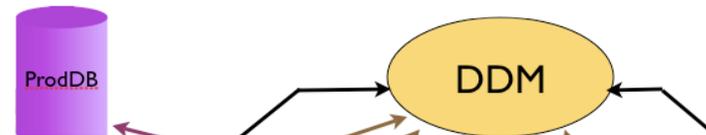
The screenshot shows the Gridview web interface. The browser address bar displays 'http://gridview.cern.ch/GRIDVIEW/same_index.php'. The left sidebar contains a navigation menu with the following items: 'Grid', 'What do you want to do?' (with sub-items: Central Service, Aggregate Site, Tier-1 Site Av, Tier-2 Site Av, Site Detail Av, SAM Test Res), 'Defining VO', 'Service', 'Tier-1 Site', 'Sites for VO', 'Tier-2 Site', 'Current Status', 'Hourly Report', 'Daily Report', 'Weekly Report', 'Monthly Report', 'From 15/4', 'To 15/3', 'Display G', and 'Set Def'. The main content area shows a table with columns: 'at 05:13:34 on 23/01/2007', '10', 'ok', 'Test Env', and 'Detail Result'. The Windows taskbar at the bottom shows the start button and several open applications: 'Gridview: Visualizatio...', 'https://gridview.cern...', 'gridview_MWG_253a...', 'gridview_chep', 'gridview', and 'CERN Alert Notification'. The system clock shows '10:31 AM'.

IV. Production Experiences

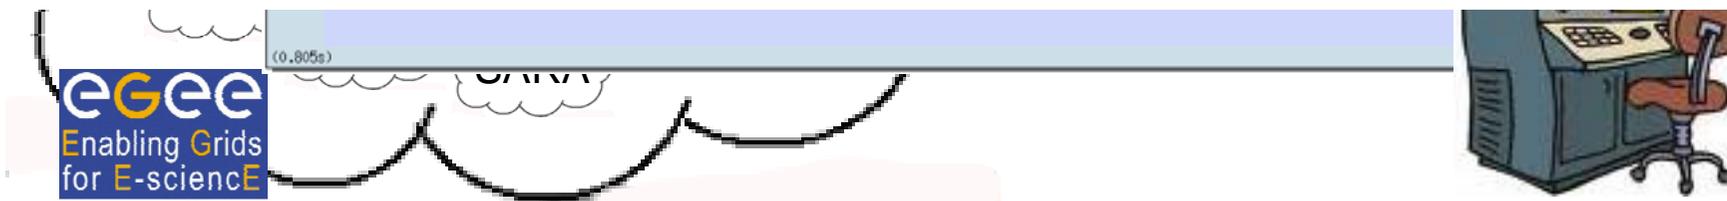
- 3D in production on LCG (Dirk Duellmann)
- ATLAS production on EGEE (Xavier Espinal)
- CMS MC production Jose Calama
- ATLAS production on OSG (Yuri Smirnov)
- CMS production system (Dave Evans)
- ATLAS T0 software suite (Luc Goosens)



The ATLAS Production System



- **Production system could cope with production ramp up challenge (Nov06-Mar07)**
 - 60M events produced in first quarter of 2007 (coexistent with the disk shortage)
- **Job and WCT efficiencies has been improving almost continuously:**
 - ~90% efficiency in WCT and ~60% for job efficiency
- **Monitoring pages are the key for successful job monitoring:**
- **Next phase is to start automation at higher level.**
- **Operations group and shift system has demonstrated to be extremely necessary.**



ATLAS MC Production

on OSG with PanDA

Walltime Usage (CPU days) since Jan 2006

30000
28000
26000
24000
22000
20000

- OSG production has achieved robust operations model
- PanDA works well
- Overall, progressing well towards full readiness for LHC data

18547, 8% 25183, 11%

(0.259s)

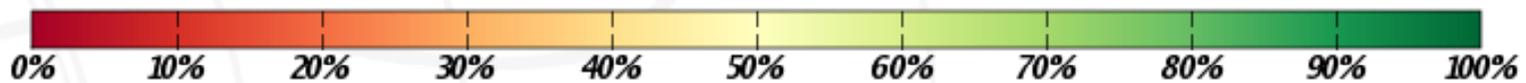
Date

Yuri Smirnov



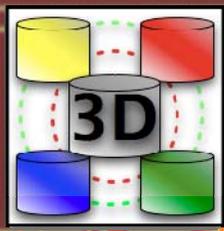
Reached scale of more than 20K jobs and 65 Mevt/month with an average job efficiency of about 75% and resource occupation ~ 50%

Production is still manpower intensive. New components being integrated to further improve automation, scale and efficient use of available resources while reducing required manpower to run the system

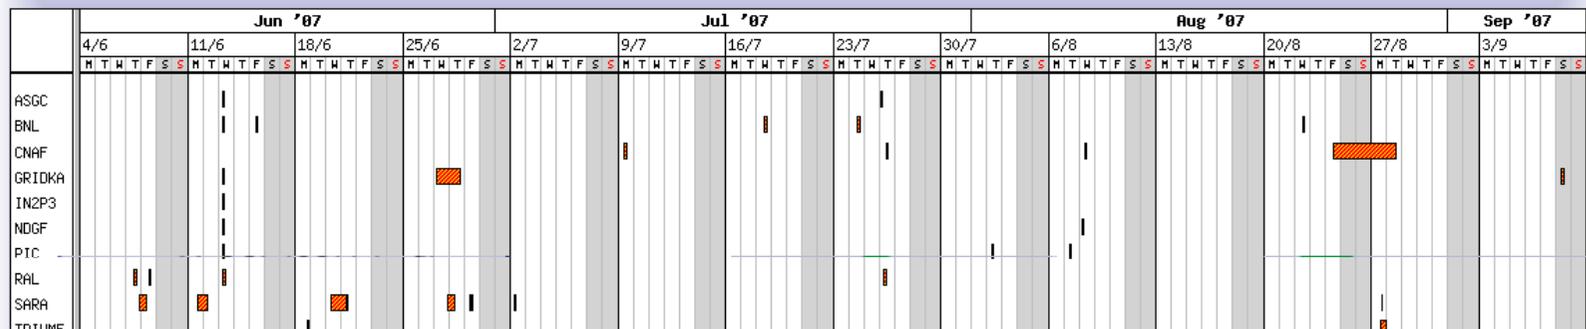


- UCSD (111/029)
- Estoma (10/7859)
- Legnaro (7/3122)
- INFN (595/39)
- WARSAW (52/1/8)
- London_RHUL (522502)
- KNU (435035)
- T2_DESY (434636)
- LIP-Lisbon (421406)
- T2_Belgium (407054)
- BUDAPEST (375924)
- Belgium_UCL (209831)
- RWTH (99463)
- Bari (57272)
- Rome1 (35939)
- Unknown (18033)
- ASGC (50)

PSS

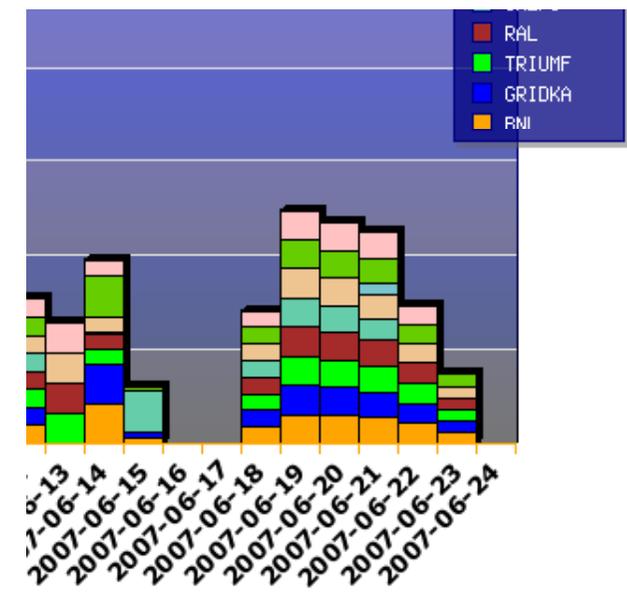
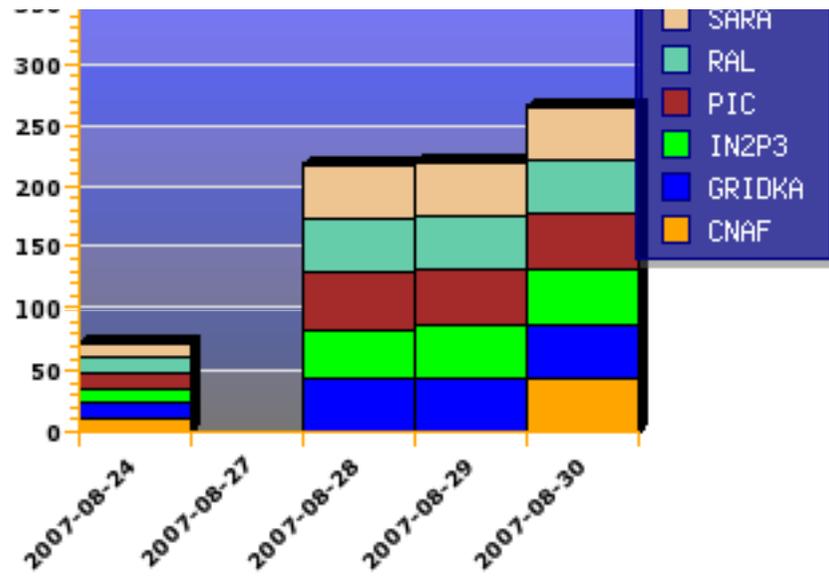
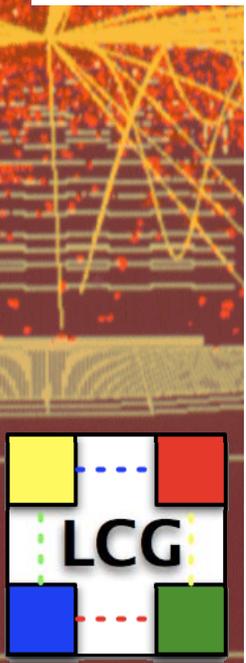


Intervention dashboard



The LCG 3D project has setup a world-wide distributed database infrastructure for LHC

- Close collaboration between LHC experiments and LCG sites
- with some 100 DB nodes at CERN + several tens of nodes at Tier 1 sites this is one of the largest distributed database deployments world-wide



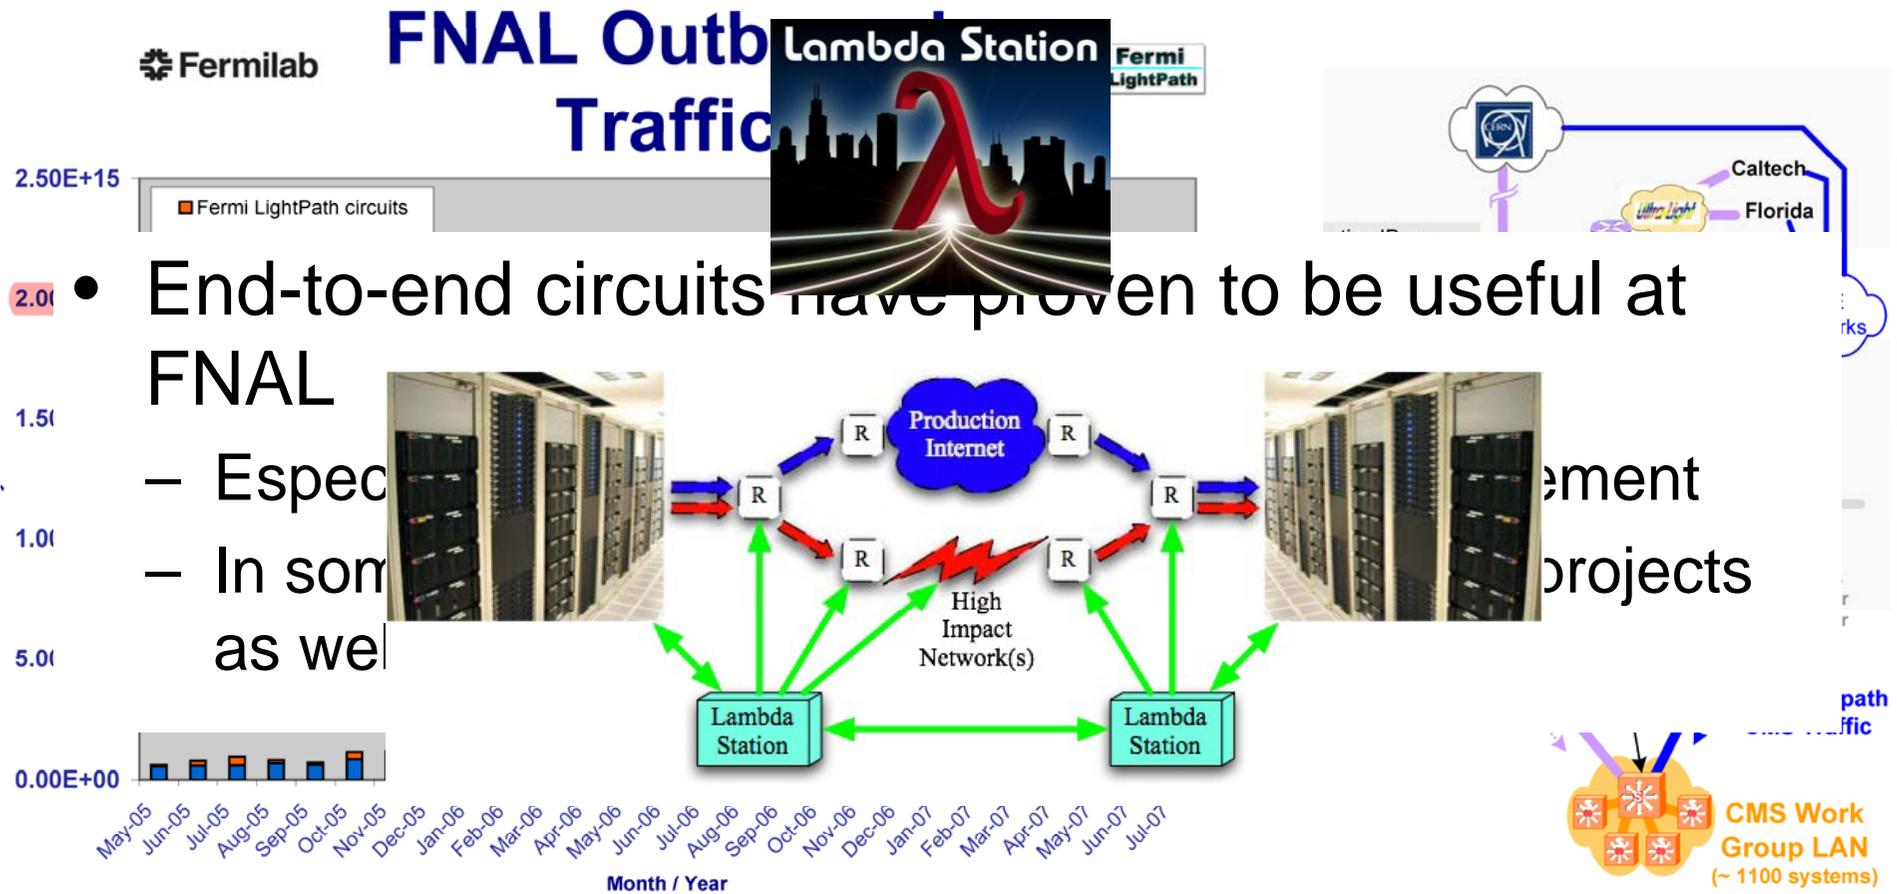
V. Networking

(Matt Crawford's session)

- Use of alternative path WAN circuits
- Lambda Station
- IPv6
- perfSonar



Making the E2E circuit routing work



- End-to-end circuits have proven to be useful at FNAL
 - Especially for experiment projects
 - In some cases as well



But When?



- IANA's last block of addresses is estimated to go 19 Mar 2010.
- Regional registries' last blocks: 10 Oct 2010 – “10/10/10.”

Will IPv4 end then? Of course not.

VI. Operations and Storage

- EELA (Lucas Nellen)
- ATLAS DDM (Alexei Klimentov)
- CMS Comp.Operations (Daniele Bonacorsi)
- Storage @CNAF (Luca dell'Agnello)
- Castor2 Operations (Miguel dos Santos)
- SRMv2.2 in dCache (Timor Perelmutov)



VII. Other things

- SunGrid for STAR (Jérôme Lauret)
- gLExec (Igor Sfiligoi) ←
- gLite WMS (Simone Campana) ←
- CE and SE in Xen (Sergey Chechelnitzkiy)
- dCache @BNL (Zhenping Liu)



Pilots and the Grid ⁽³⁾

Gatekeepers know only about the pilots



- **Pilots are becoming a reality**
 - But are introducing new problems
- **Yesterday's Grid security mechanisms are not enough anymore**
 - Site admins want fine grained control
 - Pilot owners want OS level protection
- **gLExec can help solve the technical problems**

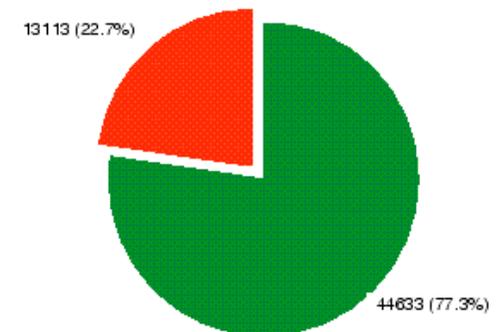
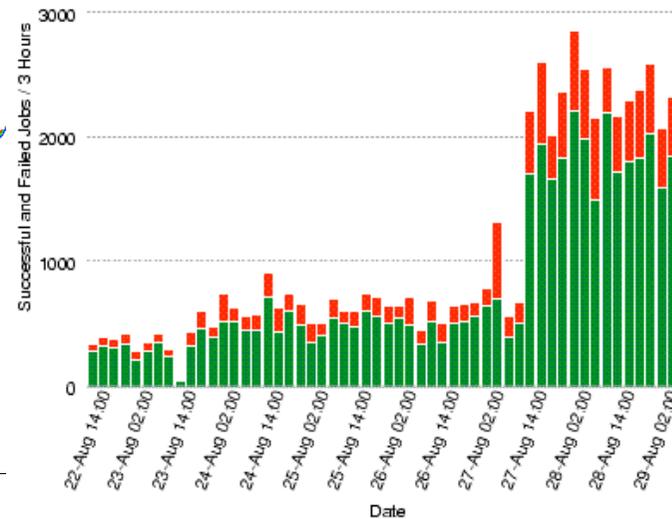
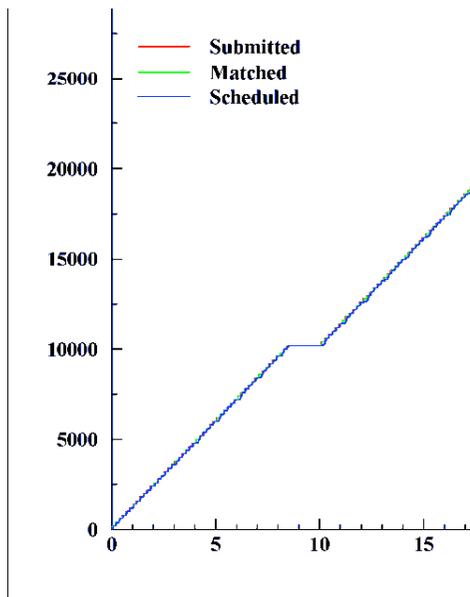
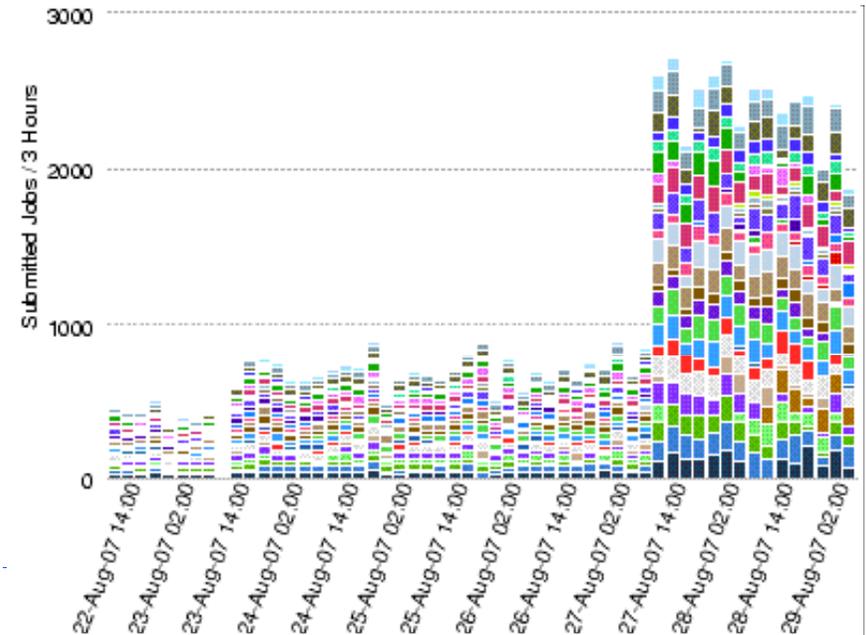


of job scheduling

Pilots and all users in different sites share the same local identity



- CMS supports submission of analysis jobs via WMS
 - Using two WMS instances at CERN with the latest certified release
 - For CSA07 the goal is to submit at least 50000 jobs/day via WMS
 - The Job Robot (a load generator simulating analysis jobs) is successfully submitting more than 20000 jobs/day to two WMS



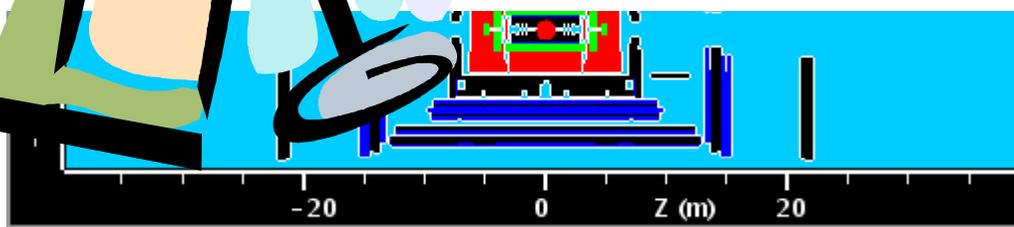
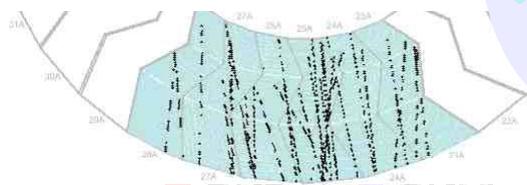
■ Number of Successful Jobs ■ Number of Failed Jobs

ATLAS M4 Cosmics Run

Tracks in the muon chambers in
And in the TRT



This last week of August ATLAS has shown to master for the first time the whole data chain: from a measurement of a real cosmic ray muon in the detector until an analysis in real-time analysis in sites in Europe and the US with all steps in



Final Remarks

- (too?) many WLCG papers
- Other grids: SunGrid, EELA
- 3 ATLAS, 3 CMS presentations
- No Alice or LHCb, 1 STAR, no Tevatron
- Much attention for services
- Storage is HOT !



Thank You !

