

RAL Service Challenge Meeting
27/01/05

Present:

Andrew Sansum(RAL), Andreas Hirstius(CERN), Roger Jones(ATLAS), Martin Bly(RAL), Ian Bird(CERN), Steve Traylen(RAL), Derek Ross (RAL), Tim Barrass(CMS), Kors Bos(SARA), Jeremy Coles(RAL), John Gordon (RAL), Lionel Schwartz(IN2P3), Laurent Caillat-Vallet (IN2P3), Pete Elmer(BABAR), Lassi Turra(CMS), Nicola Pezzi(UCL), Giuseppe Lore(INFN), Doris Ressman(FZK), Bruno Hoefft(FZK), Vlado Bahyl(CERN), Tiziana Ferrari (INFN), Stefano Zani(INFN), Nick White (CERN)

Jamie Shiers – Summary of experiment computing models

The models all represent things differently. We are looking at the real differences between different models and trying to extract commonality. For Service Challenges, start with commonality in models, and then try to include the differences.

There was a discussion of questions and uncertainties in the models that stop us building a “metamodel” right now. It is critical to clarify this now.

Roger Jones: LHC4 is trying to “smooth out” the spikes in the computing usage of LHCb and ALICE.

Lassi Tuura: Is the 7.6Gb/s the actual traffic on the line?

Jamie/John Gordon: No. That includes a x2 factor for capacity. Real traffic can peak at 50% of the stated number.

Networking – List of questions from NREN/Dante of what they’d like to see:

- Breakdown of capacity needed and on which route
- A compilation of Tier2 sites

Kors: Sent out to the GDB list a request for Tier-2s. Some info back, Italy still missing.

Relucance for Tier-2s to stand forward.

Roger: In the US the sites are still tendering to be Tier-2s so they aren’t know yet.

Network links are not ‘public internet’. Restricted to well-defined traffic on well know hosts at both ends.

Tiziana: The base assumption is that it’s a dedicated link. Why?

Jamie: There will be other public links – but they shouldn’t be confused with everything else. Guaranteed bandwidth is important.

John: If you’ve a 10Gb to a tier-1 – where does the Tier-2 work?

Jamie: Tier-2 is not included here. It’s extra.

James: Tier2 is different quality. It’s ok if the transfer to the Tier-2 slows down

Roger: But it’s important to get back the simulated data from Tier-2 to the Tier-1 for archiving.

Kors: What is the ramp-up plan from the experiments until 2007?

Jamie: It’s not in the models.

Roger: ATLAS is trying to do the intermediate planning in place this month.

John: Who will do the similar spreadsheets for Tier-2 to Tier-1

Jamie: Useful for someone else to do this too for a cross check.

John: And is it generic Tier-1 considered

Roger/Jamie: The old spreadsheet had the difference in the Tier-1’s in it.

James Casey – SC1 report

Service Challenge was run with SARA. Problems seen at many levels – hardware, network, software. Fermi declined to carry out challenge. FZK currently in progress. SARA achieved 54MB/s for 3 days. SC1 did not succeed. We have a lot of work to do for SC2 & SC3!

Kors Bos – Test Plans in Amsterdam

Want to look forward to what to do now after SC1, in order to achieve what is necessary. Replace teras in the system by a cluster from UVA (10 dual opteron). Plan to use “only” 1GE between CERN and SARA except for a test in June.

Test Schedule:

February:

- Using only 1GE connection
- Need the radiant cluster
- Goal is >800Mbit/s transfer rate
- Need better communication
- VRVS, etc...
- Send Sander Klous to CERN

Also try to spec and buy new hardware.

March:

- Disk-Disk performance tests (SC2)

Install new hardware

April:

- 4 dual CPU servers, 200GB RAID1, dual GE, dual FC

- New set-up tests (reproduce SC2)

- Try to show an integrated storage environment to the world

- Try new software: SRM on DMF (preferred) or SRM plus Dcache on DMF

May

- First tape tests

June

- 10Gbit tests - disk to HSM (not to tape) > 500MB/s

- After 10/6 switch back to 1Gb/s for July challenge

- Vancouver 10GE tests

July

- Service Challenge

Ian: What SRM

Kors: Ron Trompert has expertise in the area.

Ian: He should sit with J-P Baud, and discuss the issues.

James: Implementing SRM is not easy – we have expertise which can give you a good headstart! And test suites exist.

Reda: Why are you moving back from 64 bit to 32 bit?

Kors: the opteron is loaned hardware. And the Xeon 32 bit kit is off the shelf. It's also a test environment, so these machines aren't the ones we'll use in the end.

James: It's good to get dates specified – it would be good to get other sites to say the same thing.

Vlado Bahyl – CERN data services vs. service challenge

Different from previous presentations – it covers the existing CERN production systems, and gives ideas on how they can be used in the service challenges.

Castorgrid is main WAN data transfer service. Recently added 8 new nodes and doubled the connectivity (2 x 1Gbit per 4 nodes). Only basic TCP stack tuning.

Kors: What is this cluster – not the radiant cluster.

Vlado: This is the production system to castor right now.

It's easy to fill the pipe right now even though there isn't much usage. Now it's dedicated switches for the service and also there is extra bwidth on the routers.

Three lessons learned:

- Load balancing – DNS alias corresponds to three lowest loaded hosts.

- Certs – problems with hosts vs service certs.

 - Used to do double DNS lookup and IPs didn't match. (fixed in globus-ftp-control >

1.11)

John: In what version is this fixed

Vlado: Sometime in mid 2004.

- User mapping: all users from one experiment are mapped onto a group account.

 - Confusing for user e.g. home directory access

 - For exceptions, local map file is used (castorgrid specific)

Also there is a dcache system for experiments small files and a CERN LCG2 SE.

Network connectivity. As compared to radiant setup, there is only 2Gb external network from castorgrid as opposed to 10bit.

Tapes: to get 300MB/s in July, it would require 10 STK 9940B. This is 20% of total drives, and would be a full powderhorn required for the bandwidth.

The other option is to use new IBM 3592 drives. These can do 40MB/s in steaming mode. This gives us 320MB/s. There is room for 100 tape cartridges.

Kors: Is is disk to tape or disk to tape?

Ian: We should do disk to tier-1 tape, as well as to tier-2 tape. Also from tier-0 tape needs to be tested.

John: Problems seen transferring to CERN.

Ian: It is around the firewall.

Vlado: But it's still limited

Kors: Pleased to see the aggressive schedule, but will it effect Tier-1 schedules?

James: It shouldn't – the connection to tape is on a third system at CERN - i.e. "production radiant" cluster.

Andrew Sansum – Tier-1 local infrastructure for Service Challenges

SC2 team – shared load among several staff (~5 split across networking, LAN and hardware deployment, local system tuning, dCache and grid interfaces)

Also expected to call on support from gridpp storage and gridpp network optimization

Currently disk server with 4 fs, 2 luns are exposed per RAID device. These are exported via NFS.

Ideally deploy production capacity for service challenge. Plan B is to deploy batch nodes with extra drives. Device throughput is ~ 120-140MB/s up to 32 threads

dCache – several instances in production. About 17TB configured – 4TB used. Wish to use dCache as SRM interface to Service Challenge 2. Current deployment is dCache pools deployed on lightweight headnodes, and talk to backend stores which have NFS. Maybe NFS will be removed in the future with PNFS getting the files direct from the disk servers.

LAN – network will evolve in several stages. Choosing low cost solutions to minimize spend until it's needed. Currently using Nortel 5510 @ 1Gbit with 1Gbit interconnect. Next stage is to move to 10Gbit interconnect (also Nortel 5510) Soon 4 x 1Gbit to UKLIGHT

MSS – preparations for SC3 (and beyond) underway since August 2004. Until now, not been hitting the tape system hard. Review known limitations and see how the hardware performs. Limits in ADS tape systems, but all can be fixed in software. Good knowledge of how the system reacts to load. Estimate suggest 60-80MB/s to tape now. Putting in more cache disks will probably double it.

Robin Tasker – Networking to the RAL Tier-1

Currently 2 * 1Gbit out of site to TVN. It goes through a 8Gbit/s rated firewall 2.5Gbit in TVN and 2.5Gbit peering SJ4/Geant

UKLIGHT. Has a pop in UCL with 10G to Starlight Chi and Netherlight AMS. Also has links between UK sites. Nominal 10G capacity.

Use of UKLIGHT is for service challenge @ 2 x 1GB into UKLIGHT to go via netherlight to CERN. The service bandwidth is guaranteed when you get it, but if the network is overloaded, you negotiate when you have it!

Eventually...

Will move to "Terabit RAL Backplane Infrastructure" – 04/2005.

m*10G + n*1 G access to TVN-2

TVN-2 - 11/05

SJ5 Core – early 2006

Peered to GN-2 (2005/6)

Multi-10Gbit into CERN

Kors: When using the netherlight system – will you need 2 gig or 10 gig from AMS-CERN?

Robin: 2 gig.

Site Reports

INFN – Stefano Zani

This is a short update from Luca's presentation last time. WAN Access now up to 2Gbit/s dedicated to SC from Bologna to Milan (gigapop)

Mar 05 10 Gb/s End 05 dark fiber.

LAN – bought switch.

Hardware – bought 11 x dual opteron. Shipped next week.

At first start with internal disk (2 x 73GB U320 SCSI HD 10KRPM)

Then move to SAN

IN2P3 Lyon – Laurent Caillat-Vallet/ Lionel Schwarz

Network 1Gb/s link to Renater. Renater Backbone 2.5Gb. Geant 10Gb to CERN

Between July/September – 1Gbit to RENATER + 10Gbit dark fiber to CERN

MSS – HPSS as SRM, RFIO as front-end. Remote access through srm, bbftp, gridftp
Storing customization of HPSS for exp. (rfio, service classes).

Move to HPSS 5.1 next week.

dCache SRM for HPSS – in test for last five months.

Pros: customizable pool attractive scheme, gridftp load-balancing, async migration.

Cons: not easy to install, impossible to debug
Test status: HPSS config ok, and migration scripts ok.

Plans for July Challenge

Further dCache tests – sustained transfers, cache cleaning, HPSS breakdown simulation, Tape only COS – CERN-Lyon SRM transfer expected 02/05

Hardware infrastructure

- 2 pool nodes with 2TB each (1-day buffer)
- 9940 cartridges

FZK – Bruno Hoelt

Currently running service challenge. Running at 4Gbit/s for ~18hours. 3rd party transfers. 4Gbit/s from 24 nodes. (6 nodes get 1.9Gbit/s) When we run 1 node to 1 node we see 70MB/s with 3 processes. If we run another node at the same time, the transfer rate goes down. Writing data to SAN directly (GPFS)

Tiziana: Could be good to try with other application to see if it is a gridftp problem.

James: Would be good to rerun the iperf tests.

UK Tier-1 and Tier-2 Storage – John Gordon

Overview of gridpp and sites in the production system. Gridpp provides funding for middleware development. It includes storage. Developers at RAL, Edinburgh, Glasgow.

Goals:

Provide SRM interfaces to

- Atlas petabyte storage at RAL
- Disk (Tier 1 and 2 in UK)
- Disk Pools (for Tier 1 and 2 in UK)

Provide packaging and support

Options

- RAL-SRM – an SRM interface to ADS developed from EDG-SE
- dCache
- DRM – from LBNL
- Dpm

Short term timeline is to provide a release of SRM to disk and disk pool by end of January 2005. It was planned to coincide with gLite release. Priorities changed to production in light of service challenges - Focus on dCache. Long term strategy is possibly a dual solution.

Actively interworking with other SRMs. Doing cross testing and development. Working within SRM Collaboration and GSM within GGF.

Other deployments –

Manchester – dCache with read pools on worker nodes.

Edinburgh – 24TB of raid array exported via NFS. Looking at dCache.

Pointing other communities towards SRB.

Detailed planning for Service Challenge 2 – Jamie Shiers

Detailed planning is not too detailed! But all sites (but for Fermi) are here, so we can move on. Try to test each site individually, and aggregate at 500MB/s

Ian: Need same level of detail from all sites as for SARA.

Target 100MB/s disk-to-disk:

	Rate	Date	
NIKHEF/SARA	100MB/s	End March	
FZK	> 100MB/s	Mid Feb	
IN2P3	100MB/s	End March	
INFN	100MB/s	End March	
Fermi	> 100MB/s	End March	
RAL	100MB/s	End March	

We need drafts of plans when sites start.

Ian: We should have a plan from each site by end of next week.

CERN: We need to write the description of how the service challenge is run.

Kors: How do we improve monitoring and communication.

Jamie: perhaps sites need to send someone to CERN for part of the time during the challenge.

Jamie: Can we handle this by email?

Ian: Phone conference week after next.

Kors: One day meeting at the end of February – at CERN. 24th February.

Kors: What monitoring?

Ian: A single place with a global rate, and individual rates.

Kors: Who's writing the document.

James: We will write the monitoring bit.

RAL and UK Plans – John Gordon

Lots of CPU resources unused. Planning has been done at the Tier1 for all LHC and non-LHC experiments. Experiments request allocation, and that is broken down into quarterly allocations by the User Board. CPU isn't a hard limit, just guidance to the batch scheduler. Disk is allocated directly to the experiment in terms of disk space. Tier-2 resources are also allocated.

Issues:

- Moore's law and other hardware unknowns

- Long lead time on MSS acquisitions

- Over complication, reliability and maintainability of middleware

- Conflict between experiment demands

- And understanding the computing models

- Balance between 2008 and 2005-7

- Buying hardware that isn't used weakens the case for hardware that will be required in 2008

- Often there aren't enough resources in 2008 for what the experiments want

Kors: Which Tier-2's if we want to include them by the end of the year?

John: Tier-2s are a management entity, and all appear as individual LCG-2 sites. We should ask each Tier-2 to nominate a site.

Roger: Lancs would be interested, since they have a person who is funded to look at networking issues.

Jamie: We probably only want one from a few countries for this year, but next year we need to get thirty onboard. They'll have to come on board quickly.

John: How does the network topology work? Is the Tier-1 dedicated and there fore could it be overloaded by all the Tier-2s. RAL has a person tasked to look at the models and work out the Tier-2 load.

Roger: We can get a good first cut at the numbers for ATLAS quickly.

Jamie Shiers – Long term planning

Not a plan being presented, but the plan for the plan. It's a huge task with a short timeline. Many things need to be delivered in the next few weeks (in terms of planning documents).

Trying to visit all European sites by Easter. In April, do all the non-european sites on the back of the Taipei service challenge meeting. Also need to address Tier-2s. Try to work through Hepix?

Progress so far - High level overview, "Requirements" document, based on summary of the Computing Models, Detailed planning for SC3 is starting now, global milestone document (up to 2007)

Initial schedule is for major cycle of 3 months, driven by service challenge with monthly meeting (perhaps only for GDB reporting)

At CERN, planning has started with all the groups involved – GD, GM, ADC, FIO, CS. And discussions have started with the experiments too for SC3. Ongoing are the discussions with the sites as well.

Timeline – Official target is April 2007, but the target is earlier, since we'll have cosmics and calibration earlier. So, we've a lot less than two years. We're basically in production mode now, and can't stop 'til 2020.

Discussion

Kors: Maybe we shouldn't focus on the point SC's but get into continual running mode. As soon as a site is doing something, they shouldn't stop. We don't want to leave the network early.

Lassi: For CMS, that's what we do. Let's get it out earlier and use it as much as possible.

Jamie: They all overlap – there's no short term challenges.

James: The systems should always be trickling into the production systems.

Lassi: Yes. The functionality should move into the next level service after each service challenge.

Jamie: We all agree this is the model. What we're looking at in terms of peaks of effort for the challenges.

Kors: Do we have these services right now?

James: Yes – we have castorgrid and we have the radiant service. But they're different and we need to converge.

Lassi: But the experiment timescales are different. We need to do TBs from tape in March.

John: But how does this interact – will you use the SC stuff or is this separate? There's only one set of effort at the Tier-1s.

Lassi: We'd like the SC knowledge to move into the production system. We have a working system but it's not long-term operable system. We'd like the SC's to get the services into a state that just work – e.g. gridftp, File systems, tape systems. That will give us the reliability that we need.

Ian: What Lassi has said is exactly what James said.

Lassi: We don't need to go to 500MB/s, 50MB/s is good enough – the 3MB/s that we see is bad!

Bruno: From our view we had different infrastructures. One for SC, and one for production.

John: And at RAL, we've different lightpaths.

Ian: Lightpaths is not the problem.

Andrew: It's the software stack that is the problem.

Lassi: What was good at FNAL is that the the dCache system they used in the SC is now the production system for CMS. This is good.

Jamie: But the timescales is different. March is very aggressive.

Lassi: We've our own timescales –we need to deliver a certain amount of data to the experiments by the summer.

Jamie: Can we reset the timescales? When the timescales were set, this was before the LCG effort. Can we coordinate these timescales?

Lassi: Physics TDR is what we need to achieve by summer.

Jamie: We need to go back and re-negotiate. I think that we've the same milestones. We just need a schedule for the ramp-up.

Lassi: That's fine, but the agreement needs to be at a higher level.

John: If you can get it out of the production system right now, it's ok. But if the production infrastructure doesn't do it, there won't be manpower to fix it, but there will be a channel to ask for things later – that's the service challenge.

Lassi: We need the system by the end of 2006.

Jamie: This is what we have – full model validation at nominal data rate, scaled down to 2006 CPU is the target for April 2006. We have production service, and service challenges.

Ian: Lets me clear between middleware and service. We need the middleware in the next few months. The service might change, but the middleware will just slowly evolve. We need it right now. The service challenge is to raise the system to a new level, with the knowledge feeding back into the production system.

Lassi: We see ~30% outage of systems in PHEXED. That's gridftp, SRM, etc...

Ian: And that's what the SC should try and fix.

James: We do the bottom bit of phedex.

Lassi: You can use phedex to run the service challenge.

Jamie: We do want to do that. CMS should be one of the experiments in SC3 and should be using the PHEDEX system.

Lassi: SCs are trying to do things - Reaching 24x7 services and trying to increase the throughput. From the physics point of view, the 24x7 aspect is more important.

Ian: That's understood. We believe that too, but we need to increase the data rate too.

Jamie: We need to show we can do this at a high reliability level with minimum manpower. Also we can change the numbers if possible.

Kors: when can we have a productionv system in place.

John: We have it right now.

Lassi: The bottleneck is not castorgrid – it is in other places.

Bruno: We can duplicate what we set up since we need to think about security at our end.

John: But if we build a 50MB infrastructure what happens when all the experiments want it.

Lassi: Currently we are 98% of the total CERN infrastructure.

John: With SG-5 we won't have the network 'til 2006

Andrew: We can find bits of networking to drive the hardware right now to find bottlenecks.

Jamie: Right now, the experiments don't have the systems that need all the data, but they do need the reliability.

Kors: I think we should stop saying that we don't have the network – we need to put the schedule down for what we need, and then go and ask for it.

Andrew: Credibility issue, since we don't use the links.

James: We can use the SC's to do this – we can fill links if needed.

Kors: We need to not diverge, go to the GDB and ask for the transfer plans. Can ATLAS provide this?

Roger: We can go and ask.

Lassi: need to talk to David. We're in continual production mode now, so we can tell you what we're doing and what we expect to move.

John: What rate is babar at?

Andrew: It's a dribble – it's about 4TB a month. The rate bottleneck is at SLAC for transfers due to the tape robots. 35TB capacity came online in October, but there's still some space.

Peter Elmer: We'll have a look.

Kors: What about Tier-1 to Tier-1 and Tier-1 to Tier-2?

Jamie: If we get to the Tier-0 to Tier-1 by Feb GDB, and if it's accepted then we put in some effort after that.

Lassi: We will be doing some of this in Feb/Mar. Some problems in protocol matching (SRM/plain gridftp)

Kors: Doing some of this for D0 right now from SARA to FNAL.

Jamie: We need to spend a month to understand the model so we know how much these rates are?

Kors: Will there be any analysis at Tier-2s in this year?

Roger: We're trying to do some stuff. It would pull from Tier-1s

Jamie: This is why we need a first idea of the topology that we need.

Roger: If you take the ATLAS model, there isn't a huge traffic from individual analysis. Mostly it's from group analysis.

Jamie: To do the modeling we need to know the number of Tier-2s for a Tier-1.

Roger: We approximate to 3.

Jamie: The network guys need to know the real topology with identified links.

Lassi: We have a topology in PHEDEX right now. It's not a star topology. We also have ways to send streams to alternative sites, and then that data would get sucked when the original site comes back up.

Jamie: We need to test using this model. This is all complicated, so we have to know that our bits work under the model.

Lassi: We do have some of this right now.

Kors: Now what about the other five tier-1s. Nordugrid, PIC, BNL, FNAL, Vancouver, Taipei.

Reda: Planning to buy the disk hardware pretty soon. We don't have specific links right now.

Kors: Can you be ready for your test slot via Amsterdam?

Reda: Yes.

Kors: And you can do disk tests before June 10Gbit link.

Reda: yes.

Kors: Do we want to go to nominal rate by 2006 – since we need twice nominal for end 2006.

Jamie: Don't we need to have better numbers for the rate first. It's not clear right now.

Kors: And tape? Should we go for twice nominal to tape or to disk.

Ian : tape is important – we need to test since experiments will affect it a lot with file size.

Kors: But do we need to test in 2006?

Ian: Don't know.

Peter: will the sizes increase?

Roger: 2GB is ok for ATLAS.

Ian: ALICE is a problem, since they split.

Peter: LHCb will have the same problem as CMS with taking 1% of vents from a file.

Ian: We need to know the list of problems we'll have. File size is only one of them.

Kors: Ian & Kors talked about a supercomputing demo. We should really do the publicity.

Peter: The problem isn't the network, it's the GUI.

Roger: CERN doesn't have a presence.

Kors: It should be LCG! We should be able to get a bit more manpower. It's 12-18 November.

Roger: One problem is that we'll take effort from the other people who already have stands.

Peter: It's a lot of effort.

Ian: But it's worthwhile. SLAC and FNAL get big gains.

Kors: We should discuss it at the next GDB.

Next Meeting:

24 Feb @ CERN

15 March @ Lyon

26 April @ Taipei