High End Computing Activities in the Indian Physics Community

Outline

- Computing issues in Theoretical Physics
- Facilities for HEP Theories
- Need for HEP experiments
- Facilities for HEP Experiments
- Some work at national level
- Summary

January 13, 2011

Sunanda Banerjee

Problems in Theoretical Physics

- Computation in the field of Quantum Field theory is handled using path integral method in Euclidian space time without assuming smallness of coupling constant
- The calculations are done on a large lattice – typically $32^3 x 64$. This corresponds to 10^8 gauge fields.
- Estimation of an expectation value involves integration over these large dimensions (including all degrees of freedom) which is not possible with present technology – one needs several importance sampling techniques to get the results.
- This involves inversion of huge matrices – typically of the size $10^9 x 10^9$. This takes ~85% of the CPU usage of a Lattice Gauge Field computation
- Many of the algorithms are highly parallelizable. Can utilize high performance computing facility and highly parallelized codes

→ Several HPC clusters setup in different parts of the country

HPC Facilities in India

- Facility at Saha Institute, Kolkata:
  - Since 2005 Cray XD1 with 60 nodes
    - 2 single-core AMD opteron 2.2 GHz processor, 4 GB RAM
    - 4 TB (5.2 TB) usable (raw) disk space
    - gcc 3.3, pgi v6 compilers
    - Peak performance 528 Gflops
    - Typical efficiency for application jobs ~20%
  - Since 2009 Cray XT5 with 172 nodes
    - 2 quad-core AMD opteron 2.4 GHz (2.6 GHz) processor for compute (service) node
    - 36 TB (48 TB) usable (raw) disk space
    - Cray Seastar 2.0 interconnect
    - gcc 4.3 compiler
    - Peak performance 13.8 Tflops
    - Typical efficiency for application jobs ~30%

HPC Facilities in India

- Facility at Tata Institute of Fundamental Research, Mumbai
  - Since 2004, Cray X1 with theoretical peak performance of 200 Gflops
  - Since 2010, a rack of Blue Gene P with theoretical peak performance of 14 Tflops
- Facility at Institute of Mathematical Sciences, Chennai
  - Since 2004, a Linux cluster with Infiniband interconnect with a theoretical speed of 1.3 Tflops
- At the Indian Association of Cultivation of Science, Calcutta
  - Experimenting with a GPU cluster

Indians in HEP Experiments

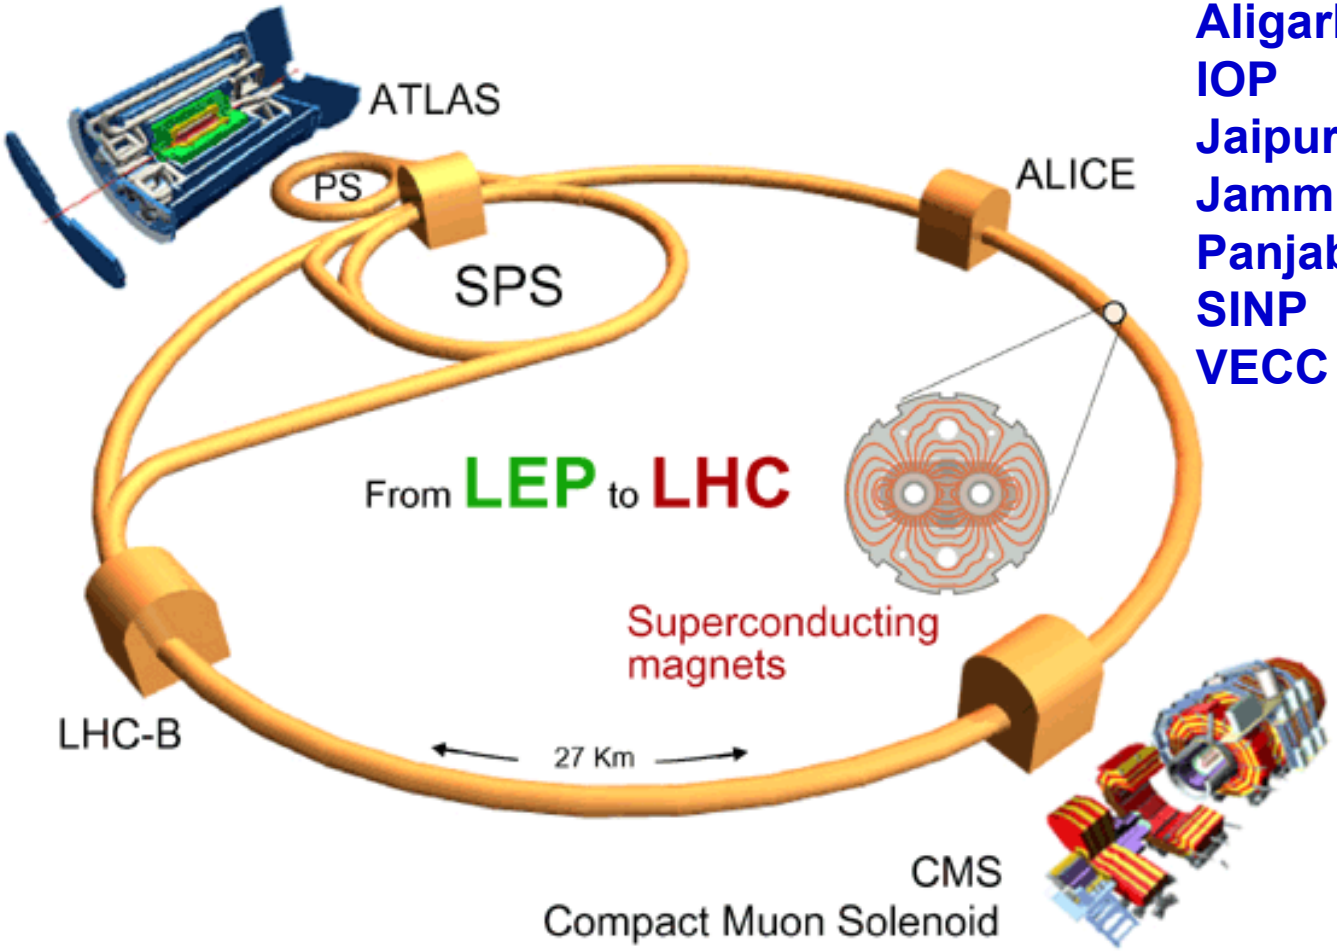

Aligarh
IOP
Jaipur
Jammu
Panjab
SINP
VECC

BARC
Delhi Univ.
Panjab Univ.
TIFR (EHEP)
TIFR (HECR)
Viswabharati

Experimental Challenge

LHC Detectors are radically different from the ones from the previous generations

High Interaction Rate

pp interaction rate **~1 billion interactions/s**

Large Particle Multiplicity

~ <20> superposed events/crossing

~ 1000 tracks stream in every 25 ns

Need highly granular detectors with good time resolution for low occupancy

⇨ **large number of channels (~ 100 M ch)**

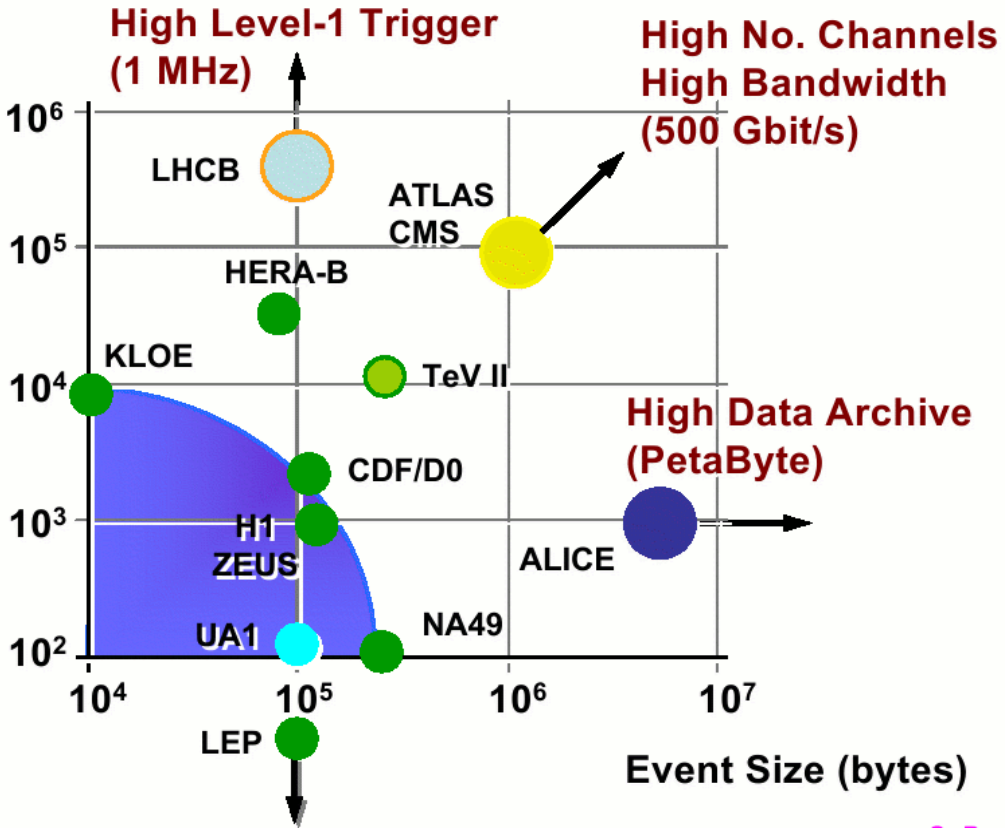

Experimental Challenge (II)

Typical requirements for a LHC experiment:

Event o/p rate	10^9 events/year
Data to tape	few petabytes/year
Life time of experiment	2-3 decades
Processing capacity	> 10^{13} instruction/s
Network bandwidth	several Gbps
Users	> 2000

Distribution of Users and developers

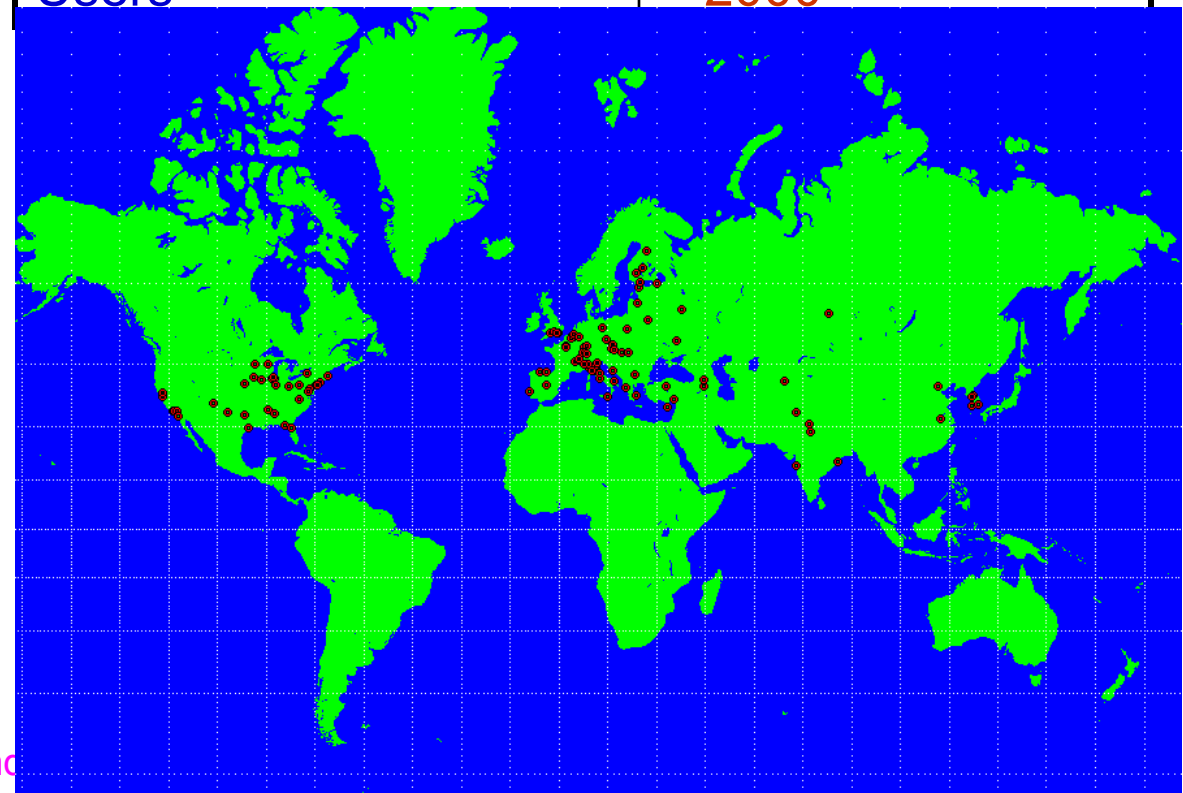

Computing Strategy

- Several level of reconstructed results
  - Results from full reconstruction (RECO)
  - Event summary data (ESD)
  - Analysis object data (AOD)
  - Event tag data (TAG)
- In addition to reconstruction of real data, there will be similar or larger number of simulated events

	Rate (Hz)	RAW (MB)	ESD/RECO (MB)	AOD (kB)	MC (MB)	MC (% of real)
ALICE (HI)	100	12.5	2.5	250	300	100
ALICE (pp)	100	1.0	0.04	4	0.4	100
ATLAS	200	1.6	0.5	100	2	20
CMS	150	1.5	0.25	50	2	100

LHC Grid Framework

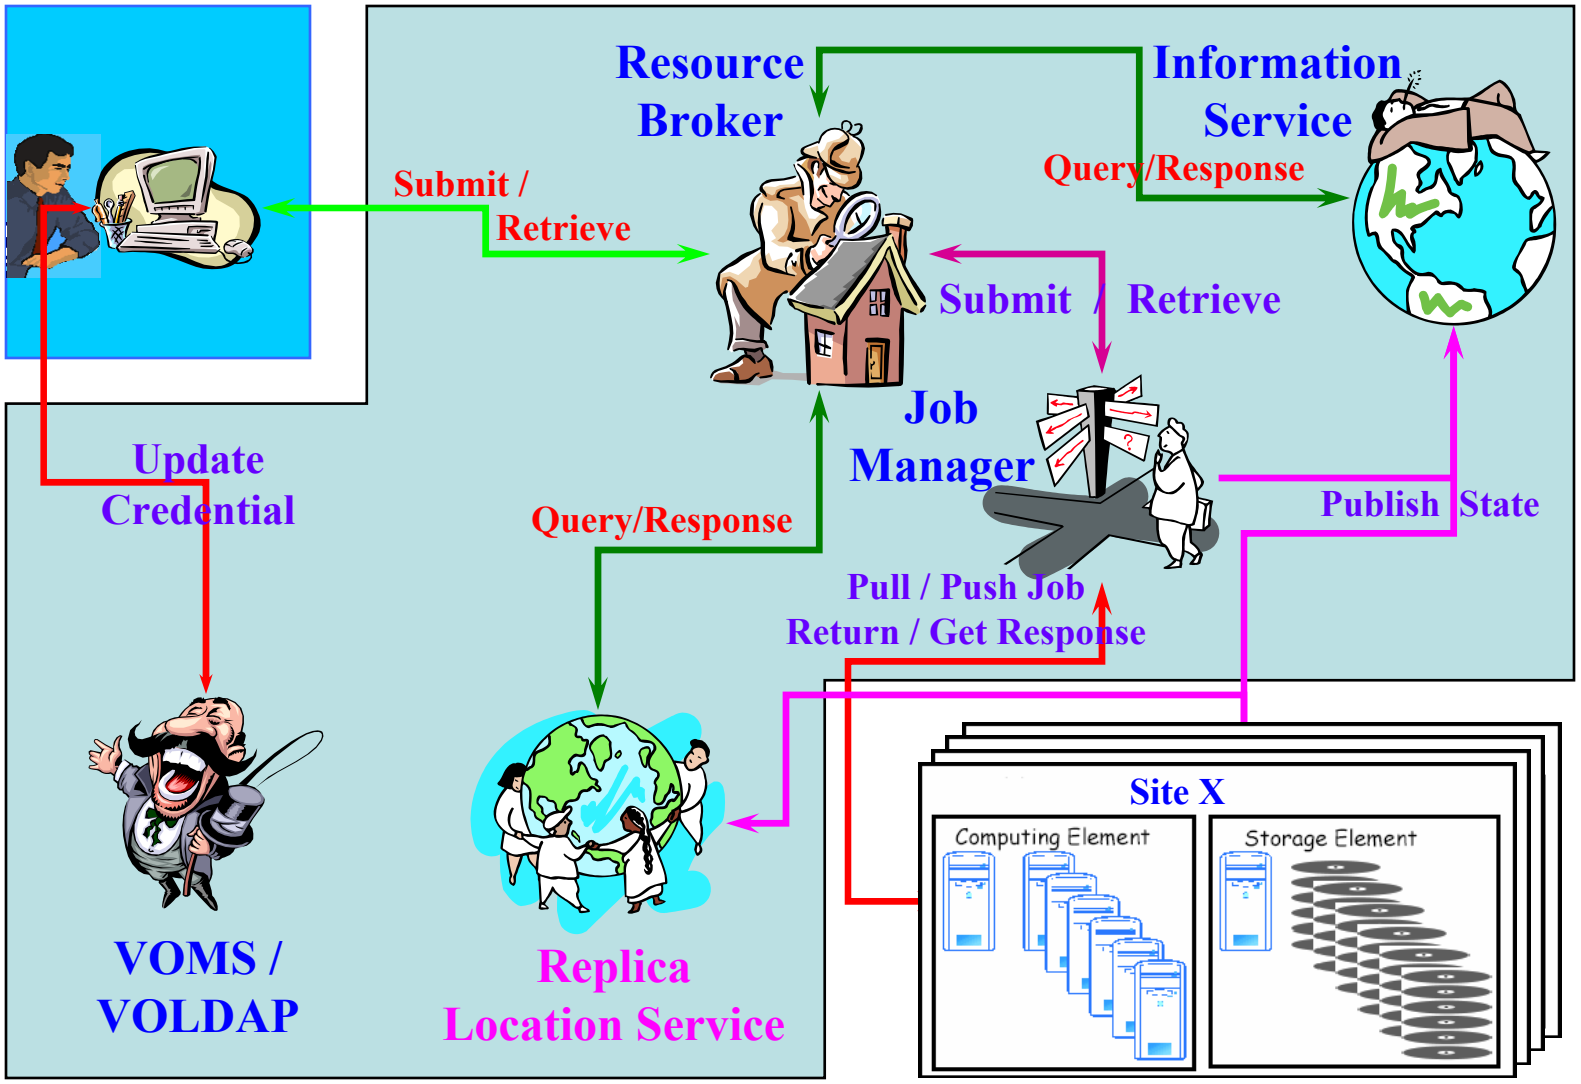

India participated in Grid Middleware development and in deployment

ALICE Computing Model

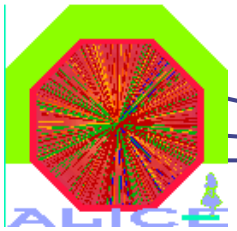

Online System

~1.5 Gb/s

Online Farm

~1.25 Gb/s

Tier 0

CERN Computer Center

622/2500 Mb/s

Tier 1

France Regional Center

Germany Regional Center

Italy Regional Center

UK Regional Center

155/622 Mb/s

Tier 2

Kolkata

Tier3

Institute

Physics data cache

Start with a strict hierarchical structure

CMS Computing Readiness

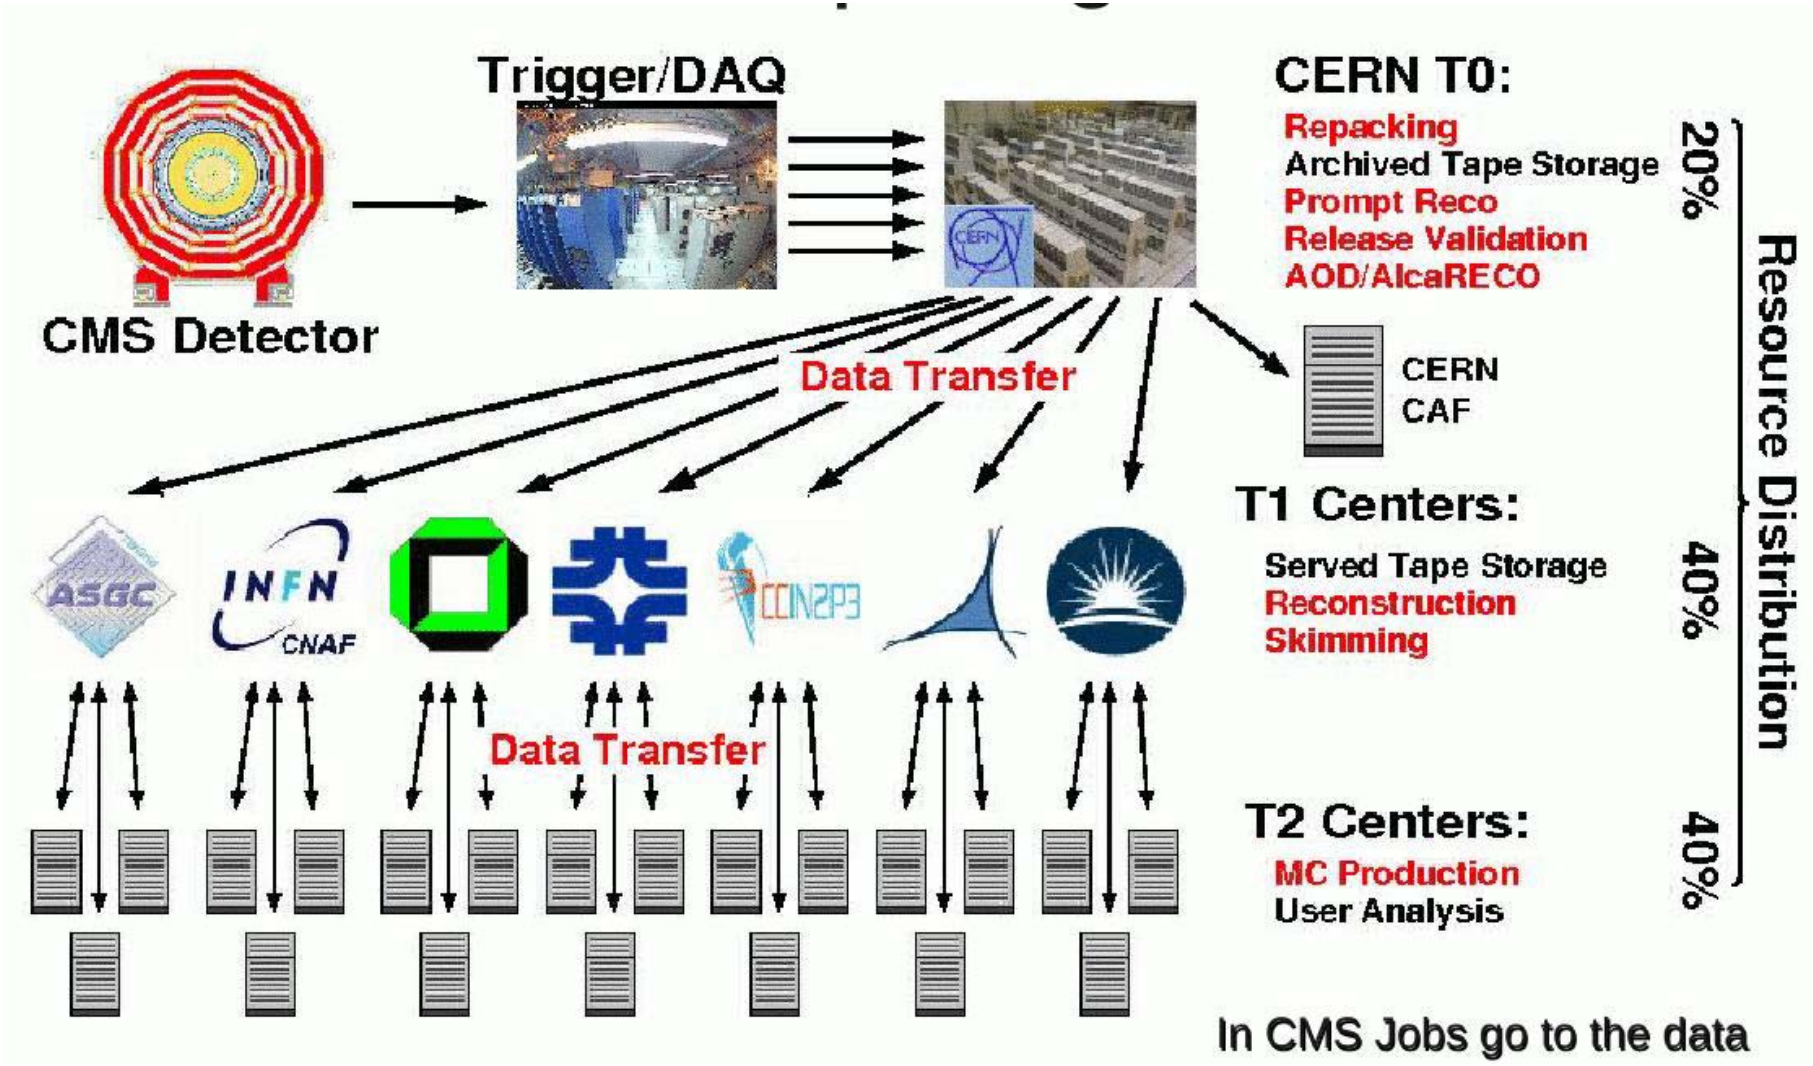

Network Connections in India

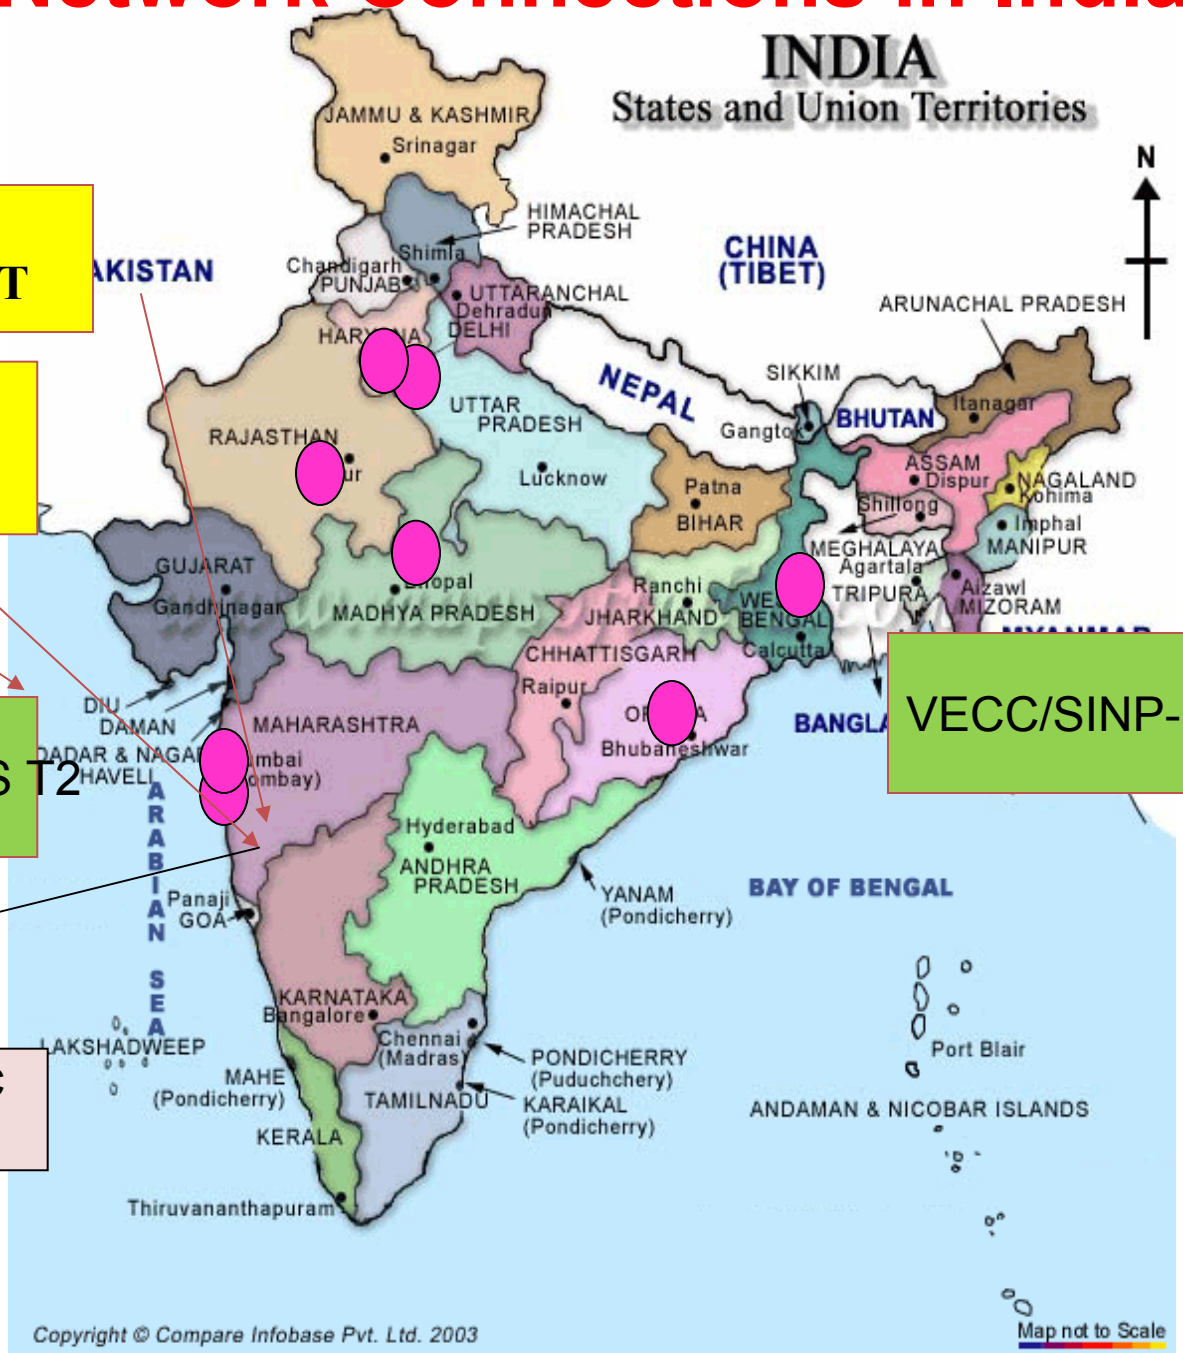

1 Gbps to CERN peered to GEANT

175 Mbps Geant Internet → 2.5 Gbps TEIN3

1 Gbps GLORIAD

TIFR-INDIACMS T2

VECC/SINP-INIDAALICE T2

100 Mbps to VECC RRCAT, IPR

ALICE Grid Sites On MonaLisa

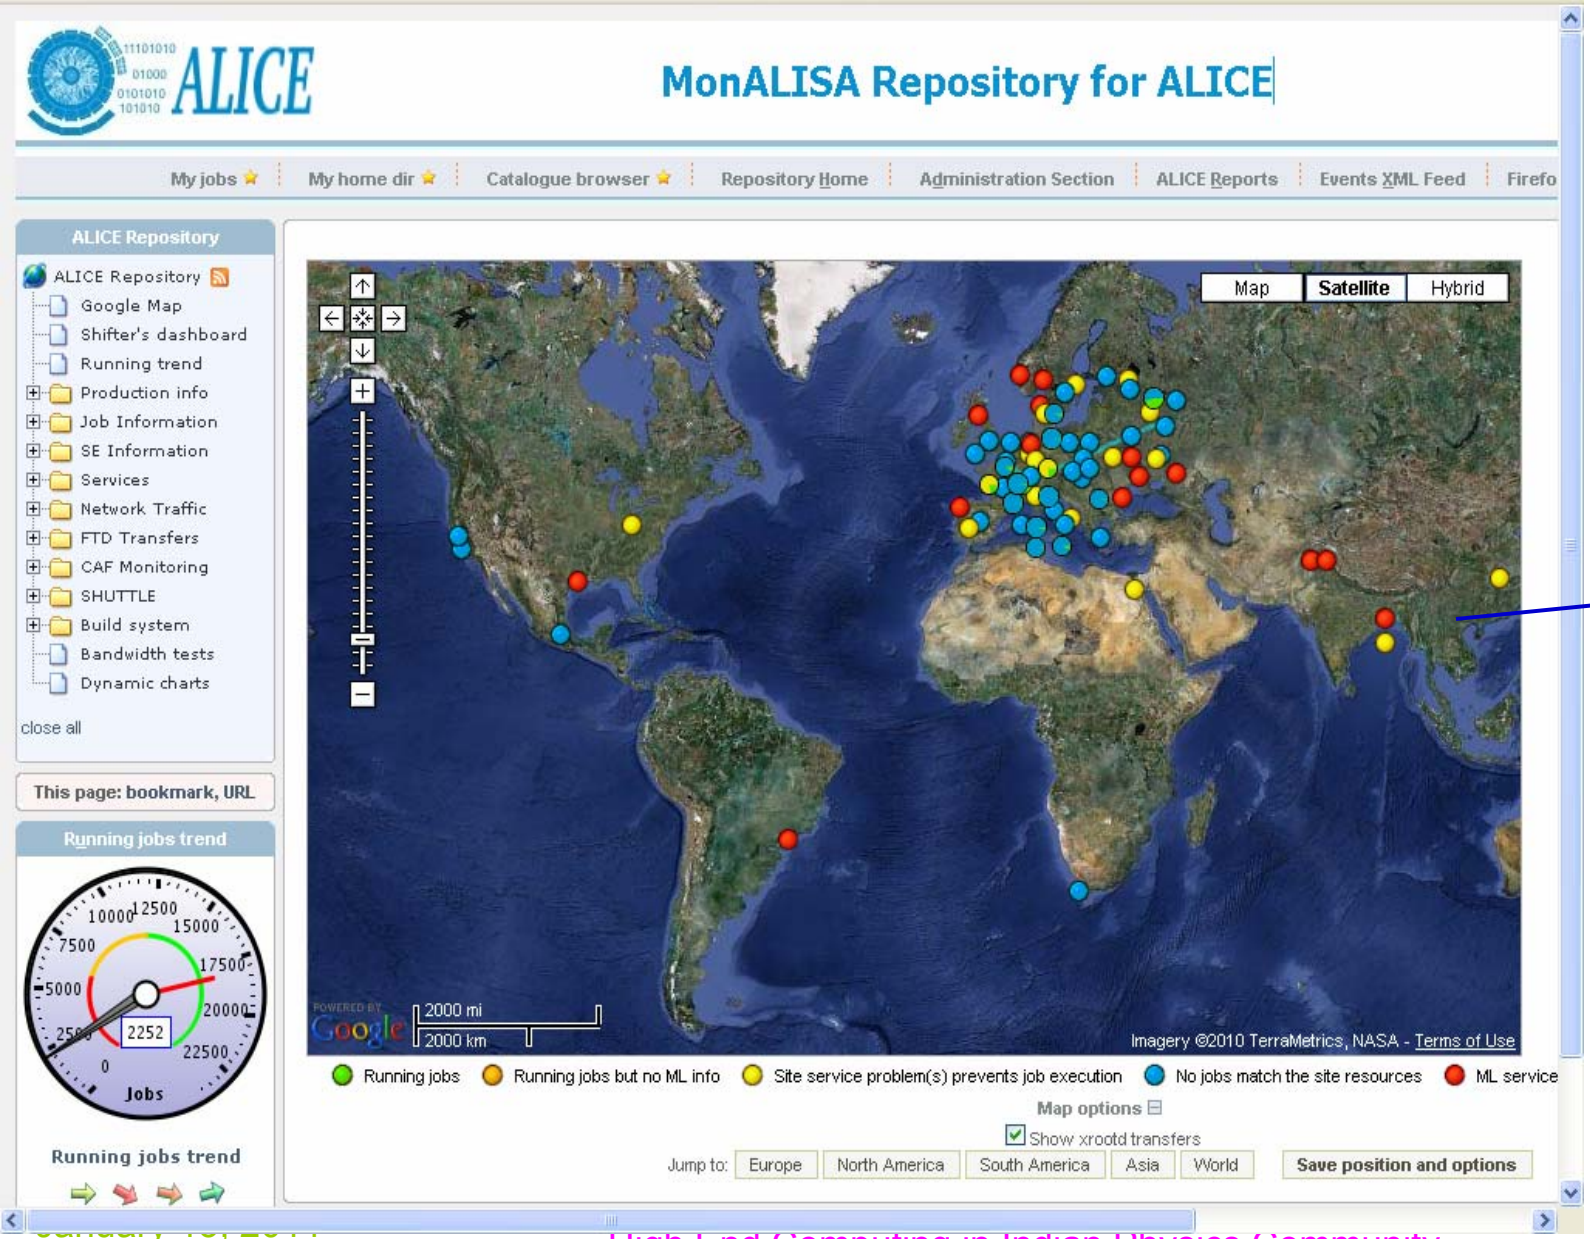

Only 64 bit Site named Kolkata-CREAM is UP and 32 bit site named KOLKATA is deprecated. (As per ALICE requirement.)

Hardware procured (VECC Tier-2)

1U Sliding LCD Monitor with 16 port KVM

- Dell(TM) PowerEdge(TM) M1000e Blade Server Chassis.
- 16 Number of Dell(TM) PowerEdge(TM) M610 High Performance Intel Blade
- Each blade has latest Nehalem based 2 * Intel Quad Core E5530 Xeon 2.4 GHz CPU with 8 MB cache.
- Each blade has16 GB RAM.
- Each blade has 2 * 146 GB Mounted as RAID1.
- Installed SLC 5.2 x86_64 OS

- AS-MD3000I - Dell(TM) PowerVault(TM) MD3000i
- 15 1TB SAS hard disks.
- 11.88TB Usable space after RAID5 and Hot Spare.

Availability & Reliability (VECC Tier 2)

Month	Site Name	Physical CPU	Logical CPU	Availability	Reliability
November-10	IN-DAE-VECC-02	94	220	77%	77%
October-10	IN-DAE-VECC-02	94	220	99%	86%
September-10	IN-DAE-VECC-02	94	220	98%	98%
August-10	IN-DAE-VECC-02	94	220	99%	99%
July-10	IN-DAE-VECC-02	94	220	100%	100%
June-10	IN-DAE-VECC-02	94	220	100%	100%
May-10	IN-DAE-VECC-02	78	156	93%	93%

Availability increased with the addition of 8 * 2 * Intel Quad Core E5520 Xeon 2.13GHz CPU with 8MB cache and 16GB RAM (64More jobs can run.)

ALICE Jobs completed (Last 6 months)

More than 40000 Jobs Completed

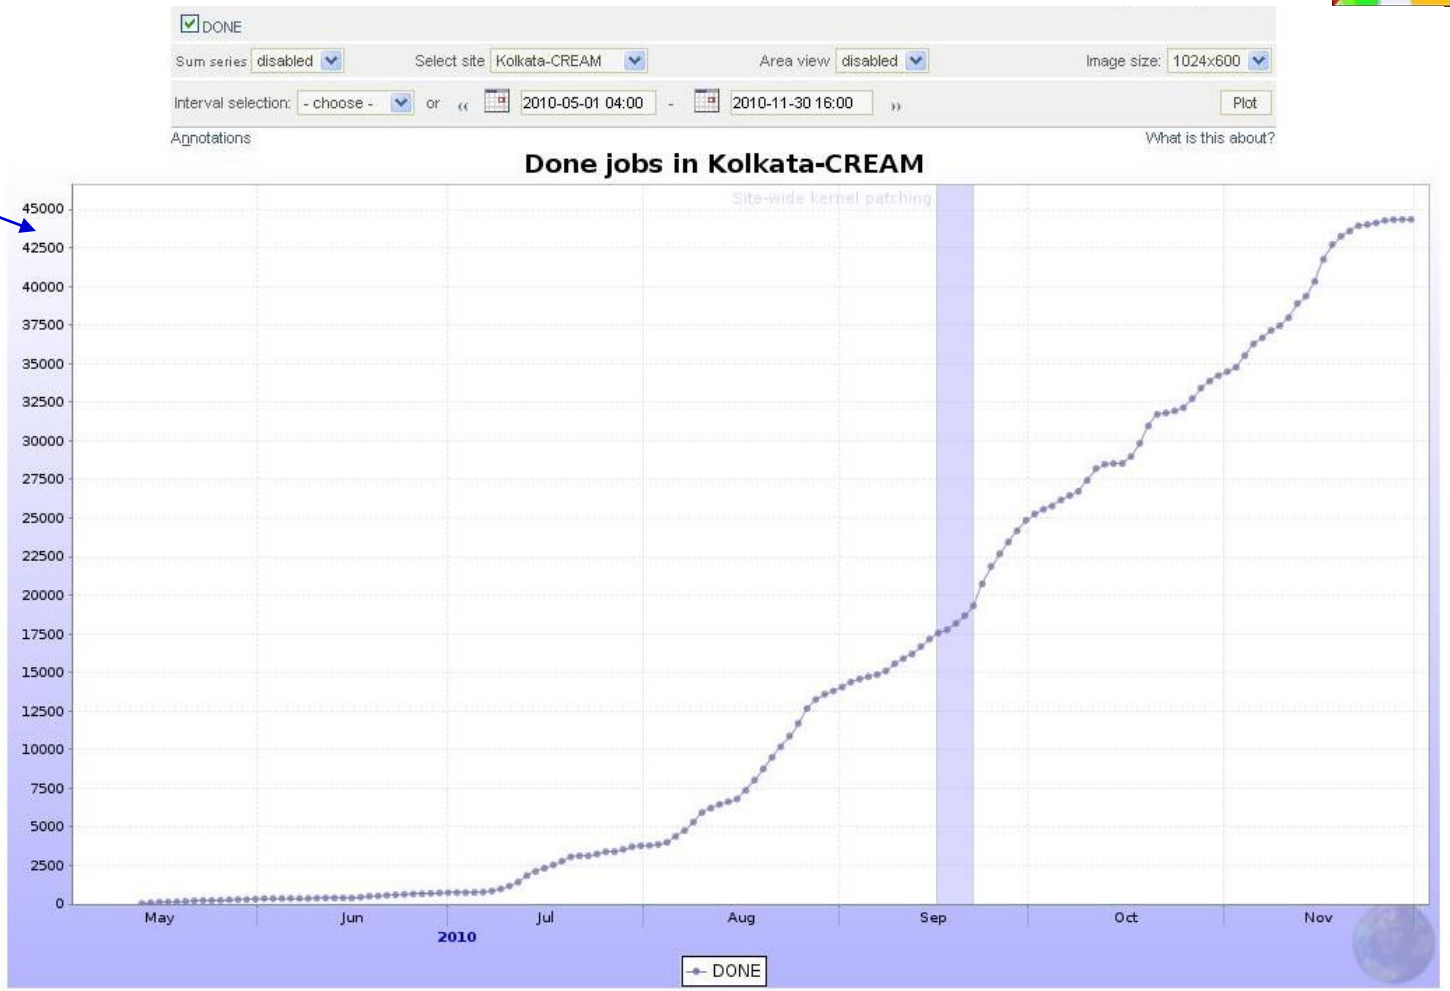

Statistics

Done jobs in Kolkata-CREAM

	Series	Last value	Min	Avg	Max
1.	DONE	44331	39	14911	44331
Total		44331		14911	

Kolkata-SE study for last 6 Months

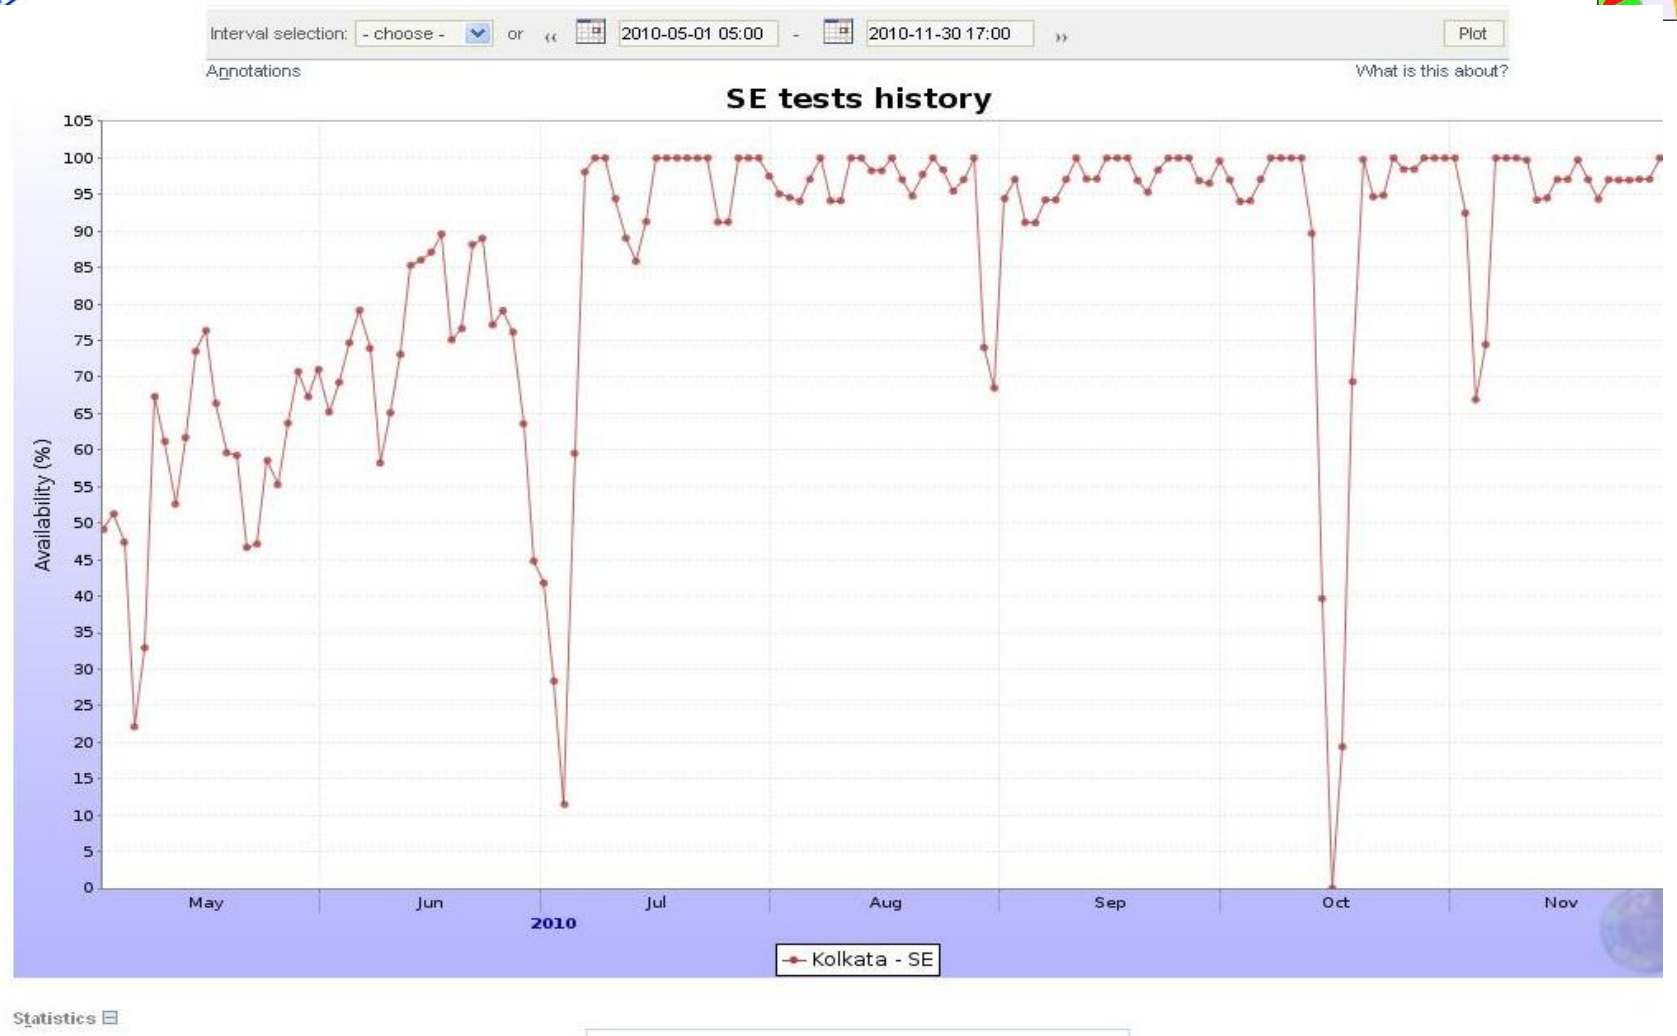

Statistics

SE tests history

	Series	Last value	Min	Avg	Max
1.	Kolkata - SE	100	0	84.57	100
	Total	**100**		**84.57**	

Tier 2 Site at TIFR

- Site resources for computing
  - 360 core CPUs
  - 48 BladeServers (Quad Core Dual Processors). Total HEPSpec2006 is ~2976
  - Each server:
    - Intel 5355@2.66 GHz (Clovertown)
    - 2 GB RAM per core (16 GB per server)
    - 72 GB, 10K SAS HDD
- Site resources storage:
  - One Storage element with DPM+SRM (18 DPM Disk Pool Nodes totalling 355 TB)
  - Write/Read: 500GB/s, 800GB/s. GsiFTP for WAN, RFIO for LAN access
  - Each Disk Pool Node:
    - Intel 5430 (Dual, Quad Core), 16 GB RAM
    - 4x1GB NIC 24x1TB HDD(7.2K SATA), RAID6

Tier 2 Site at TIFR

- Site resources for networking:
  - 1 Gbps LAN Fabric: Multiple 1 Gb/s connections between disk servers and Worker Nodes for better downloads
  - 1 Gbps WAN to 7 CMS Tier1's
  - 1 Gbps point to point connection to T1_CERN
  - 1 Gbps to other T1s by GEANT peering
  - 175 Mb/s internet via GEANT
- User Access to Tier 2
  - High end server as UI (User Interface)
  - Directly connected to Tier 2 LAN
  - Fast access to storage using RFIO
  - 20 TB local disk space to users with individual directories
  - AFS client installed
  - Dedicated job slots available

Data Transfer Quality

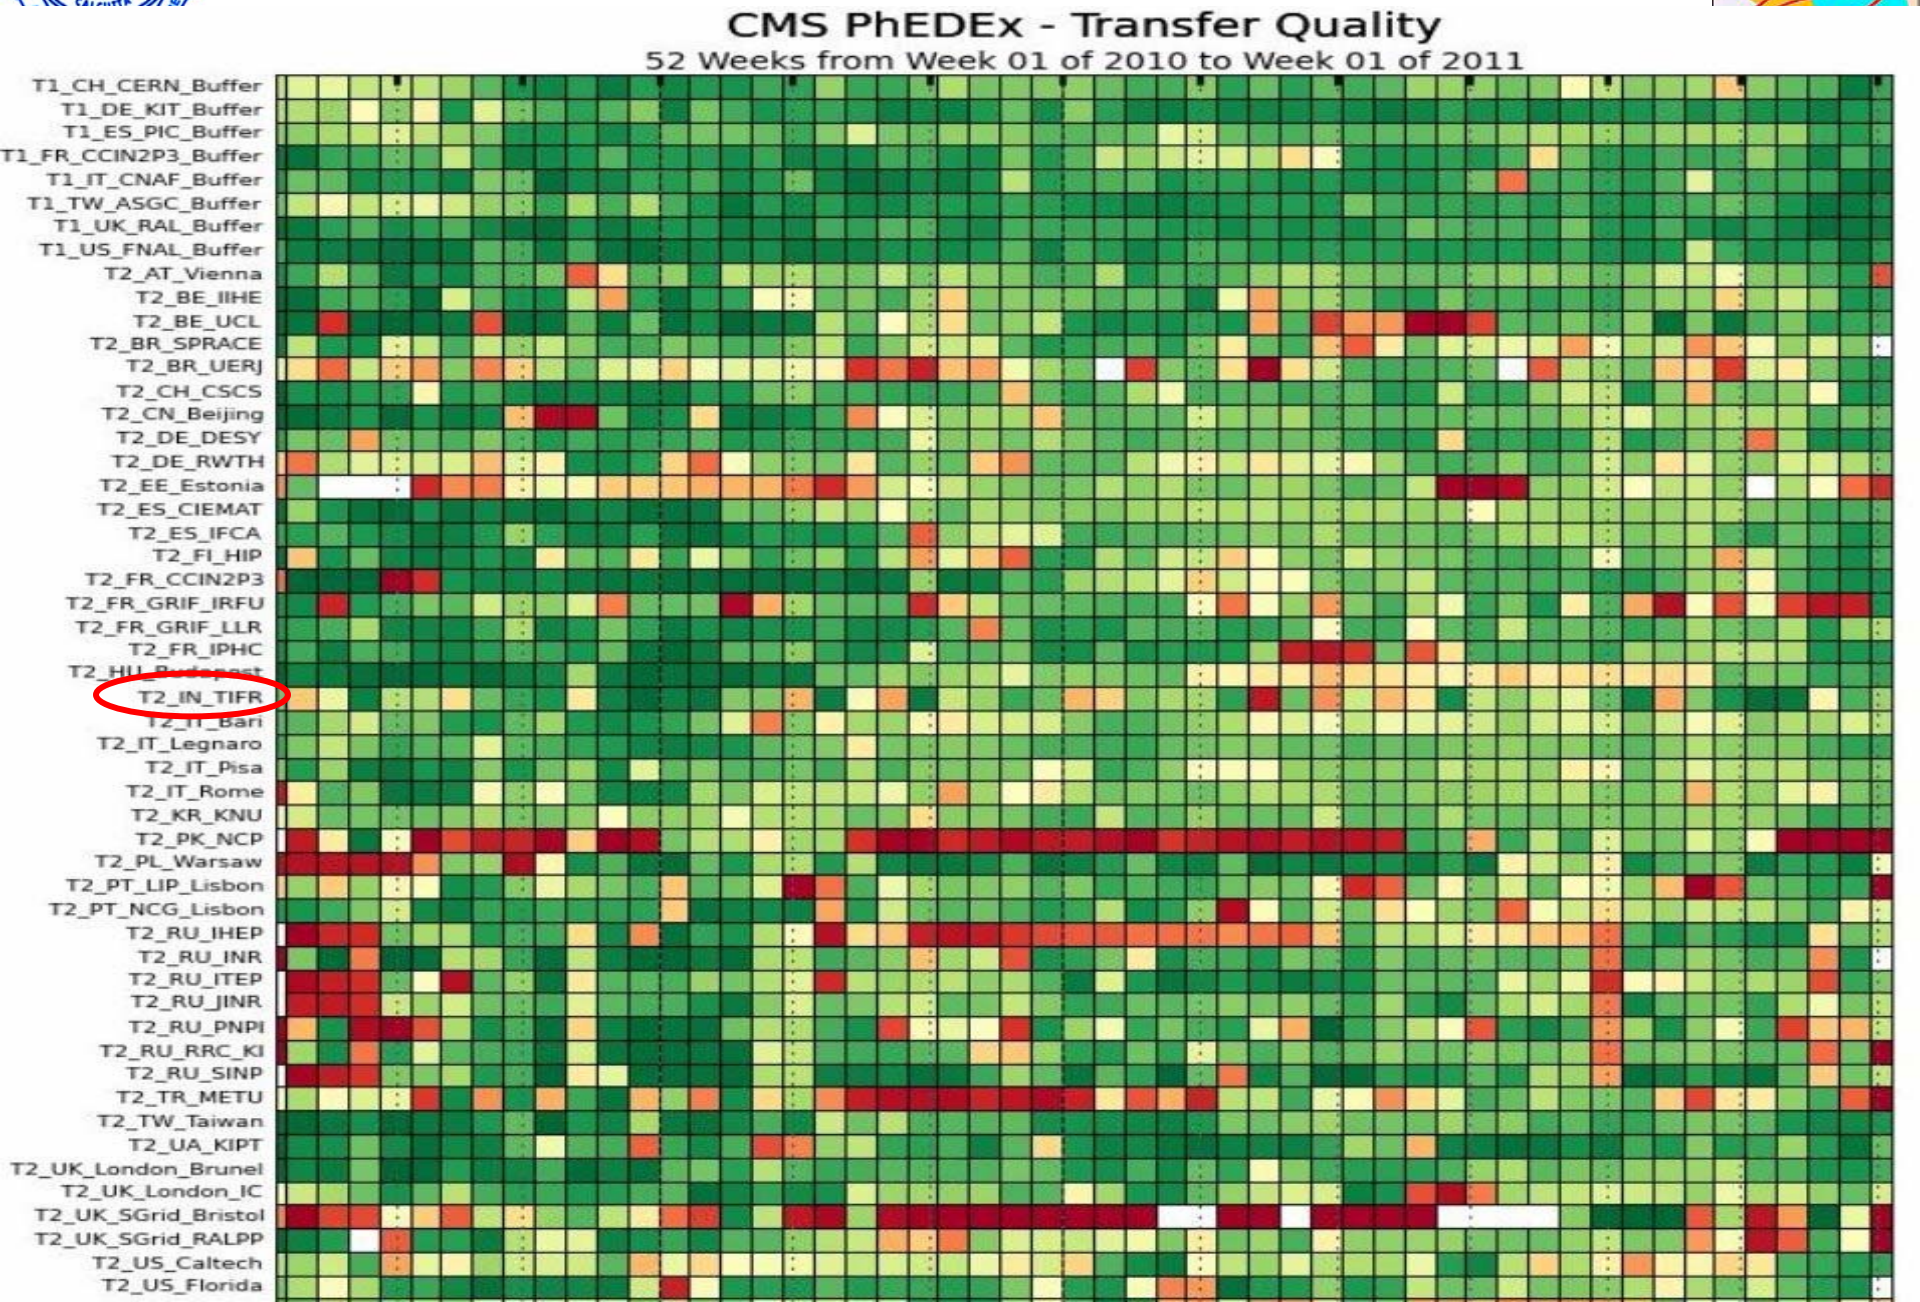

CMS PhEDEx Transfer

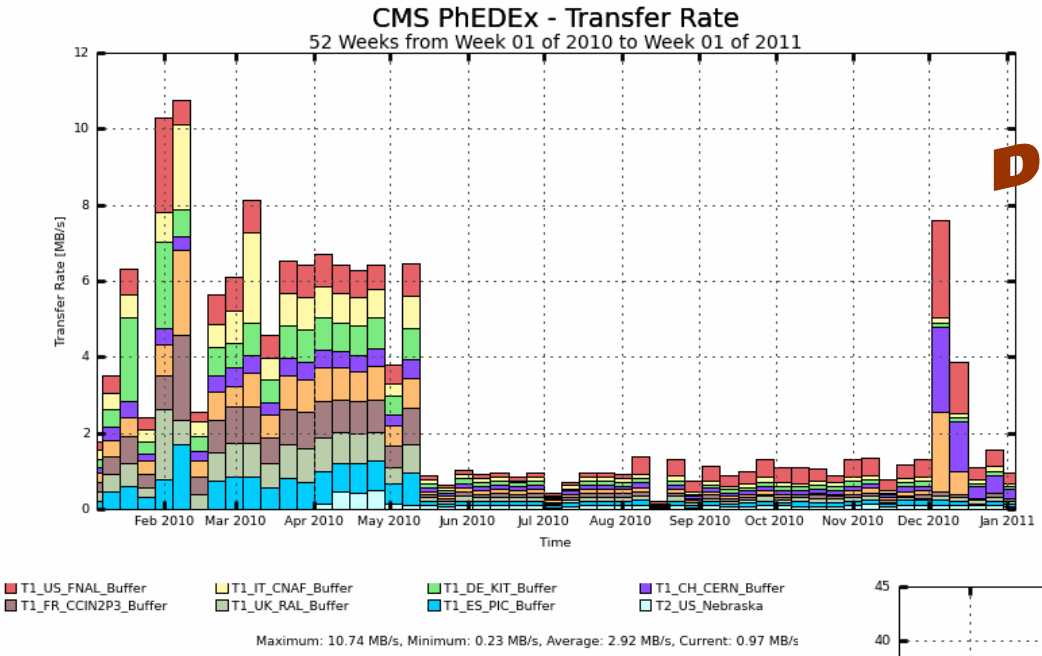

Download Debug

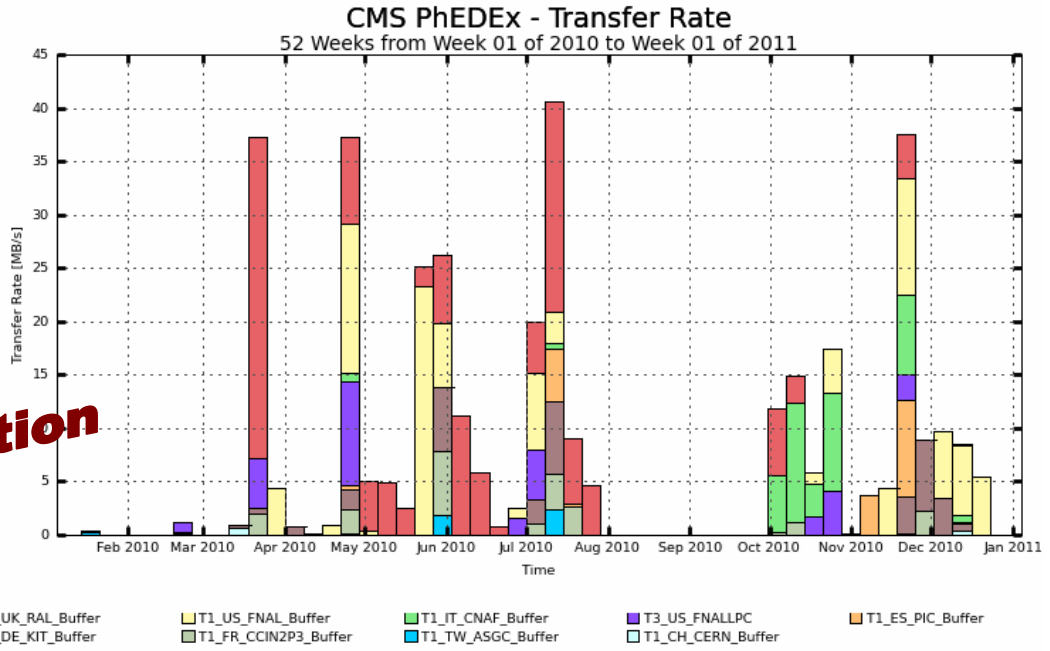

Download Production

CMS PhEDEx Transfer

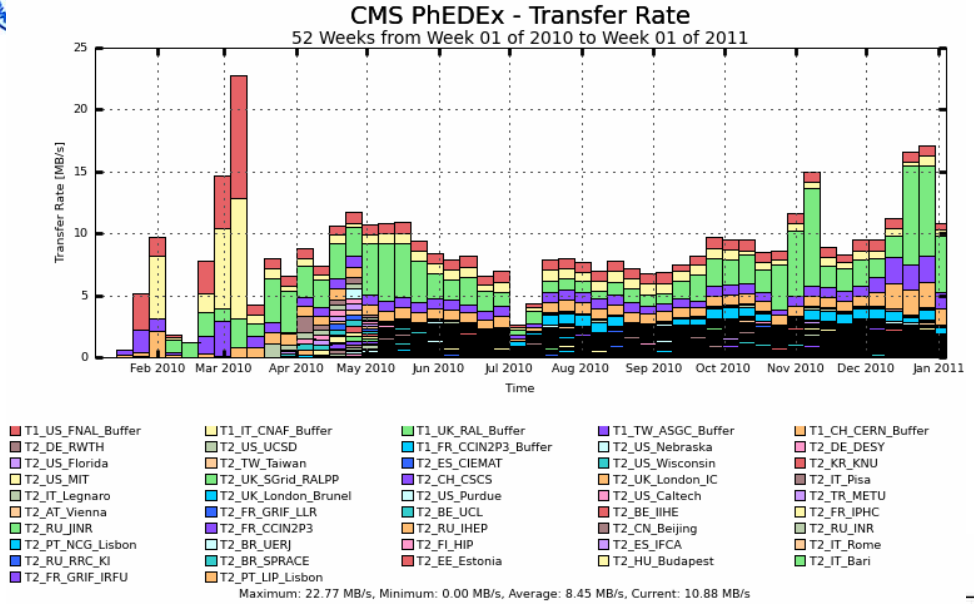

Upload Debug

Upload Production

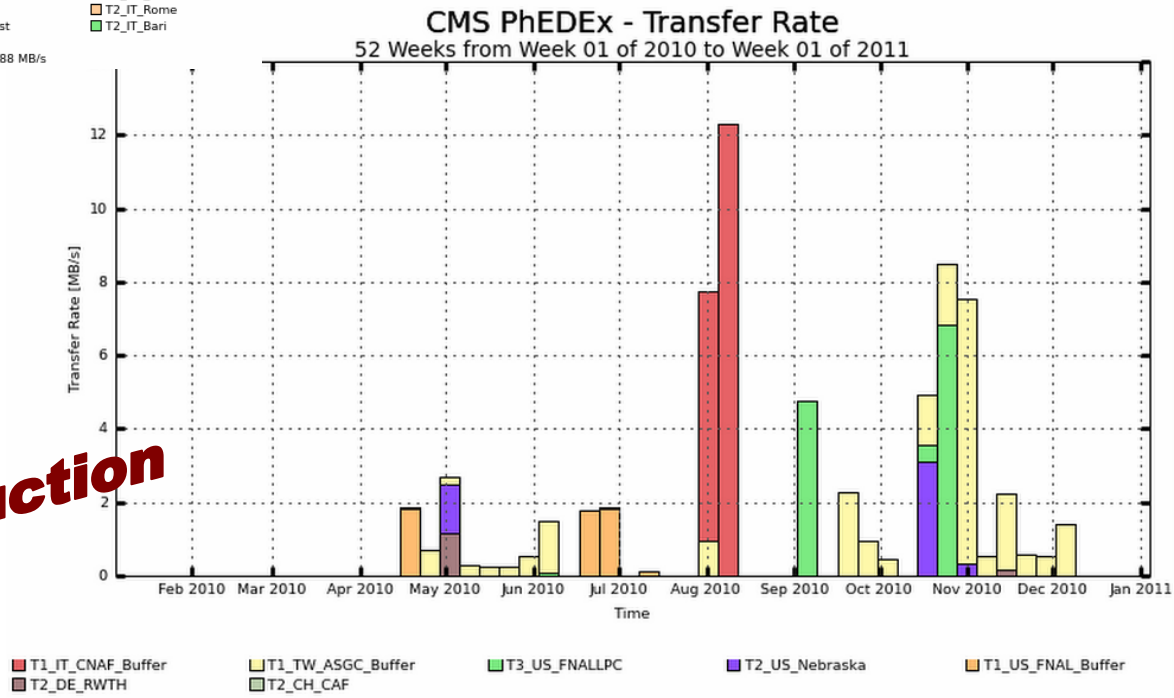

High Enc

Online Storage Utilization

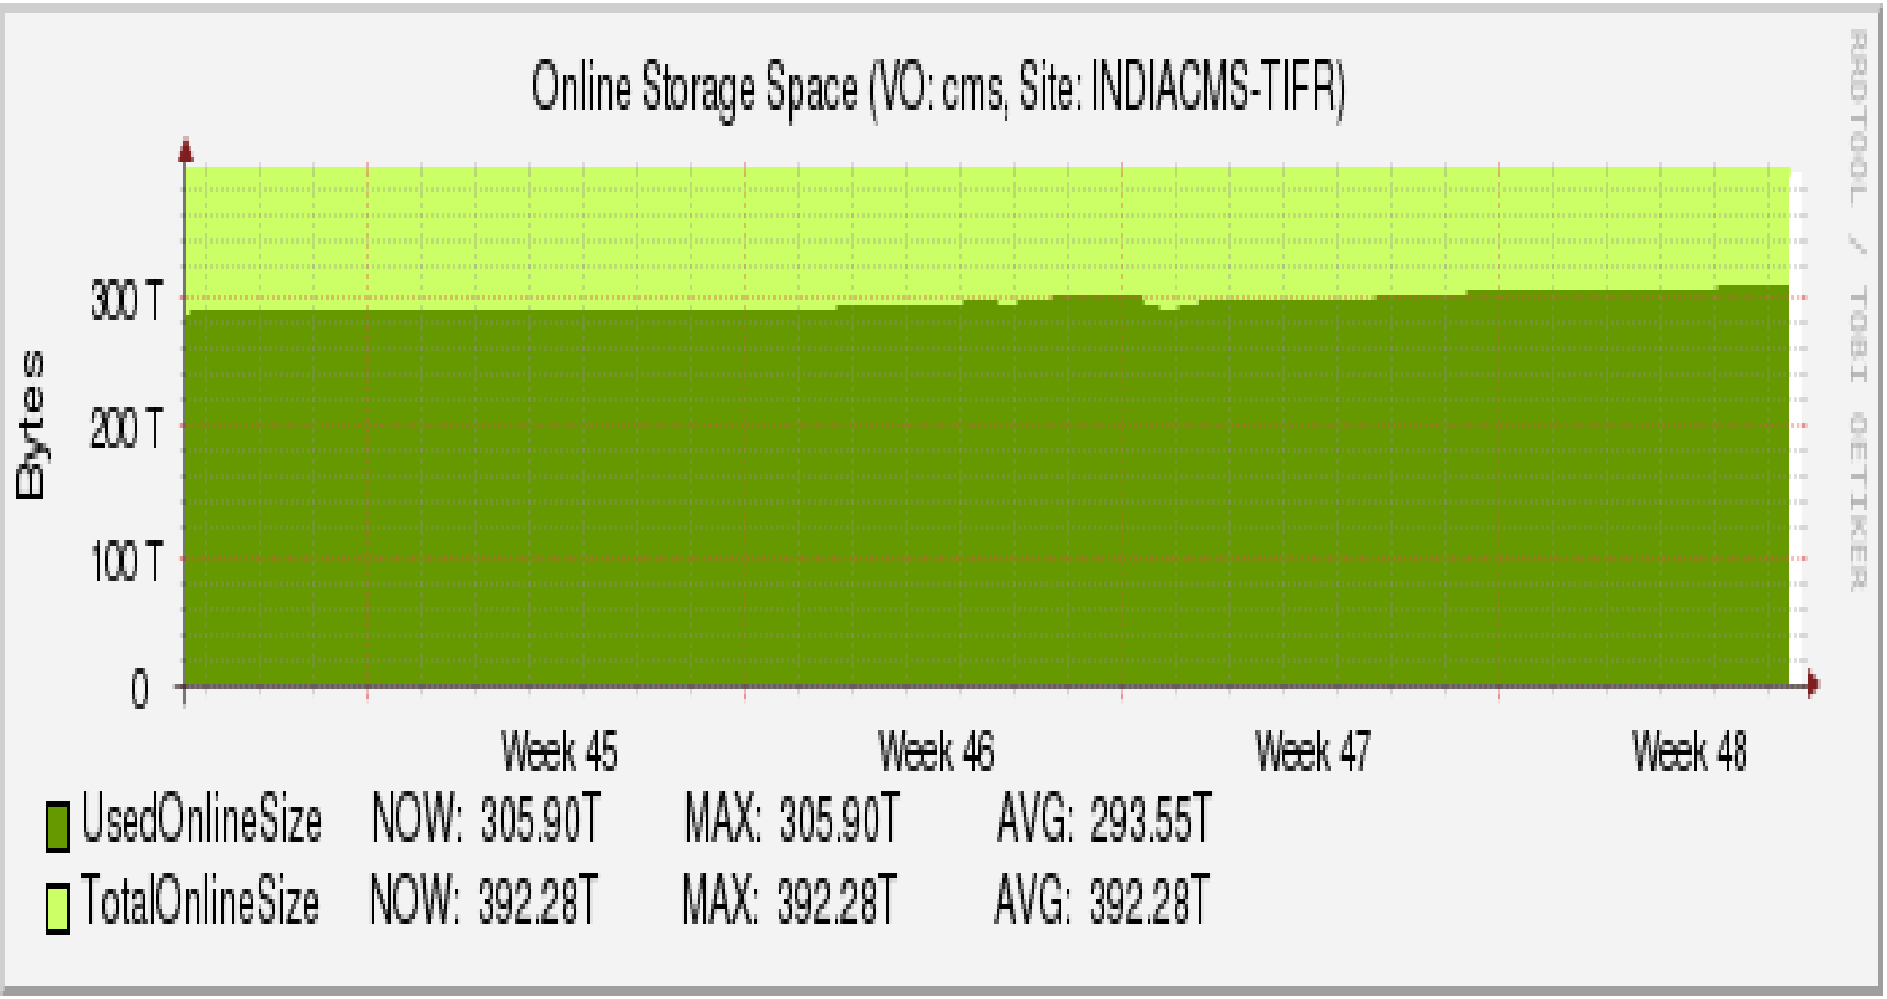

Performance of TIFR Tier-2

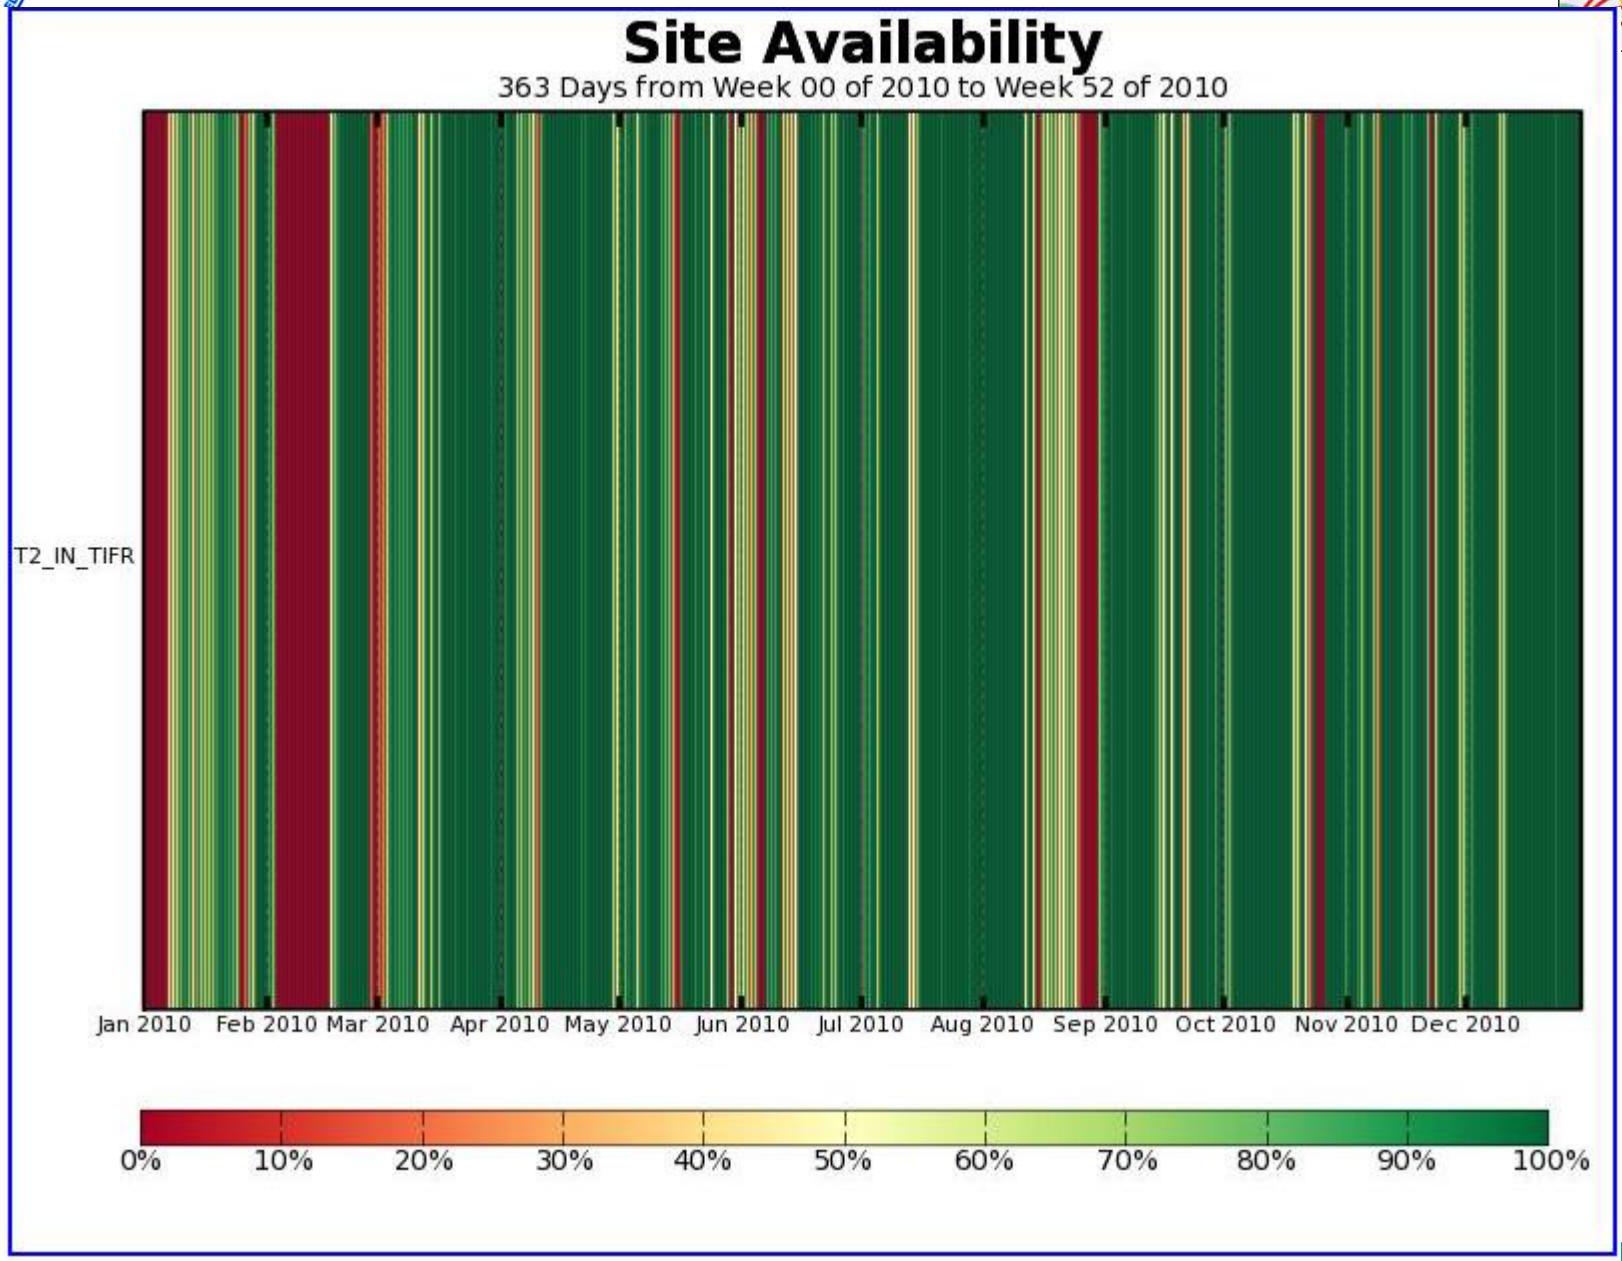

Tier3 at VECC

- Total 25 Nodes for VECC users and PMD Collaborators.
- 12 - 32bit nodes
- 13 - 64bit (2* Quad Core Nehalem with 16GB RAM) computing nodes
- 8TB of Usable storage
- 50 + active users (across India)
- 20 + active users (in VECC)
- Tier-3 is now extensively used
- Analysis jobs are running on 24*7 by different users on 64bit machines
- Jobs for all major results presented from PMD have been executed here

Tier3 at Panjab University

- Group part of the CMS Collaboration
  - 4 nodes of SUNFire x4140 (8 cores) based cluster (CentOS 5.3)
  - 32 GB RAM
  - 1 TB of Storage on the Head Node
  - 300 GB working space on each Compute Node
  - Connected via private 1 Gbps LAN
  - Head Node is on public IP + Lab LAN
  - 15 Desktop PC's as access/user nodes
- Group part of the ALICE Collaboration
  - Two Clusters
    - SUNFire x2200 (Scientific Linux CERN 4) : 4 Nodes, 32 Cores total, 32GB RAM
    - Dell PowerEdge 2950 (SLC 5) : 6 Nodes, 48 Cores total, 48GB RAM
  - DELL Power Vault NF600: 10TB + 10TB storage
- Outside connectivity
  - Via ERNET-Airtel 10 Mbps VPN
  - Via NKN – 1 Gbps through University Network (via proxy)

Departmental Facilities

- Centre at Delhi University Physics Department:
  - 32 nodes of Intel Xeon quad core dual processor of 3 GHz E5450 with 12 MB L2 cache 1333 MHz FSB
  - 32 GB RAM
  - 4 Gigabit ethernet card
  - 40 Gbps Infiniband switch for interconnect
  - Raid6 disks with 20 TB raw capacity
  - 20 Mbps shared line University wide – now upgraded to 32 Mbps with additional 100 Mbps like from knowledge network
- Centre at Panjab University Physics Department:
  - 3 SUN intel Xeon (Nehalam) 16 core nodes
  - 24 GB RAM per node
  - Two 1 Gbps connection using NKN
  - Peak performance ~400 Gflop

HEP groups at Delhi/Chandigarh benefit from these computing facilities

Indian GRID Project

- GARUDA is a collaboration of science researchers and experimenters on a nation wide grid that aims to provide the technological advances required to enable data and compute intensive science for the 21st century. One of GARUDA's most important challenges is to strike the right balance between research and the daunting task of deploying it.

- The Department of Information Technology (DIT), Government of India has funded the Centre for Development of Advanced Computing (C-DAC) to deploy the nation-wide computational grid 'GARUDA' which will connect 17 cities across the country in its Proof of Concept phase.

GARUDA

GARUDA Project

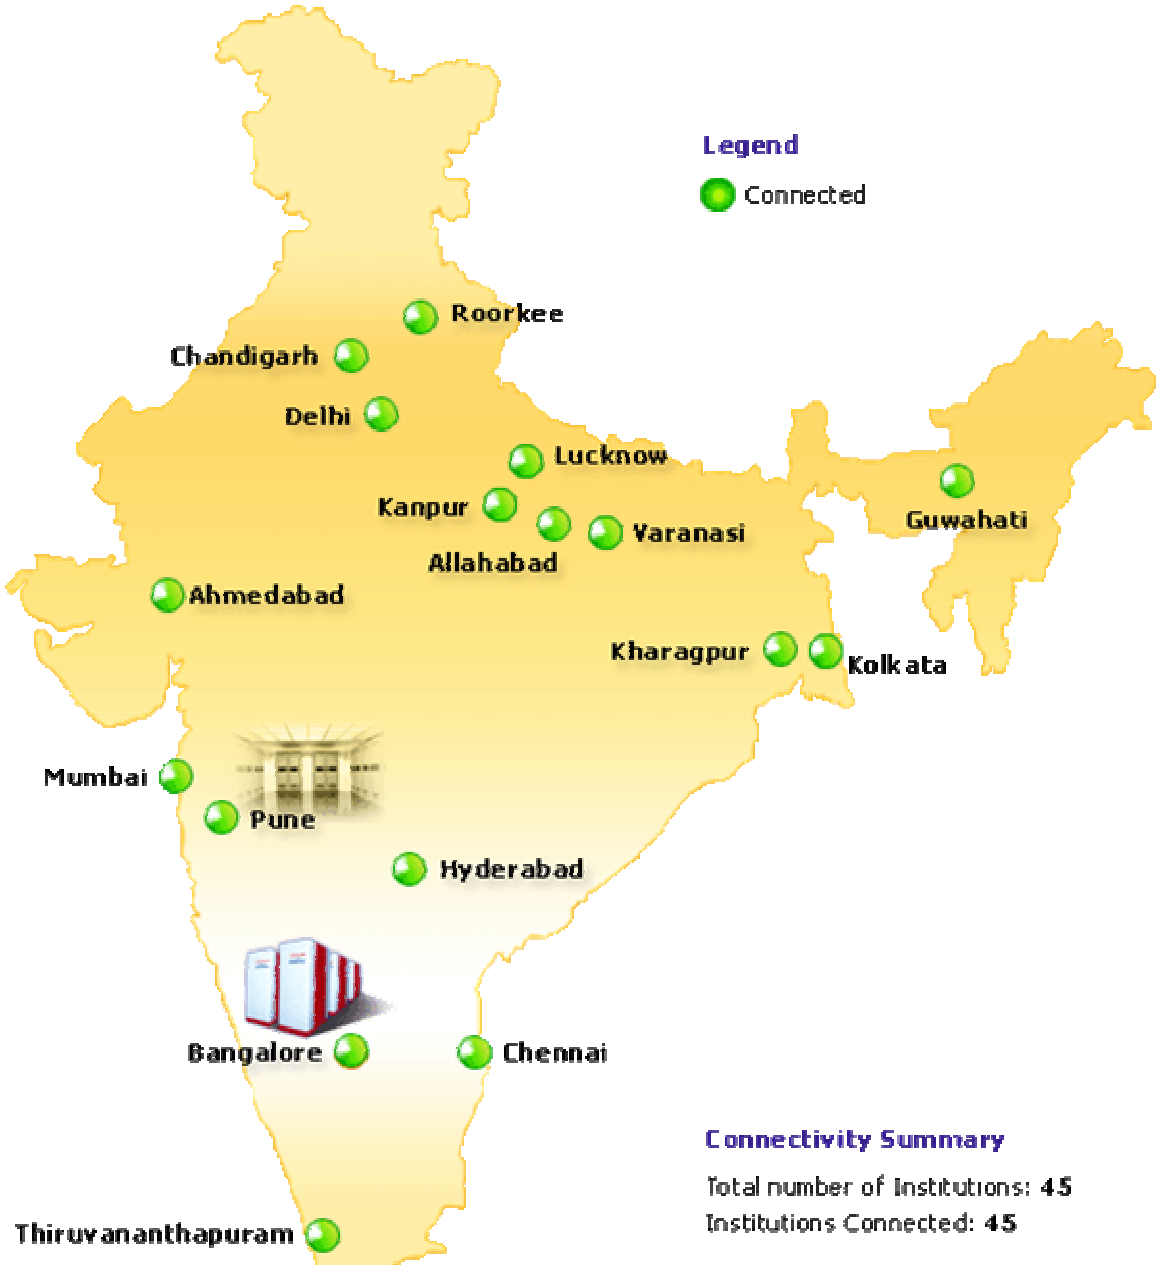

GARUDA: Objectives

- GARUDA aims at strengthening and advancing scientific and technological excellence in the area of Grid and Peer-to-Peer technologies. The strategic objectives of GARUDA are to:
  - Create a test bed for the research & engineering of technologies, architectures, standards and applications in Grid Computing.
  - Bring together all potential research, development and user groups to develop a national initiative on Grid computing Create the foundation for the next generation grids by addressing long term research issues in grid computing

GARUDA:Deliverables

- The following key deliverables have been identified as important to achieving the GARUDA objectives:
  - Grid tools and services to provide an integrated infrastructure to applications and higher-level layers
  - A Pan-Indian communication fabric to provide seamless and high-speed access to resources
  - Aggregation of resources including compute clusters, storage and scientific instruments
  - Creation of a consortium to collaborate on grid computing and contribute towards the aggregation of resources
  - Grid enablement and deployment of select applications of national importance requiring aggregation of distributed resources

EUIndia Grid

- What is EU-IndiaGrid?
  - EU-IndiaGrid supports interconnectivity between the European Grid infrastructure, EGEE, and the Indian Grid infrastructures, (Garuda India Grid - and Department of Atomic Energy Grid), to build a common infrastructure to support data processing for e-Science application areas, with a particular focus on Biology, High Energy & Material Science and Earth and Atmospheric Sciences
- Main Objectives
  - To build a dynamic research, academic and industrial community that may benefit from the use of Grid technology
  - To promote the use of advanced Grid technologies via pilot applications in the focused areas
  - To publicize EGEE Grid technology achievements in India and leverage on Indian Grid experiences and skills

EUIndia Grid

- The partners of this project are prominent actors on the European and Indian e-Infrastructures scene.
  - Europe
    - Italian National Institute of Nuclear Physics, INFN
    - Metaware SpA,
    - Italian Academic and Research Network (GARR)
    - Cambridge University
  - International
    - Abdu Salam International Centre for Theoretical Physics (ICTP)
  - Indian
    - Bhabha Atomic Research Centre, Mumbai
    - Indian Education and Research Network (ERNET)
    - Variable Energy Cyclotron Centre, Kolkata
    - University of Pune, Pune
    - Saha Institute of Nuclear Physics, Kolkata
    - Centre for Development of Advanced Computing , Bangalore
    - Tata Institute of Fundamental Research including NCBS

Summary

Indian Physics Community is maturing for high performance computing including Grid Technology

Acknowledgement:

Vipin Bhatnagar
Satyaki Bhattacharya
Subhasis Chattopadhyay
Asit De
Pandurang Deshpande
Sourendu Gupta
Tapas Samanta
Jasbir Singh

Procurement Under process (VECC)

- Purchase order sent for 8 * 2* Quad Core Nehalum (200 KSI2k) HP blade server inside HP 7000c Blade chasis for Kolkata-Cream.
- Committee recommended for 2 * 25TR (W+ S) cooling solution with de-Humidifier and Micron Filter and full rearrangement of Grid Facility.
- Cooling solution Public Intent is now under process
- Planning to procure 200 Core and 150 TB of raw storage. Complete all Pledged resources for Kolkata-Tier2 this year.
- Planning to procure 2 Layer-3 Gigabit Switch and 2 Layer-2 Gigabit Switch.

Tier2 IndiaCMS Network Connections

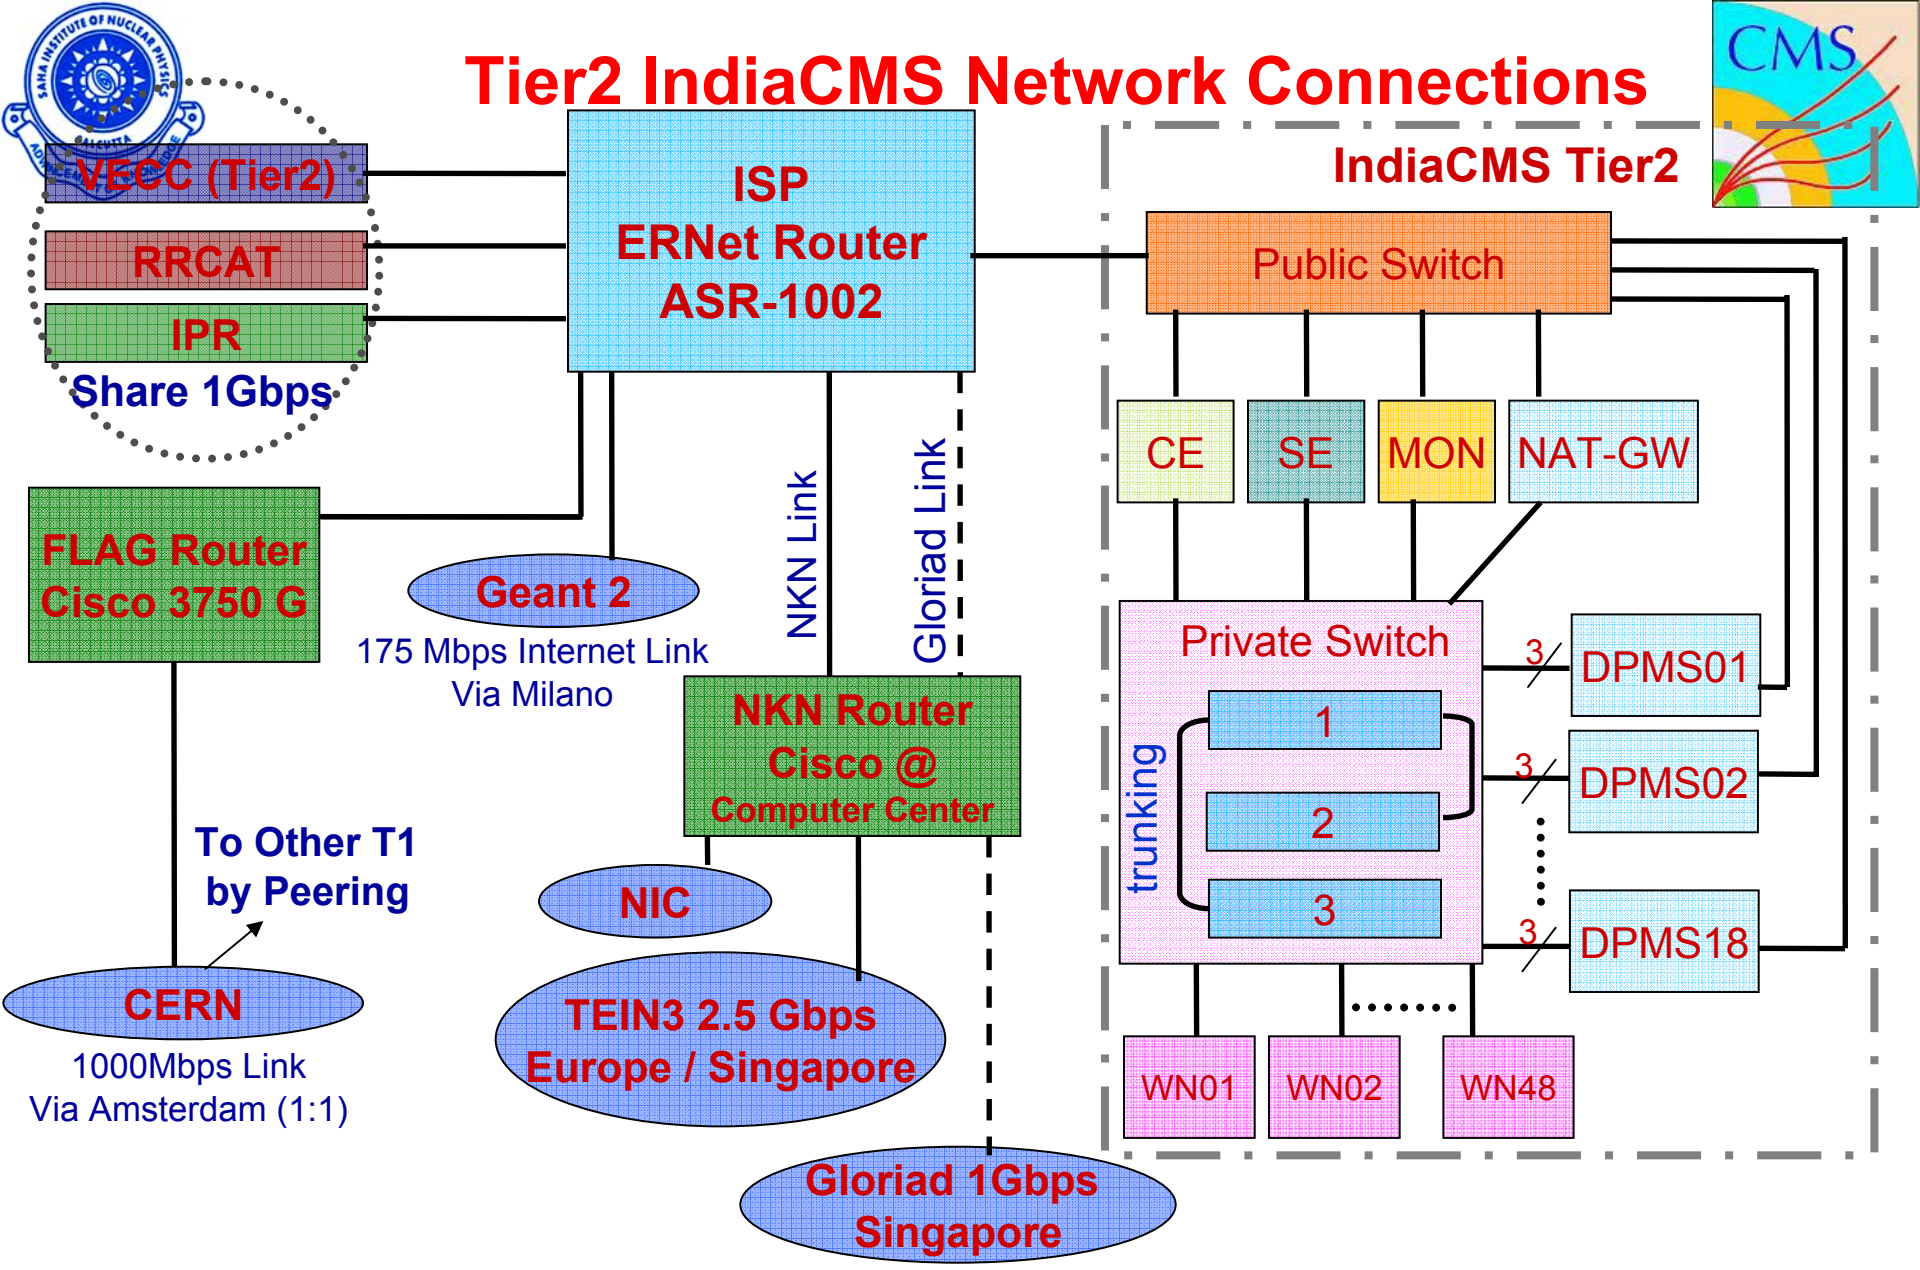

Typical Data Flow