



Conference Conclusions

LATBauerdick/Fermilab

Computing In High Energy Physics 2004

Interlaken, Switzerland



Conference Conclusions

LATBauerdick/Fermilab

Computing In High Energy Physics 2004

Interlaken, Switzerland

Plan For The Talk



- ◆ NOT a summary of the summaries
 - ◆ a big thanks for the excellent summary talks!
- ◆ Try to present selected highlights and messages
 - ◆ charged to “interview whoever you want — it is also good to get input from the floor, ie young researchers might have a different view”
 - ◆ my personal observations and views expressed to me and discussed in the Blog and Forum — apologies for the “LHC perspective”

blog

n : a shared on-line journal where people can post diary entries about their personal experiences and hobbies [syn: [web log](#)]

Source: WordNet ® 2.0, © 2003 Princeton University

blog

On the Internet, a blog (short for [weblog](#)) is a personal journal that is frequently updated and intended for general public consumption.

The essential characteristics of the blog are its journal form, typically a new entry each day, and its informal style.

The author of a blog is often referred to as a blogger.

Source: WhatIs.com

Plan For The Talk



- ◆ NOT a summary of the summaries
 - ◆ a big thanks for the excellent summary talks!
- ◆ Try to present selected highlights and messages
 - ◆ charged to “interview whoever you want – it is also good to get input from the floor, ie young researchers might have a different view”
 - ◆ my personal observations and views expressed to me and discussed in the Blog and Forum – apologies for the “LHC perspective”

blog

n : a shared on-line journal where people can post diary entries about their personal experiences and hobbies [syn: [web log](#)]

Source: WordNet ® 2.0, © 2003 Princeton University

blog

On the Internet, a blog (short for [weblog](#)) journal that is frequently updated and intended for general public consumption.

The essential characteristics of the blog are its journal form, typically a new entry each day, and its informal style.

The author of a blog is often referred to as a blogger.
Source: WhatIs.com

Plan For The Talk



- ◆ NOT a summary of the summaries
 - ◆ a big thanks for the excellent summary talks!
- ◆ Try to present selected highlights and messages
 - ◆ charged to “interview whoever you want — it is also good to get input from the floor, ie young researchers might have a different view”
 - ◆ my personal observations and views expressed to me and discussed in the Blog and Forum — apologies for the “LHC perspective”

blog

n : a shared on-line journal where people can post diary entries about their personal experiences and hobbies [syn: [web log](#)]

Source: WordNet ® 2.0, © 2003 Princeton University

blog

On the Internet, a blog (short for [weblog](#)) journal that is frequently updated and intended for general ...

The essential characteristics of the blog are its journal form, typically a new entry each day, and its informal style.

The author of a blog is often referred to as a blogger.
Source: WhatIs.com



CHEPtalk "partially successful"



◆ ~1000 hits, ~50 postings **Thanks to everyone who participated!!**



at Sep 30, 17:00

Logged in: LATBauerdick
[Log-out](#)
[Your Profile](#)
[Members](#)

[ChepTalk](#)



STATS

Total entries: 13
Total comments: 30
Total members: 7

0001024

September 2004

[Day 3](#)

[Day 2...](#)

[Video Conference](#)

[About who asks questions in plenaries](#)

["Grid Security"](#)

[Day 1](#)

[Single-Vendor Lock In?](#)

[Finding the Common Themes](#)

[Deploying a farm in 3 months](#)

[Forefront of Computing?](#)

[50 years of computing](#)

[Engineer](#)

[First Blog Entry!](#)

FORUM SEARCH

[POST NEW TOPIC >>](#)

General Discussion						
	TOPIC	STARTER	VIEWS	REPLIES	POSTED ON	LAST REPLY
	Common Theme -- Successful Systems?	LATBauerdick	31	1	Sep 27, 04 3:17 pm	Sep 30, 04 11:24 am
	Common Theme -- How Good are the Grids Now?	LATBauerdick	85	5	Sep 27, 04 3:22 pm	Sep 29, 04 6:04 pm
	Theme: Innovation - response to the Industry talks	Ruth Pordes	16	0	Sep 29, 04 11:23 am	-
	Common Theme -- Interfacing to what others provide	LATBauerdick	33	1	Sep 27, 04 3:40 pm	Sep 29, 04 11:16 am
	SMTP server?	ktf	43	1	Sep 27, 04 11:31 am	Sep 27, 04 3:10 pm
	Forefront of Computing?	LATBauerdick	31	0	Sep 27, 04 10:17 am	-



CHEPtalk "partially successful"



◆ ~1000 hits, ~50 postings **Thanks to everyone who participated!!**



at Sep 30, 17:00

Logged in: LATBauerdick
[Log-out](#)
[Your Profile](#)
[Members](#)

[ChepTalk](#)



STATS

Total entries: 13
Total comments: 30
Total members: 7

0001024

September 2004

[Day 3](#)

[Day 2...](#)

[Video Conference](#)

[About who asks questions in plenaries](#)

["Grid Security"](#)

[Day 1](#)

[Single-Vendor Lock In?](#)

[Finding the Common Themes](#)

[Deploying a farm in 3 months](#)

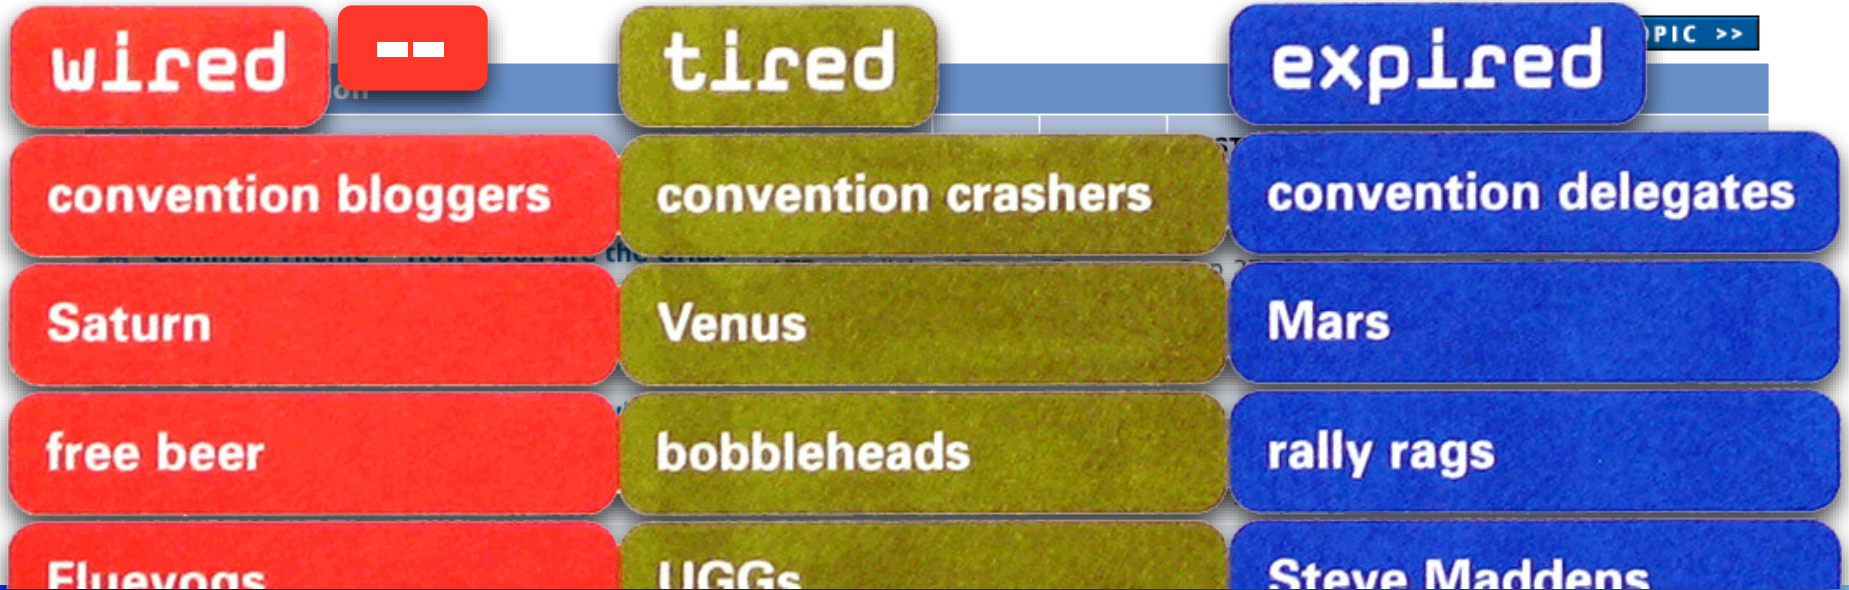
[Forefront of Computing?](#)

[50 years of computing](#)

[Engineer](#)

[First Blog Entry!](#)

550 CHEP participants...



Wolfgang's Welcome Talk



Aim of the Conference

- Learn from running experiments
- Get ready for LHC by making the Grid a powerful and reliable computing resource
- Stay in touch with other sciences
- Have a peek into the future
- Talk to each other, learn from each other and help each other !!
- Celebrate CERN's 50th birthday
- Enjoy Interlaken and its nice environment



Wolfgang: "Learn From the Running Experiments"



Vital Statistics



Vital Statistics	CDF	D0
Raw Data Size (kbytes/event)	205	250-300
Reconstructed Data Size (kbytes/event)	180	200 (20→60)
User formats	25-180	20-40
Reconstruction Time (Ghz-sec/event)	(5)10	50(120)
Monte Carlo Chain	fast	full Geant
user analysis times (Ghz-sec/event)	1 (3)	1
Peak Data Rate(Hz)	75(+)	50(+)
Persistent format	RootIO	D0om/dspack

→ 300 Hz
~2006

Both collaborations continue to evaluate and evolve data formats in response to analysis needs and computing constraints

**D0 computing has a strong production focus
CDF computing has a strong analysis focus**

Amber Boehnlein, FNAL



RunII is Now

Vital Statistics



Approaching the Scales of Data Management and Data Access Problems expected at the First Years of LHC!

Vital Statistics	CDF	D0
Raw Data Size (kbytes/event)	105	250-300
Reconstructed Data Size (kbytes/event)	180	200 (20→60)
User formats	25-180	20-40
Reconstruction Time (Shr-secs/event)	(5)10	50(120)
Monte Carlo Chain	Fast	Full Chain
user analysis times (Ghz-sec/event)	1 (3)	1
Peak Data Rate(Hz)	75(+)	50(+)
Persistent format	RootIO	D0 m/corpack

300 Hz
~2006

Both collaborations continue to evaluate and evolve data formats in response to analysis needs and computing constraints

First Years of LHC!

D0 computing has a strong production focus
CDF computing has a strong analysis focus


Amber Boehnlein, FNAL



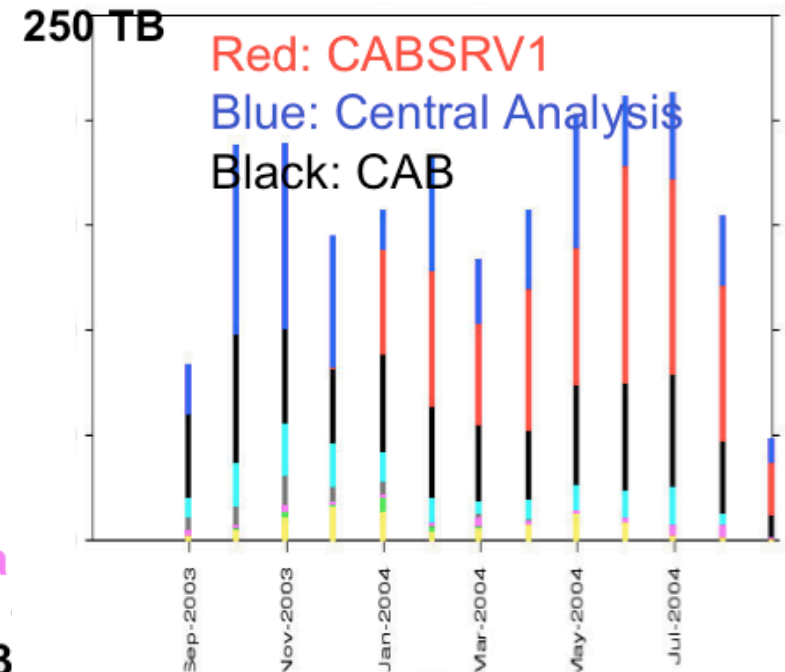
- ◆ End-to-end Data and (structured) Meta-Data Handling
- ◆ High-throughput Disk caching

Direct dCache access at CDF

- ◆ 60 TB/day movement at peak
 - ◆ Currently provides primary access to data
- 60 TB read by CDF clients on 06 (no discernible file delivery errors) <3



nBytes Read Per Day



D0 ALL Stations GB/month

Oct 2003-Sept 2004
 CDF: 1.5 PB 12B events
 DO: 2.1 PB; 50B events

Amber Boehnlein, FNAL



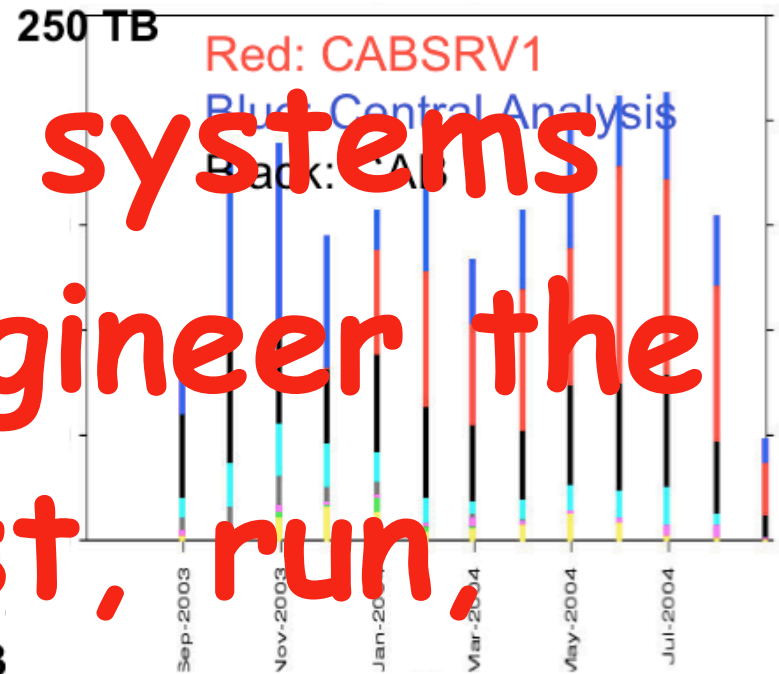
- ◆ End-to-end Data and (structured) Meta-Data Handling
- ◆ High-throughput Disk caching

Engineer the systems

— and also engineer the

tools to test, run,


maintain, monitor them!



Direct dCache access at CDF

- ◆ 60 TB/day movement at peak
- ◆ Currently provides primary access to data

60 TB read by CDF clients on 06/05/04
(no discernible file delivery errors) <3



nBytes Read Per Day

DO: ALL Status: CB Month

Oct 2003-Sept 2004
CDF: 1.5 PB 12B events
DO: 2.1 PB; 50B events

Amber Boehnlein, FNAL

Lessons from BaBar



Peter Elmer:

New Computing Model – CM2



- Large changes in BaBar offline computing
- New eventstore
- New data content
- New analysis model
- New bookkeeping/data distribution/etc.

◆ "BaBar tried the "full frontal assault" approach to computing.

We now try to be smarter."

Discussing with BaBar, RHIC



- ◆ Lively discussion on CHEPtalk and over coffee:
- ◆ "BaBar's talk seemed to be the start of a theme -- getting out of databases. Or at least, complex ones (be they either Oby or RDB)." (Gordon Watts)
 - ◆ Peter Elmer:

"learned how to **use** databases — keep info at the logical layer, and defer the physical details to the lowest possible layer -- instead knowing "all things about all places at all times" — a battle you can never win"
 - ◆
- ◆ "what makes a system successful?"
 - ◆ Jerome Laurent (RHIC):

"running experiments to go first for the practical, usable and scalability"

 - ◆ early requirements estimate (present and projected)
 - ◆ small mixed technical/experiment groups (see Belle)
 - ◆ prototype in use at an early stage
 - ◆ decoupled service [no dependence in other service but possible plugin]
 - ◆ interface standardization
 - ◆ consolidation and scalability in real life scenario, constant feedback ...

And Of Course, the LHC DCs...



- ◆ provide a focal point — we can't do without,
 - ◆ but we can't really do with either....
- ◆ Let's be skeptical and self-critical about what we achieved
 - ◆ but let's try to be very constructive, too!
- ◆ Next: run analysis (input-data-intensive) on the data we have produced on the Grid
 - ◆ Up-front dialog about the goals, requirements
 - ◆ Data Management and Data Handling,
 - ◆ End-user access to remote resources, etc
- ◆ Now it's time to define the concrete computing models for the LHC



How get the lessons into the LHC?



- ◆ Now that the LHC seems to become real:
Great opportunity through working with people at the running experiments that are joining the LHC program!
 - ◆ Involve them in the requirements process
 - ◆ Opportunities: shared participation in LHC and running experiments
- ◆ Try to understand, synthesize, re-factor, interface, copy, use, ... ideas, approaches, components — (Q: re-use code?)
- ◆ CHEP'04 was a great opportunity for that
-- hope to see the results in CHEP'06

Note Added in Proof:



Rick StDenis on CHEPtalk:

"I have heard about learning from the running experiments, but the numbers in attendance where I have been (session 4) seem to be about zero or at best support for the home team.

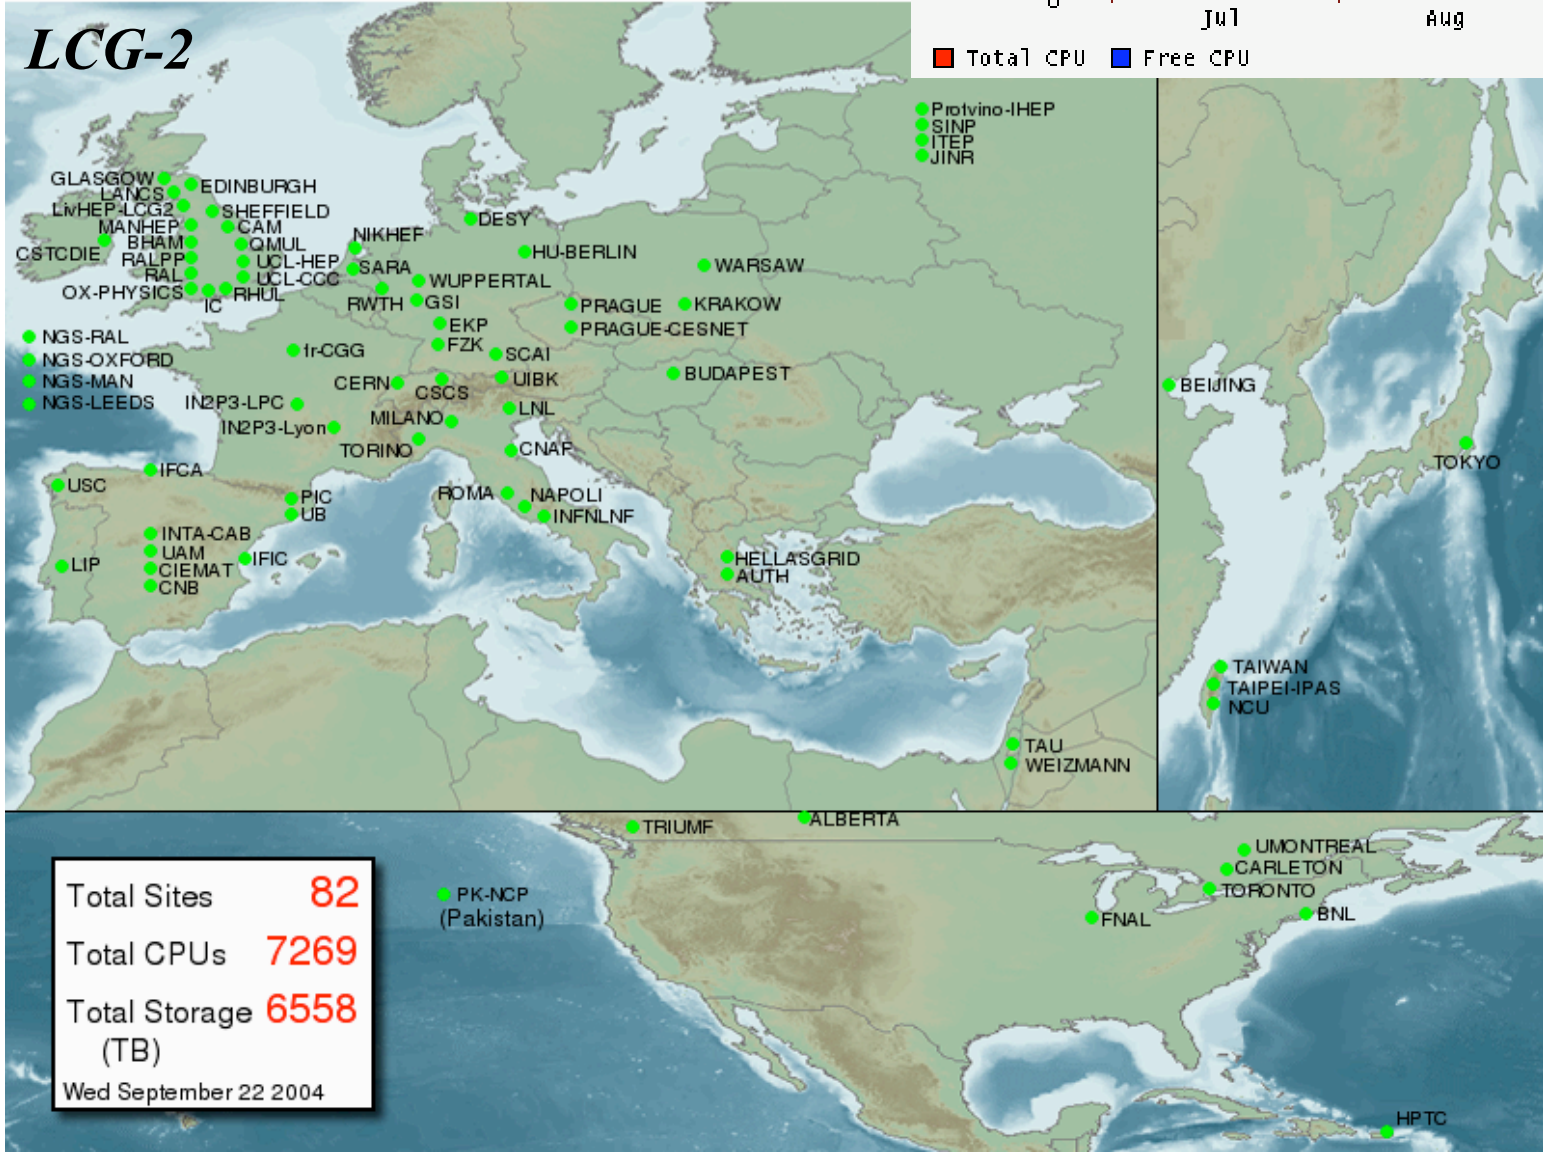
If I take this as voting with your feet, then there is a clear negative vote"



Wolfgang: "Get Ready for LHC by
Making the Grid a Powerful &
Reliable Computing Resource"





LCG-2

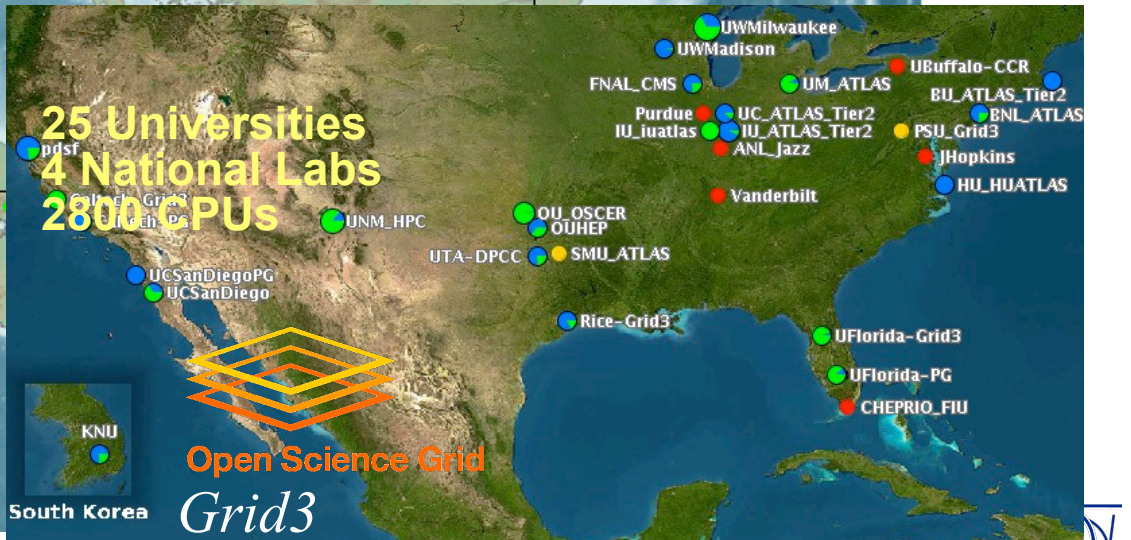


Total Sites **82**
 Total CPUs **7269**
 Total Storage **6558**
 (TB)
 Wed September 22 2004





Total Sites	82	● PK-NCP (Pakistan)
Total CPUs	7269	
Total Storage (TB)	6558	
Wed September 22 2004		







LCG-2

GLASGOW
LANCS
LIVHEP-LCG2
MANHEP
CSTCDIE
RALPP
OX-PHYSICS
EDINBURGH
SHEFFIE
CAM
QMUL
UCL-I
UCL-4
RHUL
IC

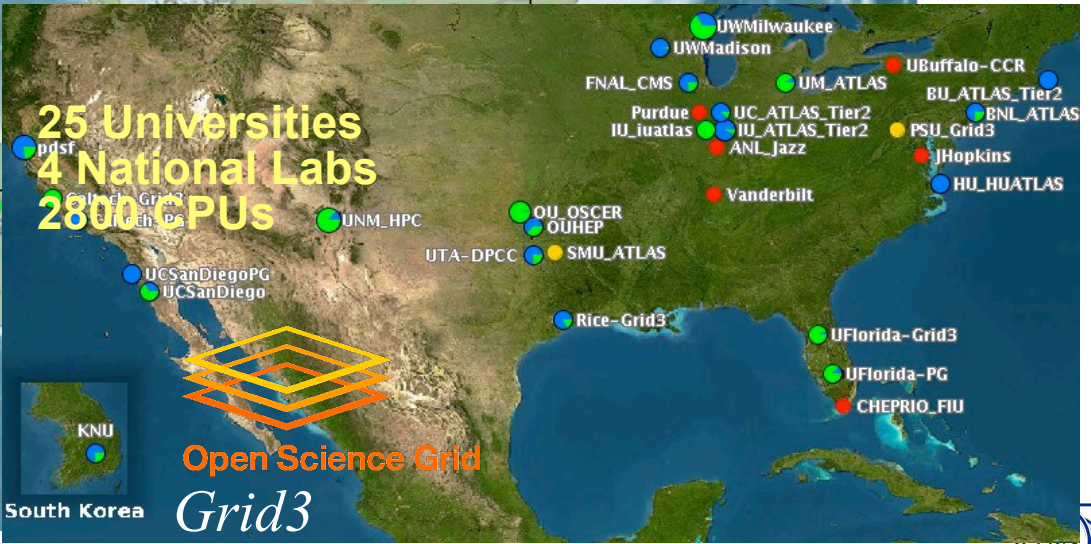
NGS-RAL
NGS-OXFORD
NGS-MAN
NGS-LEEDS


IN2P3-LPC
IN2P3-Lyon
CERN
MILANO
TORINO

USC
IFCA
LIP
INTA-CAB
UAM
CIEMAT
CNB
PIC
UB
IFIC



Total Sites	82	● PK-NCP (Pakistan)
Total CPUs	7269	
Total Storage (TB)	6558	
Wed September 22 2004		





Our Grid Systems are Successfully Enabling Broad Participation

Total Sites	82	● PK-NCP (Pakistan)
Total CPUs	7269	
Total Storage	6558 (TB)	
Wed September 22 2004		



How Successful Are Our Grids?



- ◆ LCG 80 sites, Grid3 30 sites, NorduGrid 30 Sites
 - ◆ typically 65% efficient — is that a good or bad?
 - ◆ Ian Bird: "Grids were successful in making many small problems into one big problem!"
 - ◆ Les Robertson: Many "nodes" on the Grid are in small groups who definitely do their best, but are not (yet) "professionally managed".
 - ◆ David Stickland: classic computer center problems
 - ◆ Rob Gardner (?): "the three principles of the Grid: configuration, configuration, configuration"
 - ◆ Gordon Watts: "GRID doesn't seem that hard — the only hard part seems to be the scaling and the problem of moving data" ...
- ◆ How will experiment(er)s interact with Grid operations "machine"?
 - ◆ Les: "Grids need to serve *our* needs" what are the expectations?
 - ◆ EGEE, OSG, operations centers, user support, MoUs...

How Successful Are Our Grids?



- ◆ LCG 80 sites, Grid3 30 sites, NorduGrid 30 Sites
 - ◆ typically 65% efficient — is that a good or bad?
 - ◆ Ian Bird: "Grids were successful in making many small problems into one big problem!"
 - ◆ Les Robertson: Many "no" definitely do their best,
 - ◆ David Stickland: classic
 - ◆ Rob Gardner (?): "the th configuration, configura
 - ◆ Gordon Watts: "GRID doesn't seem that hard — the only hard part seems to be the scaling and the problem of moving data" ...
- ◆ How will experiment(er)s interact with Grid operations "machine"?
 - ◆ Les: "Grids need to serve *our* needs" what are the expectations?
 - ◆ EGEE, OSG, operations centers, user support, MoUs...

Rob Gardner on CHEPtalk: What people are not talking about is a metric which measure how many CPUs are unused because people are worried about looking bad with failure rates and efficiency. Grids and production frameworks should get dinged for wasting resources, not just achieving 98% end-to-end "efficiency". The metric we focused on was the number of successful jobs per day, not the number of failures



- ◆ CHEP'03: The Grid will be successful if we can make it simple
 - ◆ Manuel Delfino: You can't treat the Grid like a cluster, like you could not treat a cluster as a single computer
- ◆ Interoperability (Les Robertson: Cohabitation) between Grids
 - ◆ posters and talks about federating Grids
 - ◆ several approaches to cope- including LHC data challenges
 - ◆ technical issues like mapping heterogeneity through information services
 - ◆ data management: PhEDEx, SRM2SRB, ...
 - ◆ (even endorsed by EU - see Max Lemke's presentation)
- ◆ Accounting: Bo Anders Ynnerman - this is what makes a Grid tick
 - ◆ for sites who's scientific mission != that of their customers
- ◆ Meta Data and catalogs was much the theme of the last months
 - ◆ Ken Peach: Data -> Information -> Knowledge
 - ◆ The power of meta analysis lies in the quality of meta data



Wolfgang: "Stay in Touch With Other Sciences"



Stay In Touch With Other Sciences



- ◆ Actually - how about our own CHEP science community?
 - ◆ RunII running on LCG, RHIC running on Grid3, etc
 - ◆ joint projects in the US: PPDG, iVDGL, GriPhyN, OSG TG -- Europe?
- ◆ make our Grids available to other (smaller) groups
 - if they even want it — see Bo Anders Ynnerman's talk
 - ◆ A driving force for EGEE and US Grids
- ◆ Driver to require flexibility in Grid configuration
 - VOs on demand, dynamic workspaces:
 - ◆ approaches, principles and technologies
 - ◆ but will help autonomy of smaller entities within large experiments!!
- ◆ Several presentations and posters show how to make Grids and HEP computing more accessible to smaller groups -- "inclusion"



Wolfgang: "Have A Peek Into the Future"

Have A Peek Into The Future

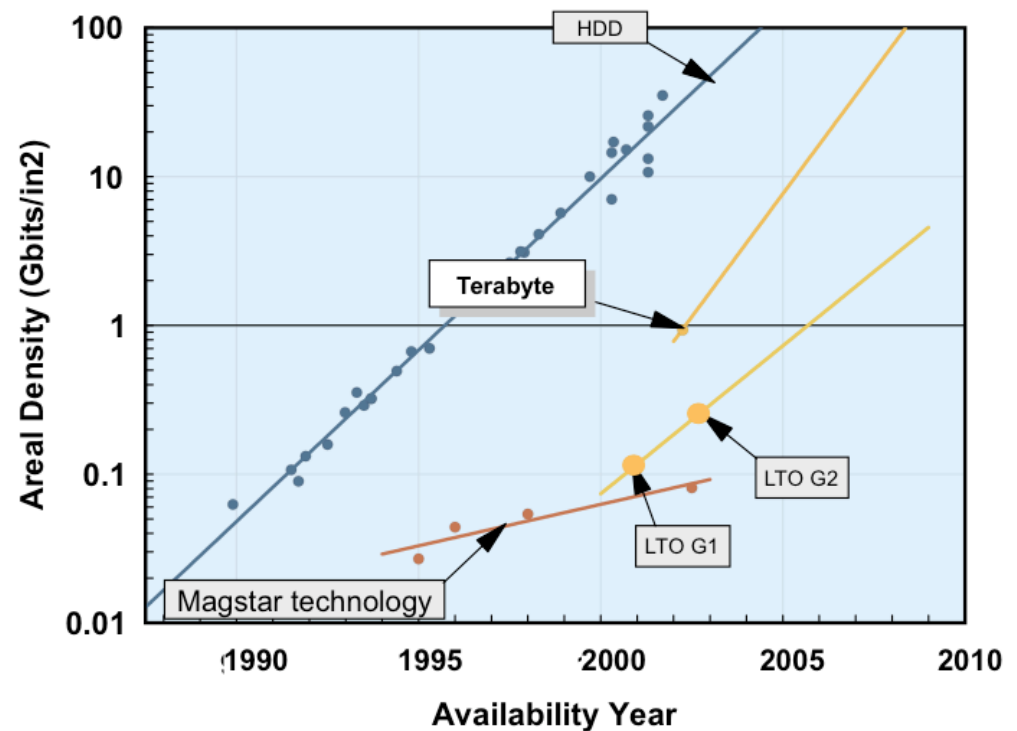


- ◆ I liked most industry presentations — a panel to put perspective?
- ◆ Throughputs and Bandwidth:
 - ◆ RHIC 500MB/sec archiving, Castor tape throughput 450MB/sec
 - ◆ dCache disk data serving CDF 60TB/day ~ 10Gb/sec
 - ◆ IPv4 Networks record 7.1 Gb/sec over 10k
 - ◆ end-to-end service challenges to put it all together
- ◆ 2.5-10Gb/sec networks "everywhere"
 - this side of the digital divide!
- ◆ new opportunities for the experiment's computing model
 - pushing data, remote archives, remote data access, ...
- ◆ -- but how do we include the other (dominant) part of the world?
 - ◆ clearly this has to be part of the experiments computing model
 - ◆ requires pro-activeness for computing managers, funding agencies, ...



- ◆ Jai Menon/IBM
tapes density roadmap
- ◆ disk error critical at 1PB
 - ◆ mirrored: 2/year data loss
 - ◆ new RAID factor 2 disk
- ◆ storage system management
 - ◆ costs 2-3 of procurement
 - ◆ current: ~TB/admin

Tape Roadmap



How About System Performance?



- ◆ CPU performance lags, except for 64 bit architectures
 - ◆ Sverre Jarpe: "Are your codes 64-bit clean?"
-> work for the experiments!
- ◆ Multi-core architectures - even more high density cooling
 - ◆ below 90nm, passive power starts overtaking active power
-- hope for another "disruptive trend"?
 - ◆ more parallelization required
 - ◆ optimize performance on the whole system and software stack!
- ◆ "Artificial Intelligence", "Ambient Intelligence", "Intelligent Systems and Agents", "Learning Systems"
 - ◆ it requires an intelligent approach...
 - ◆ control systems, e.g. BTeV, JLab run control, etc



Wolfgang: "Talk to each other,
learn from each other and help
each other !!"

We Require Better Communication



- ◆ Communication in a distributed community
 - ◆ LCG—EGEE—Experiments—IT—US Grids—...
 - ◆ need to embrace anything that help better communication
 - ◆ technology can help! It does cost money!! Please put it into your budgets!!!
 - ◆ communicating with remote peers requires to be pro-active:
 - ◆ pick up the phone, get an AIM screen name, write down your thoughts in a blog, document you work, write little proposals
 - ◆ but also: preserve the bandwidth...!
 - ◆ “connected to a room full of people” syndrome
 - ◆ we should not shy away from controversial discussion
 - ◆ with people you have met and know reasonably well...
 - ◆ M.Delfino’s blog: questions from the floor: always the same? first 10 rows?
- ◆ we want and need to collaborate in an atmosphere of trust!

Too Much Politics?

CHEPtalk sound bites:



- ◆ “Not to name names, but I also heard the second time that a design choice was made “for political reasons”” (Gordon Watts).
- ◆ “Keyhole approach” to building systems
 - ◆ is it lack of trust or inability to get engaged, or genuinely optimized choices, that system builders go down to the low-level interfaces, instead of the higher-level middleware services ?
 - ◆ was the ARDA RTAG “end-to-end” recommendation misunderstood?
- ◆ “We spend effort interfacing at low levels of the software stack we may not have the opportunity to think about and address interesting and germane problems at the higher levels - the lightweight dynamic VOs; semantic methods for defining, evolving and managing meta-data as examples.” (Ruth Pordes)

We require respect and appreciation for quality

Too Much Politics?

CHEPtalk sound bites:



- ◆ “Not to name names, but I also heard the second time that a design choice was made “for political reasons”” (Gordon Watts).
- ◆ “Keyhole approach” to building systems
 - ◆ is it lack of trust or inability to get engaged, or genuinely optimized choices, that system builders go down to the low-level interfaces, instead of the higher-level middleware services ?
 - ◆ was the ARDA RTAG “end-to-end” recommendation misunderstood?
- ◆ “We spend effort interfacing at low levels of the software stack we may not have the opportunity to think about and address interesting and germane problems at the higher levels - the lightweight dynamic VOs; semantic methods for defining, evolving and managing meta-data as examples.” (Ruth Pordes)

We require respect and appreciation for quality

Remember: Trust is “earned” and “invested”!

Too Much BS?



- ◆ David Williams, being asked about the role of HEP at the forefront of computing:
 - ◆ "[Computing in HEP] is mostly about actually getting it to work [HEP] is a most wonderful laboratory for computing ... [in HEP] you don't get away with BS!"
- ◆ and people here seem to agree -- on the long run!
- ◆ I believe, despite good achievements, LHC computing is in a crisis (-- and maybe that's normal and expected)
 - ◆ I'm not sure it's a clear scientific mission (yet) that is driving us
 - ◆ and there seems a lack of overall vision how to achieve our goals

David Williams:

People and their interactions

- I suspect that the LHC experiments are running at the limit of what is feasible...
 - Not the amount of funds that can be assembled
 - Nor the complexity of the detectors
 - But the possibility of keeping such a large number of very smart people working enthusiastically and actively towards a common scientific goal
-
- Until the mid-1980s HEP's "computing problem" was often thought to be about **obtaining enough processor power**
 - Then we worried about **storage capacity**
 - The real problem has always been, in my opinion, **getting people to collaborate on a solution**



Miron Livny:

Can we do it?

Does the scientific community (scientists from different disciplines and the funding agencies) have the know how, the resources and the will to develop, maintain, document, evolve and support a common (and shared) suite of production quality middleware that meets the needs and expectations of the HEP community?

www.cs.wisc.edu/condor



Condor

Middleware More "Accessible"



- ◆ Condor flying high --- beyond Miron's talk...
 - ◆ Appreciate concrete and real engagement of Condor team
 - ◆ Many concrete examples for use of Condor as middleware glue
 - ◆ often with help from Condor team and changes: CDF CAF, SAM/JIM, ... see talks and posters
 - ◆ HEP has direct access to this technology
- ◆ gLite/EGEE middleware prototype, people get engaged
 - ◆ architecture paper was "a good read" (almost...)
 - ◆ first experiences were presented
- ◆ OSG blueprint activity in the U.S.
 - ◆ principles and practices, reference to gLite/EGEE architecture



Miron Livny:

Yes,
we can and should
do it!



- ◆ More discussions on CHEPtalk...

First Experiences with gLite...





Wolfgang: "Celebrate CERN's 50th Birthday"


50 years of computing



- ◆ CERN Computing Generations...

David Williams:

Happy Birthday,
CERN!!




50 years of computing



- ◆ CERN Computing Generations...

David Williams:

Happy Birthday,
CERN!!





Wolfgang: "Enjoy Interlaken
and its nice environment"



Interlaken and its Wonderful Surrounding...



We actually walked up here...!



...Its many Activities...



I was just doing
some research on
Random Numbers!



... its Sometimes Surprising Sightings...



- ◆ motto: "delightfully tacky, yet unrefined"...



... its Sometimes Surprising Sightings...



- ◆ motto: "delightfully tacky, yet unrefined"...



see <http://www.hooters-interlaken.ch/EN/aboutHooters/>



... and Just a Wonderful Place
to have the CHEP Conference!



Thank you, CERN!



- ◆ Stimulating program and well-thought out schedule
 - ◆ including many very interesting poster
 - ◆ Superb organization, facilities, networks, food, drink
 - ◆ Wonderful "birthday party" on Wednesday
 - ◆ Thanks to all that were involved
- For A Very Successful Conference!

Thank you, CERN!



- ◆ Stimulating program and well-thought out schedule
 - ◆ including many very interesting poster
- ◆ Superb organization, facilities, networks, food, drink
- ◆ Wonderful "birthday party" on Wednesday
- ◆ Thanks to all that were involved
For A Very Successful Conference!

See you next time at CHEP'06