



# Report from HEPiX Meeting

Fall 2008 Taipei (Taiwan)

<http://indico.twgrid.org/conferenceTimeTable.py?confId=471>

Thomas Finnern  
DESY-IT



# Event Outline



- **GRID Camp Sa/Su**
- **HEPiX Meeting Mo-Fr**
  - Key Note Fred Baker (CISCO)
  - Site Reports
  - Technical Topics
  - Numbers
    - 100 Participants (5 DESY)
      - 33 from Asian Countries
      - 66 from Western Countries
      - 1 from Australia
  - DESY Talks and Sessions
    - “Site Report” from Wolfgang Friebe
    - “Virtualization Experience with Hyper-V” from Reinhard Baltrusch
    - “Virtual Network Services” from Thomas Finner
    - “Usage of Request Tracker at HEP” from Sven Sternberger
    - Virtualization Session convened and chaired by DESY

# HEPiX Group Picture



- **Workshop and Intro to Taiwan Grid**
  - by Academia Sinica Grid Computing ...
    - WeiLong Ueng, YungChung Chang, Hain Yu (Rachel), Chen Ethan Lin, Jims Hsu (Cheng Hsin Hsu)
  - ... and Italian Grid
    - Giuseppe La Rocca
- **1st Day**
  - EGEE and gLite Middleware: Authorization and Authentication, Workload Manager, Job Submission, Data Management, Metadata Catalog, API: AMGA, DMS, Grid File Access)
- **2nd Day**
  - Tools Suite: Parallel Processing (MPI, Distributed Analysis Environment DIANE, ...), Job Definition and Management GANGA and Grid Application Platform GAP

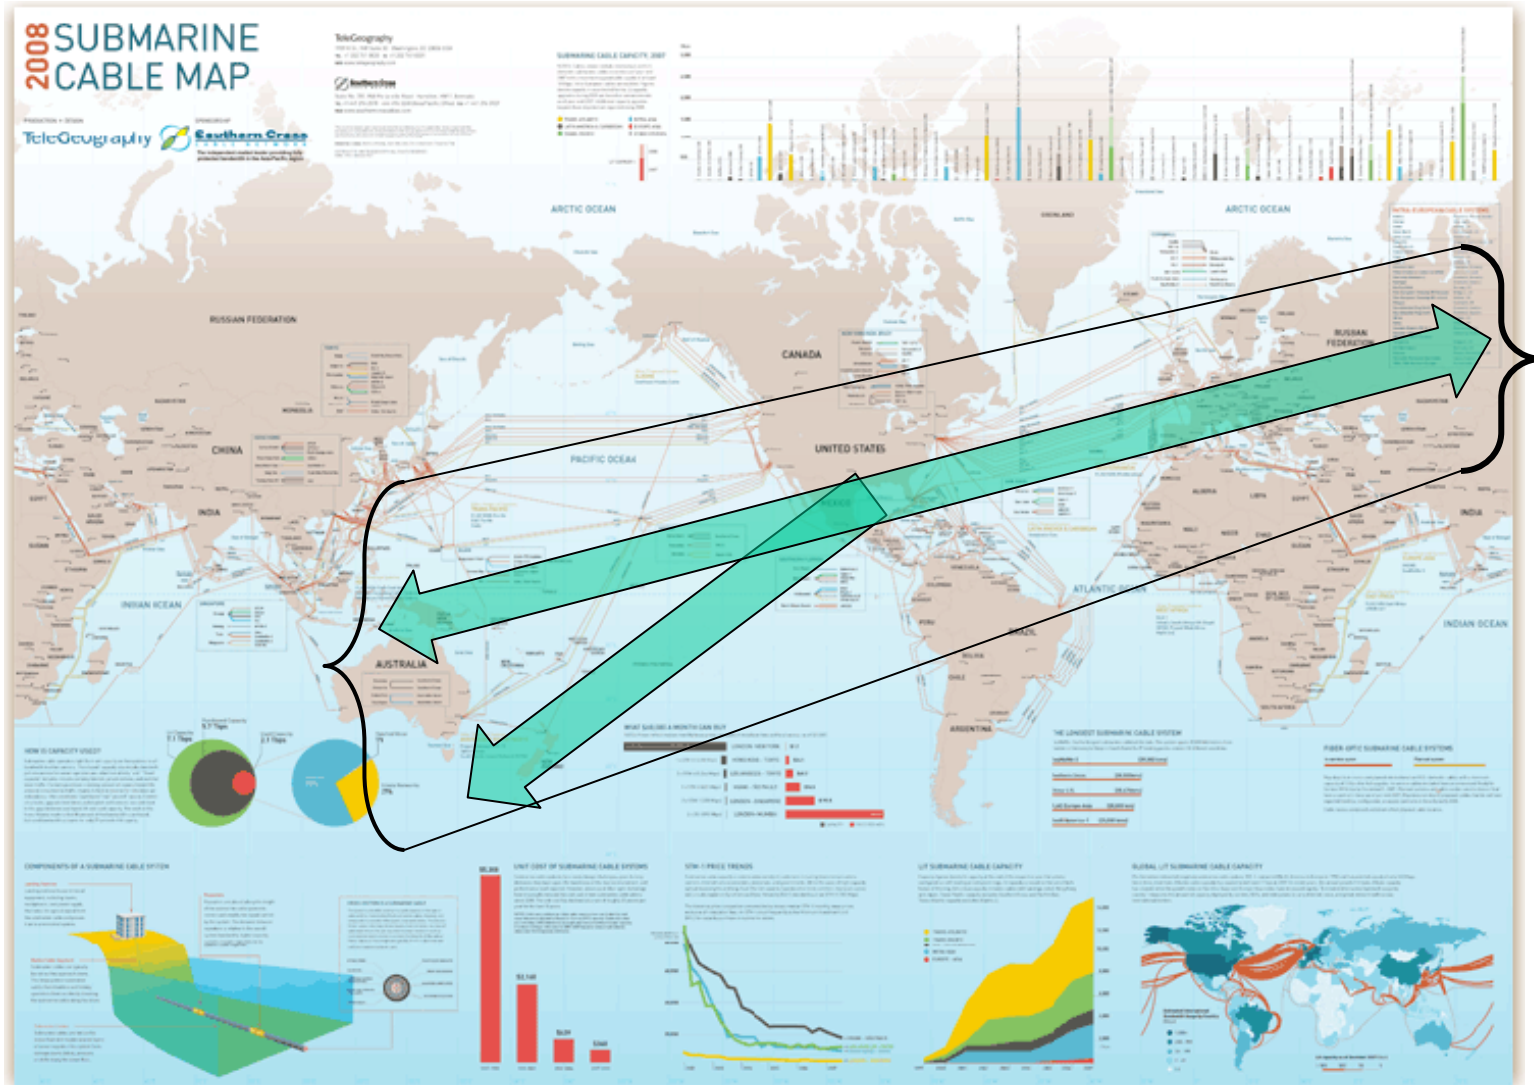


- **Fred Baker (Asian Supporter from US)**
- **Computing Trend History**
  - 1950's: Computers often standalone systems
  - 1960's: Batch Computing, Service Bureaus
  - 1970's: Time Sharing
  - 1980's: Departmental computing, Workstations
  - 1990's: Age of the laptop
  - 2000's: Data centers
  - 2010 and beyond – Service Centers?
- **Today's GRID is Cloud Computing**
  - “A style of computing where massively scalable IT-enabled capabilities are delivered as a service to external customers using Internet technologies.”
- **Later: ip4 to ip6 transition**

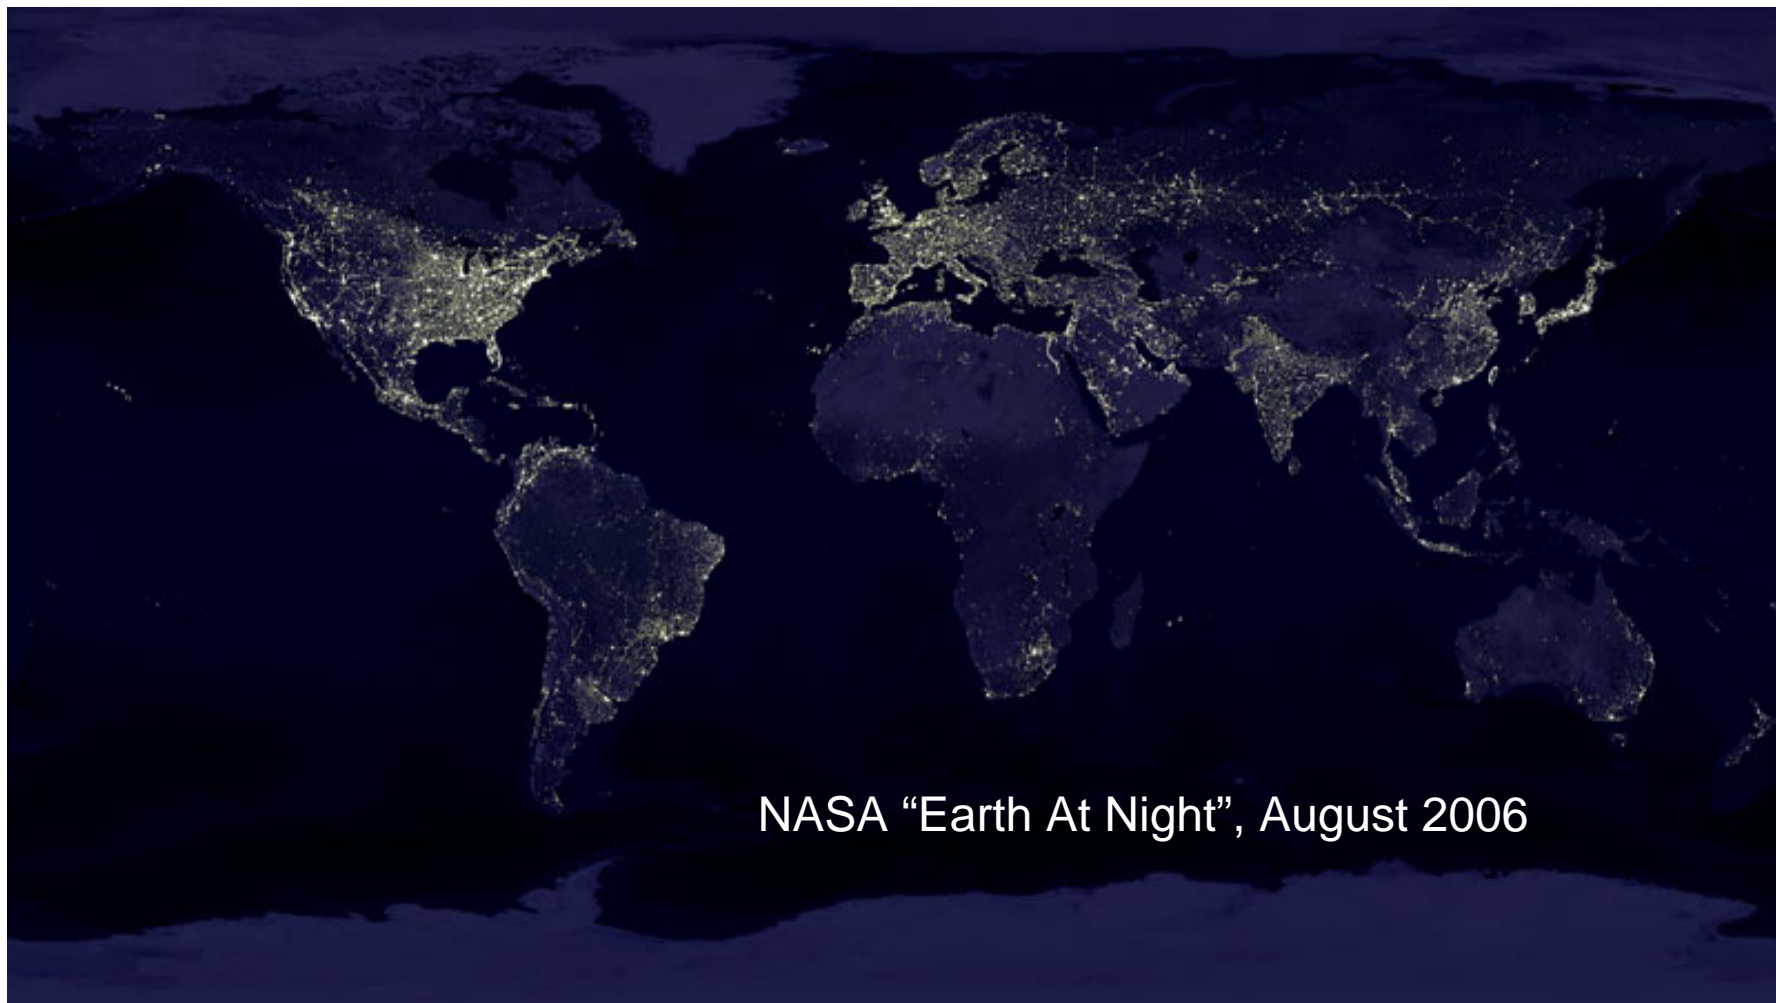
# Where is the broadband Internet today?

## The Europe/America/East Asia/ANZ fiber corridor

Map copyright 2008 TeleGeography



Power, and by extension money, throughout the world



NASA "Earth At Night", August 2006

# IP Addresses throughout the world today



# Site Reports Summary



- **More Data**
- **More „Green Computing“**
- **More Consolidation**
- **More Standardisation ?**
- **Lot of 10 kW Racks (lower edge)**
- **ITIL(IT Infrastructure Library) + ISO20000:**
  - Differing Versions
  - FERMI, CERN?, ASGC

# Site Reports I



- **CERN**
  - Linux Only now
  - Central Registry
  - New CC in Plan
- **FERMI**
  - CC Consolidation
  - Upcoming 10th anniversary
  - Still Strong Kerberos
  - ISO 20000
  - New Exchange Server
- **ASGC**
  - Academia Sinica
  - Host
  - T1/T2
- **ICEPP (Uni Tokio)**
  - T2 / SLC / Quattor / LSF 6
- **INFN**
  - T1/T2 LSF
  - CC Consolidation
  - IPv6 WG
- **UNI Melbourne**
  - Tim Dyce
  - T2 / T3
- **SLAC**
  - Alf Wachsmann
  - Now: SLAC National Accelerator Laboratory
  - N.N. -> Alf -> Randy (-> Chuck)
  - Photon Science
  - Black Box #2
- **RAL**
  - Martin Bly
  - [J.user@stfc.ac.uk](mailto:J.user@stfc.ac.uk) -> [Joe.User@stfc.ac.uk](mailto:Joe.User@stfc.ac.uk)
  - Vmware 4 linux on laptops
  - dCache -> Castor
  - New: 24x7 service
  - CC extending (1.8 PByte, ...)
- **KEK**
  - Shigeo Yashiro
  - Tokyo, Tsukuba, Tokai
  - CC expanding
- **CCIN2P3**
  - Philippe Olivero
  - T1/T2
  - 75 People (40% nonpermanent)

# Site Reports II



- **DESY**
  - Wolfgang Friebe
  - GRID active
- **GSI**
  - Walter Schön
  - CC out of scale(-money,heat+)
  - Lustre 4 all farms
- **KISTI**
  - Christophe Bonnaud
  - ALICE Korea
- **TIFR**
  - Nagaraj Panyam
  - CMS India
- **IRFU (<- DAPNIA)**
  - Pierre Micout
  - GRIF (see LAL)
- **LAL**
  - Michel Jouvin
  - GRIF
- **NDGF**
  - Mathias Wadenstein
- **KIT**
  - Manfred Alef
  - GridKa
  - RTOI in 1 y 4 old PCs due 2 P/€
- **NERSC**
  - Jay Srinivasan
  - PDSF
  - Nagios with RSS feeds
- **AGLT2**
  - Shawn McKee
  - Atlas Great Lakes (Uni Michigan)

# Technical Topics (Storage I)



## ■ Infotrend (Ted Pang)

- Against Silent Data Corruption
- End-To-End Data Protection
  - ☹ ECC, RAID, CRC, ...
  - ☺ DIF (T10/T13)
    - Performancd
    - 8/512 Data
- Securiyty
  - New: IEEE P1619
  - Data at Rest
- ZFS
  - Good, but single host ...

## ■ DTS Inc. (Hiro Takahashi)

- L3/L4 Cache Layers

## ■ HEPiX Storage WG

- Andrei
- HEPCCC -> HEPiX
- May 2008 End of Comparison
- -> next Test more practical (FZK)

## ■ Lustre Roadmap Update

- Peter Bojanic (Director Lustre Group @ Sun)
- 100 GB/sec with 1000+ Nodes
- Under Development
  - pNFS exports
  - Solaris Server/Client
  - ZFS Support (end to end ...)
- Lustre 2.0 next year
  - Kerberos
  - Replication
  - Commit on share
- Lustre 3.0 later
  - OSS write cache
  - HSM + HPSS
  - Windows Native Client

# Technical Topics (Storage II)



## ■ CERN Storage Update

- Dirk Duellmann ROOTD
- CASTOR/XROOTD
- Consolidating
  - (more) POSIX File Access
  - (try) Storage is Filesystem
  - (some time) NFS 4.1

## ■ Lustre @ GSI

- Thomas Roth
- Success Story
- Get id of NFS
- Scary Remove OST (shut all)
- Stress + ls (reboot)

## ■ Optimizing Common Storage Systems

- Uni Michigan
- Read 1500 MB/s
- Write 680 MB/s

## ■ Finding a distributed Filesystem

- AFS,  
NFSv4.1, Lustre, Xrootd, dCache
- Tables Strength/Weakness
- Winner: dCache

## ■ KEK GPFS/HPSS

- Satomi Yamamoto (IBM?)
- Worldwide

## ■ KEK Data Grid System

- Yoshima Lida
- Gridftp <-> SRB

# Technical Topics (Data Centers)



## ■ (New) CERN Computer System

- Tony Cass
- Short Term
  - (>) 2.5 MWatt (+)
- Long Term: New CC
  - (>) 5 MWatt
  - < 30 % / y growth
  - 4 company design < 28.11.

## ■ Push/Pull Registry WEB-Service

- Farooq Ahmed
- 500 MUser on 3 Bpages
- Semantic WEB
  - RDF, OWL
  - Formal Description 4 Dynamic Usage

## ■ Thermal CC Simulation

- Enrico Mazzoni (INFN)
- Simple Modell ?

## ■ Green CC

- Giovanni Organtino (INFN + Uni Roma)
- Cooling Calculations

## ■ High Density Displays

- Sebastian Lopienski (CERN)
- State, flags, property from lemon
- Colors and Icons in squares

# Technical Topics (CASTOR)



## ■ Status and Plans

- Sebastien Ponce (CERN)
- Secure, Database Centric, Robust, Scalable
- T1:ASGC,RAL,CNAF
- Continuously improving ...

## ■ SRM2 Interface + Monitoring

- Guiseppe Lo Presti (CERN)
- SRM2: (||->=) Castor
- Mon: Prototype

## ■ Operations

- Miguel Coelho dos Santos (CERN)
- Diskserver / Experiment
- Mostly HA active/active
- ITIL courses done

## ■ Increasing Tape Efficiency

- Steven Murray (CERN)
- Repack (Final Tests)
- Small Files are slow (3M/F)
- NEW + AUL Tape Formats

# Technical Topics (Operating Systems and Applications I)



## ■ SVN Pilot

- Alan Silverman (CERN)
- Centralized OpenSource CVS/SVN
- 300 Projects
- 3000 Developers
- Replace CVS by SVN
  - Fast, ACL's, DB
  - Online Backup
  - Offline Working
  - SVN Online End 2008
  - CVS down End 2009
  - https/ssh
  - NFS  $\frac{3}{4}$ , AFS
  - Pre/Post-Script-„jails“

## ■ High Performance Cryptographic Computing

- Chen-Mou Cheng (Uni Taiwan)
- Secure + Fast/Parallel
- GPU Computing
- Cuda Compiler + Threads

## ■ Scientific Linux Update

- Troy Dawson (FNAL)
- 3.0.9 S-Fixes  $\leq$  2010
- 4.4 Most Popular
- 5.2 Add: xfs, yum-xxx,...
- 4.7: Firefox 3, SL3/4 in VM
- Fastbugs once /week
- 5.3 in 03/2009
- 4.8 in 06/2009
- 6.0 in 01/2010

# Technical Topics (Operating Systems and Applications II)



## ■ Scientific Linux Plenary Discussion

- Troy Dawson
- SL4 to 10.10.2010
- Hardwaredriver
- Java 1.6 End of the Year

## ■ Update on Mail Services at CERN

- Rafal Otto
- 19000 User, Exchange 2003
- 3,5 TB, 7000 Mail List
- 100 billion spam in internet
  - 90 % overall
  - 99 % at CERN
  - Very low Error Rate
- New Spam Filter
- Migration to Exchange 2007
  - Redesign (Pub Folders, ...)

## ■ Overview and recent performance work in the ARC CE

- Mattias WADENSTEIN (NDGF)
- Resource Manager ...
- gLite || arc
  - CMS Interoperability done

# Technical Topics (Virtualization I + II)



- **New possibilities and ideas around Hyper-V virtualisation and Virtual Desktop Infrastructure (VDI)**
  - Dr. Rafal OTTO (CERN)
  - Get your personal VM “from” WEB form
  - Vision on Distributing Thinclients 4 central Windows Support
- **Grid Infrastructure at SiNET and Chrooted Systems**
  - Dr. Andrej FILIPCIC (Research) .si
  - Grid Virtualizations in “changeroot” environments
- **Bridging Grid and Virtualization**
  - Dr. Fubo ZHANG (Platform Computing)
  - LSF Company
  - Integration of VMs in LSF 4 Grid Farms
- **Virtualization Experience with Hyper-V on Windows Server 2008 System**
  - Reinhard BALTRUSCH (DESY)
  - More on Hyper-V
- **Virtual Network Services**
  - Thomas FINNERN (DESY)
  - HA Services and more with F5 Load Balancer
  - Application Delivery Networking

# Technical Topics (Networking and Security)



- **Activity for construction of secure infrastructure in KEK**
  - Dr. Kiyoharu HASHIMOTO (KEK)
  - Secure Authentication Network
  - Smartcard, WLAN, VPN
- **Vulnerability Management by the Integration of Security Resources and Devices with DBPowder**
  - Tadashi MURAKAMI (KEK)
- **MCU (Multi-point Control Unit) services switch for the quality**
  - Dr. Teiji NAKAMURA (KEK)
  - Discuss Cross-Site Compatibility
- **Grid Security Update**
  - Dr. David KELSEY (STFC/ RAL)
  - Overview Middleware
  - mostly gLite
  - T1 (+T2) CERT Test
- **IPv4/IPv6 Transition Status and Recommendations**
  - Dr. Fred BAKER (CISCO)
  - Chinese (UNI) driven
  - Pilot jobs traceability
  - 2010 no ipv4 addr anymore
  - „World pictures“
  - Turn it on in your ipv4 net
- **Cyber Security Update**
  - Tony CASS (CERN)
  - Ssh keys attacks, DNS, ...
  - Incidents > tighter policies
  - More Attacker, different styles

# Technical Topics (Benchmarking)



- **CPU Benchmarking at GridKa – Update**
  - Manfred ALEF (KIT)
  - SPECint: 32 <= 32@64 <= 64
- **Report from the HEPiX Benchmarking Working Group**
  - Helge MEINHARD (CERN)
  - Specs old 1 core
  - -> SPEC\*2006
- **Power Efficiency of Servers**
  - Helge MEINHARD (CERN)
  - CERN Data Center 0.1 permille of World Consumption
  - 6.5 € Malus per VA
  - Factor 9 better in 4 y
  - 2nd CPU: No (500GB Disk Server)
  - E Chips are sometime L(ow Power)

# Technical Topics (Miscellaneous)



- **Site Monitoring and Site Commissioning for the LHC experiments**
  - Dr. Julia ANDREEVA (CERN)
  - Nagios with Grid Extensions
  - SAM
  - Backimport to sites (?)
  - Inhomogenous
- **Usage of the Request Tracker in HEP institutes**
  - Dr. Sven STERNBERGER (DESY)
  - 300.000 Tickets / 5 y
  - Grid Integration
- **Using cfengine for Grid Site Management**
  - Dr. Jan KUNDRAT (Institute of Physics of the ASCR)
  - Better than || ssh
  - Private svn rep
  - Testing, Testing, Testing, ...
  - Transparency of Config (grep ...)
- **Evaluation of GPFS Connectivity over High-performance Networks**
  - Jay SRINIVASAN
  - Start: High performance, low latency for e.g. MPI internode comm -> HPFS
  - Linux kernel tuning 4 10GE or IB
- **Grid Interoperations**
  - Laurence FIELD
  - Move real to virt orgs
  - No common interface to org entry
  - + mon, account, env, policies, ...
  - Pilot VO
  - GLUE 2.0, BES, JSDL
  - GIN (Grid Interops Now)

# Wrap UP



- Ordinary Meeting + non-HEP
- Contact to Asia (growing)
- Chuck Boeheim to Cornell (will join HEPiX ?)
- Next: Umea (Sweden) 25-29 or 11-15 May
- Continue tracks except benchmarking
- Next Fall Berkeley ?
- HEPiX tracks:
  - General
  - Mailing
  - Virtualization emphasis without WG
- Sandy new US co-chair

# Thank you for listening (謝謝)



Questions (問題)?

Answers (答覆)!