Design considerations for the Indigo Data Analysis Centre.

Anand Sengupta, University of Delhi

Thanks to
- Maria Alessandra Papa (AEI)
- Stuart Anderson (LIGO Caltech)
- Sanjay Jain (Delhi University)
- Patrick Brady (UWM)
- Phil Ehrens (LIGO Caltech)
- Sarah Ponrathnam (IUCAA)

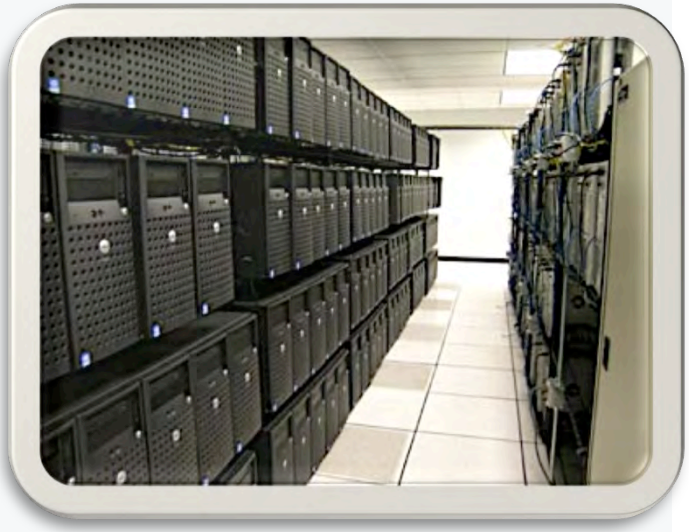

Network of gravity wave detectors

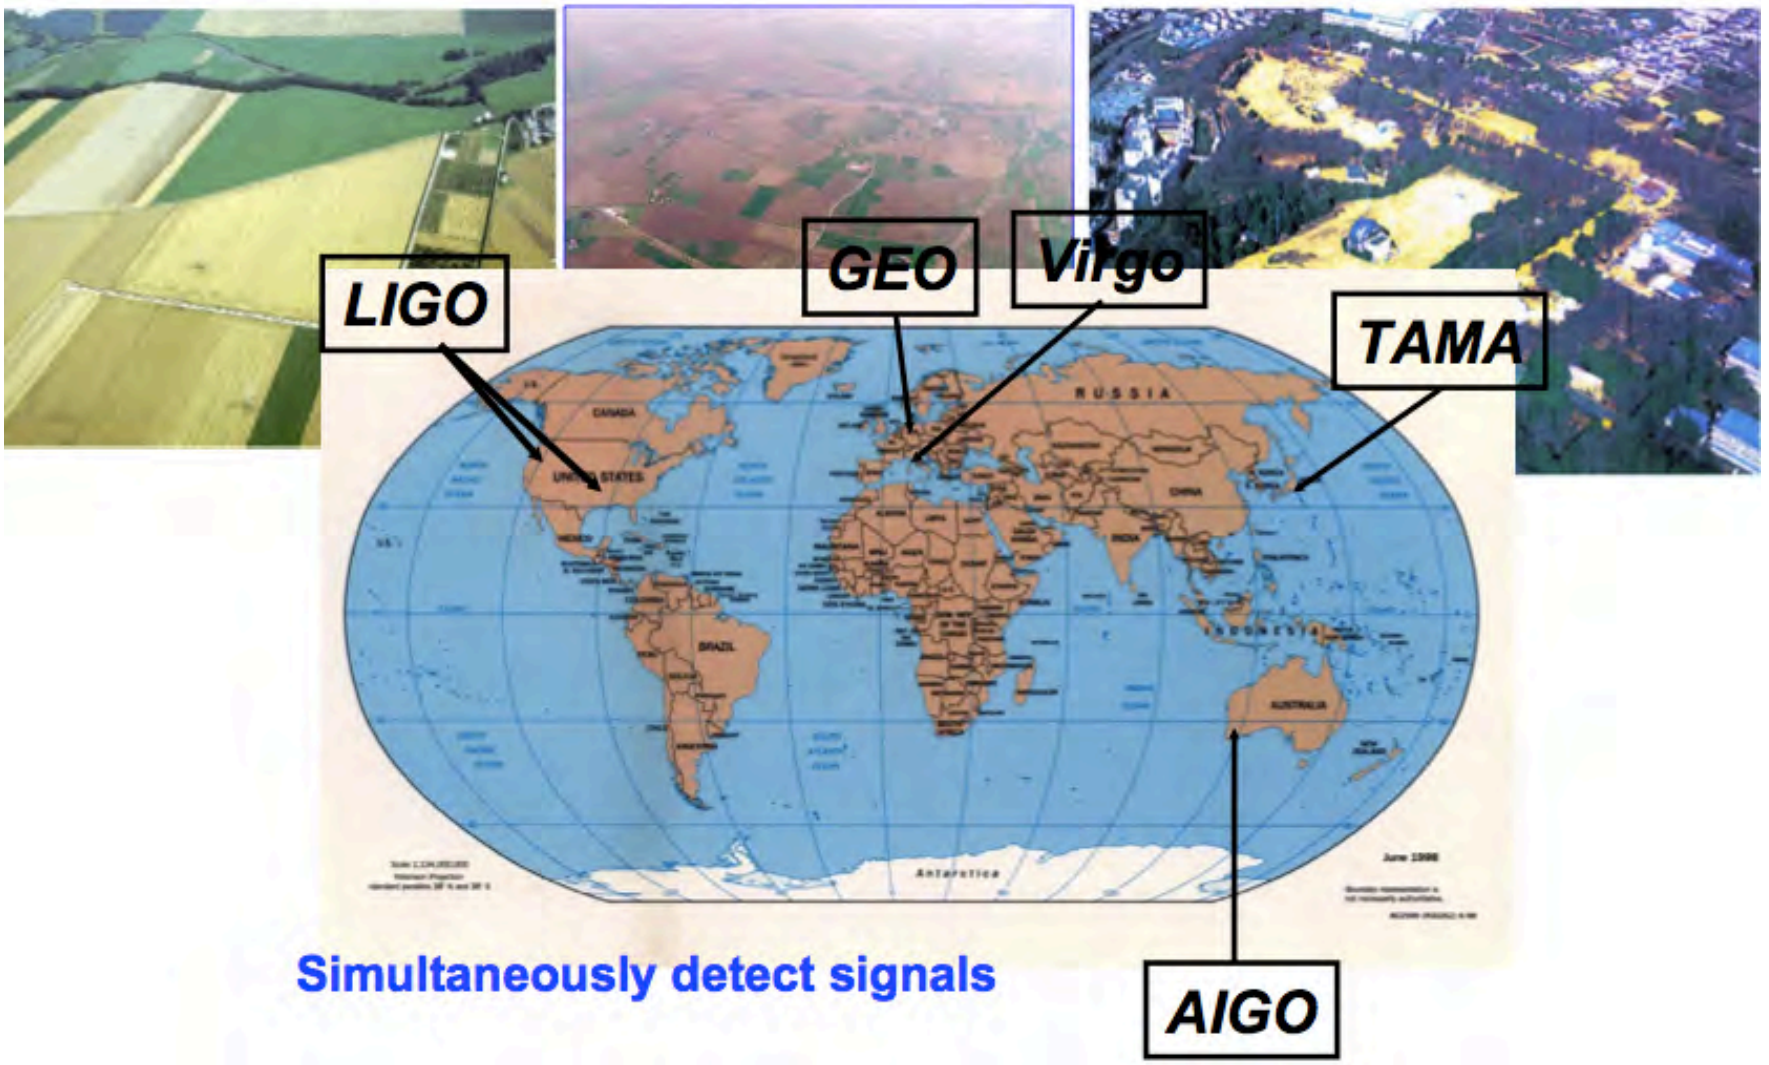

Simultaneously detect signals

Data from gravitational wave experiments

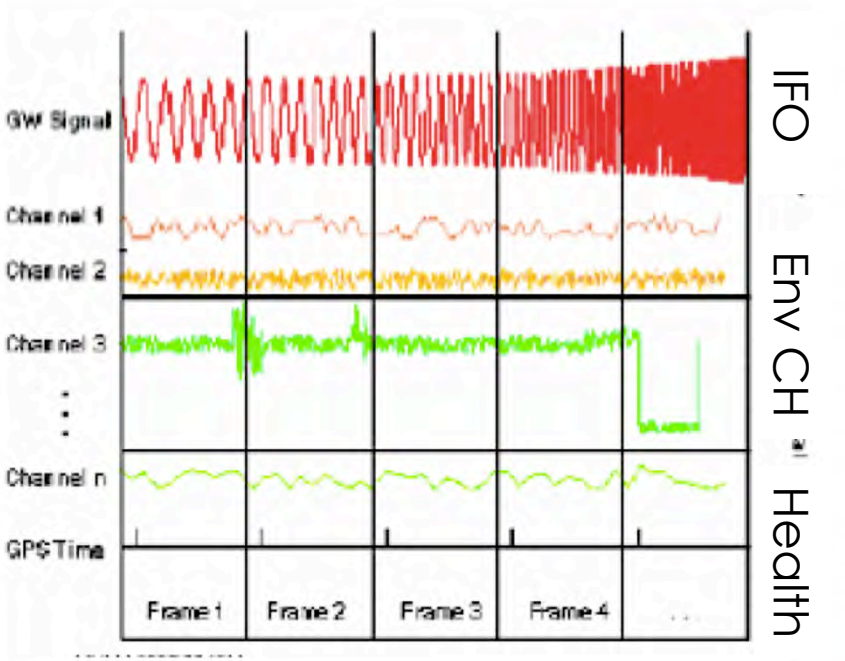

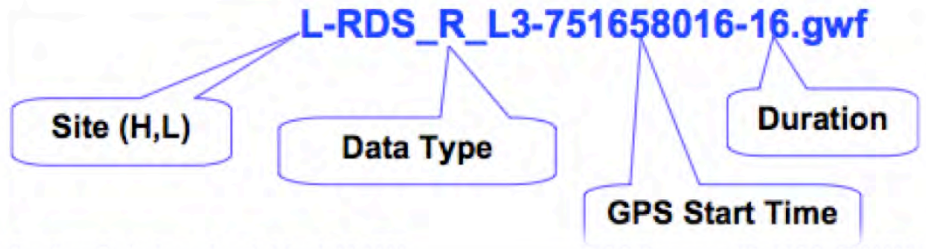

- Data comprised of
  - Gravitational wave channel (ASQ)
  - Environmental monitors
  - Internal engineering monitors
- Multiple data products beyond raw data
  - Reduced data sets
  - Level 1: gravity wave and environmental channels
  - Level 3: only gravitational wave data.
  - Different sampling rates

1TB of raw data per day!

The IndIGO data analysis centre

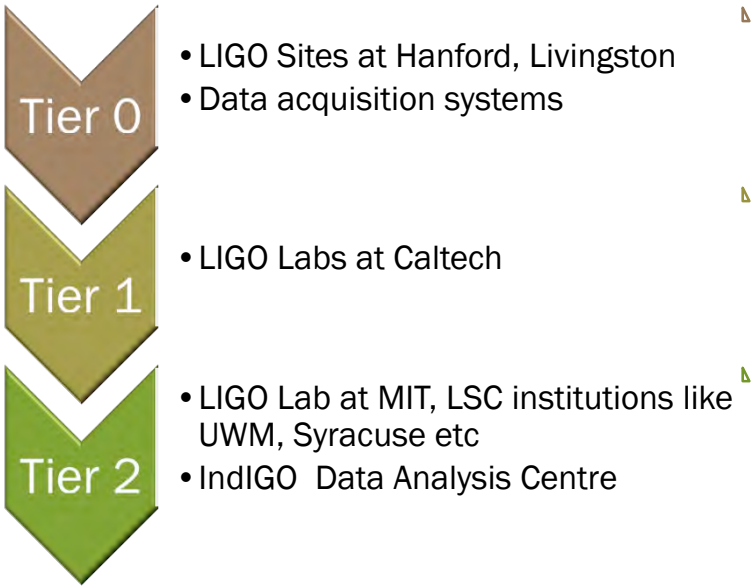

- Would like to propose for a high-throughput Computation and GW Data Archival Centre.
- Tier -2 centre with data archival and computational facilities
- Inter-institutional proposal for facility
- Will provide fundamental infrastructure for consolidating GW data analysis expertise in India.

Main objectives of the data centre

LIGO Data Grid as a role model for the proposed IndIGO Data Analysis Centre.

What is the LIGO Data Grid?

- The combination of LSC computational and data storage resources with grid-computing middleware to create a distributed gravitational-wave data analysis facility.
- Compute centres at LHO, LLO, Tier-1 centre at LIGO Caltech, Tier-2 centers at MIT, UWM, PSU, Syracuse.
  - Other clusters in Europe: Birmingham and the AEI
  - IndIGO Data Analysis Centre
- Grid computing software
  - E.g Globus, GridFTP and Condor
  - Tools built from them

LIGO Data Grid Overview

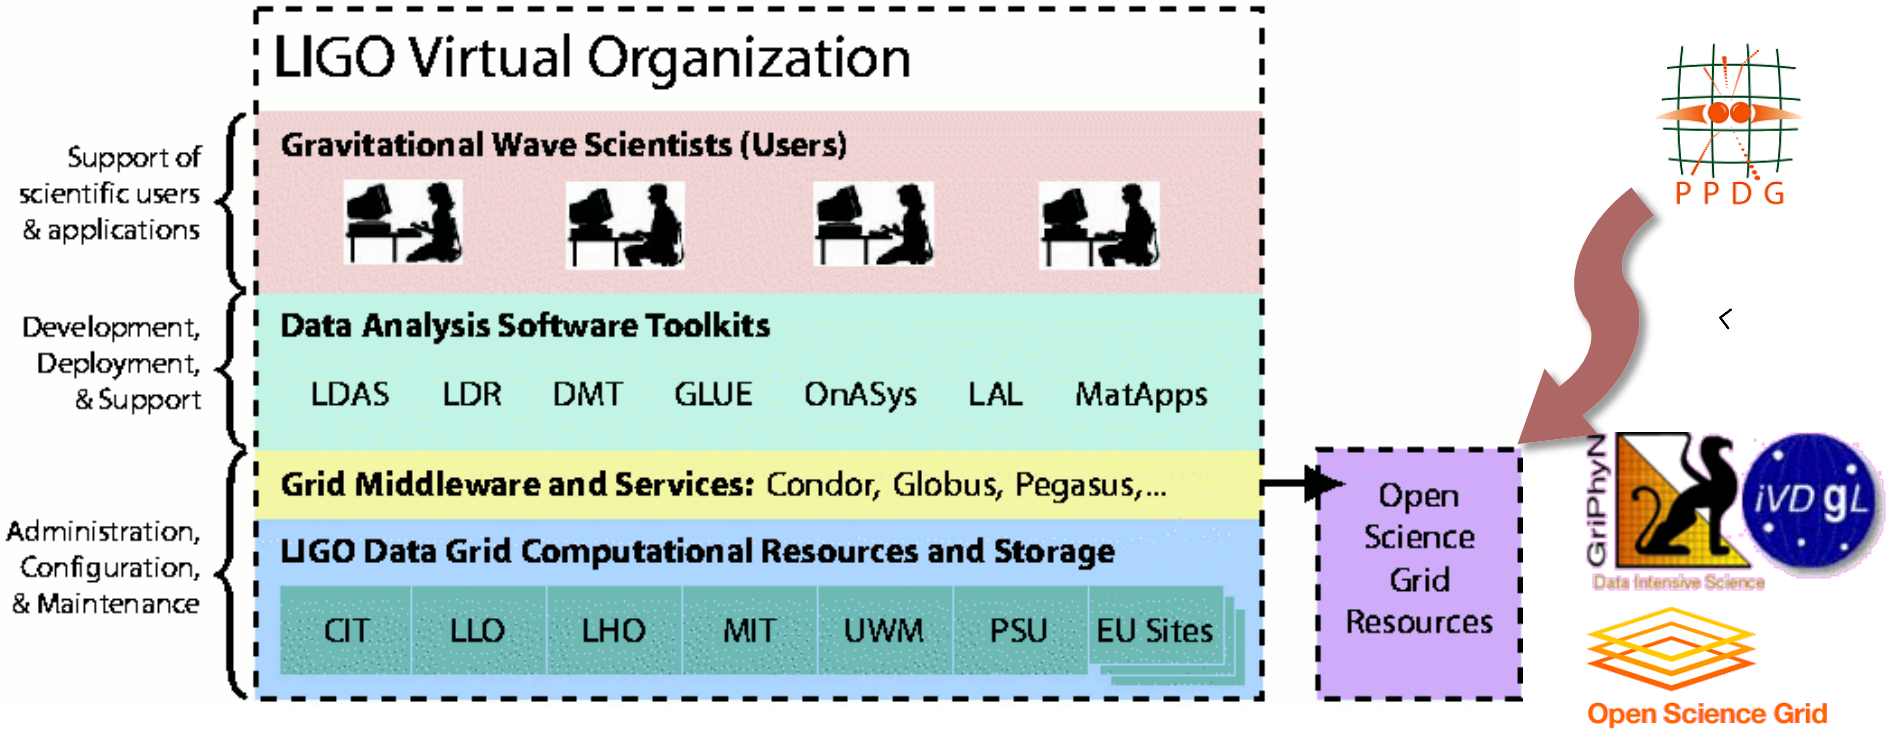

- Cyberinfrastructure
  - Hardware - administration, configuration, maintenance
  - Grid middleware & services - support, admin, configuration, maintenance
  - Core LIGO analysis software toolkits – support, enhance, release
  - Users - support

Condor, Globus, VDT and all that

- IndIGO Data Centre is envisaged to be a high throughput compute facility: (data volume driven)
  - Opportunistic scheduling, Condor
  - NOT a high performance computational facility, although one can imagine a synergy between GW users and other scientific users sharing the resources. Traditionally, MPI community requires dedicated scheduling.
- The Globus Toolkit is a collection of grid middleware that allows users to run jobs, transfer files, track file replicas, publish information about a grid, and more.
- All of these facilities share a common security infrastructure called GSI that enables single sign-on. Users can select any subset of the Globus Toolkit to use in building their grid. The VDT includes all of Globus.

Typical Work flow in inspiral pipeline

GLUE: LSC has developed in-house toolkit to write out work-flows as Condor DAGs

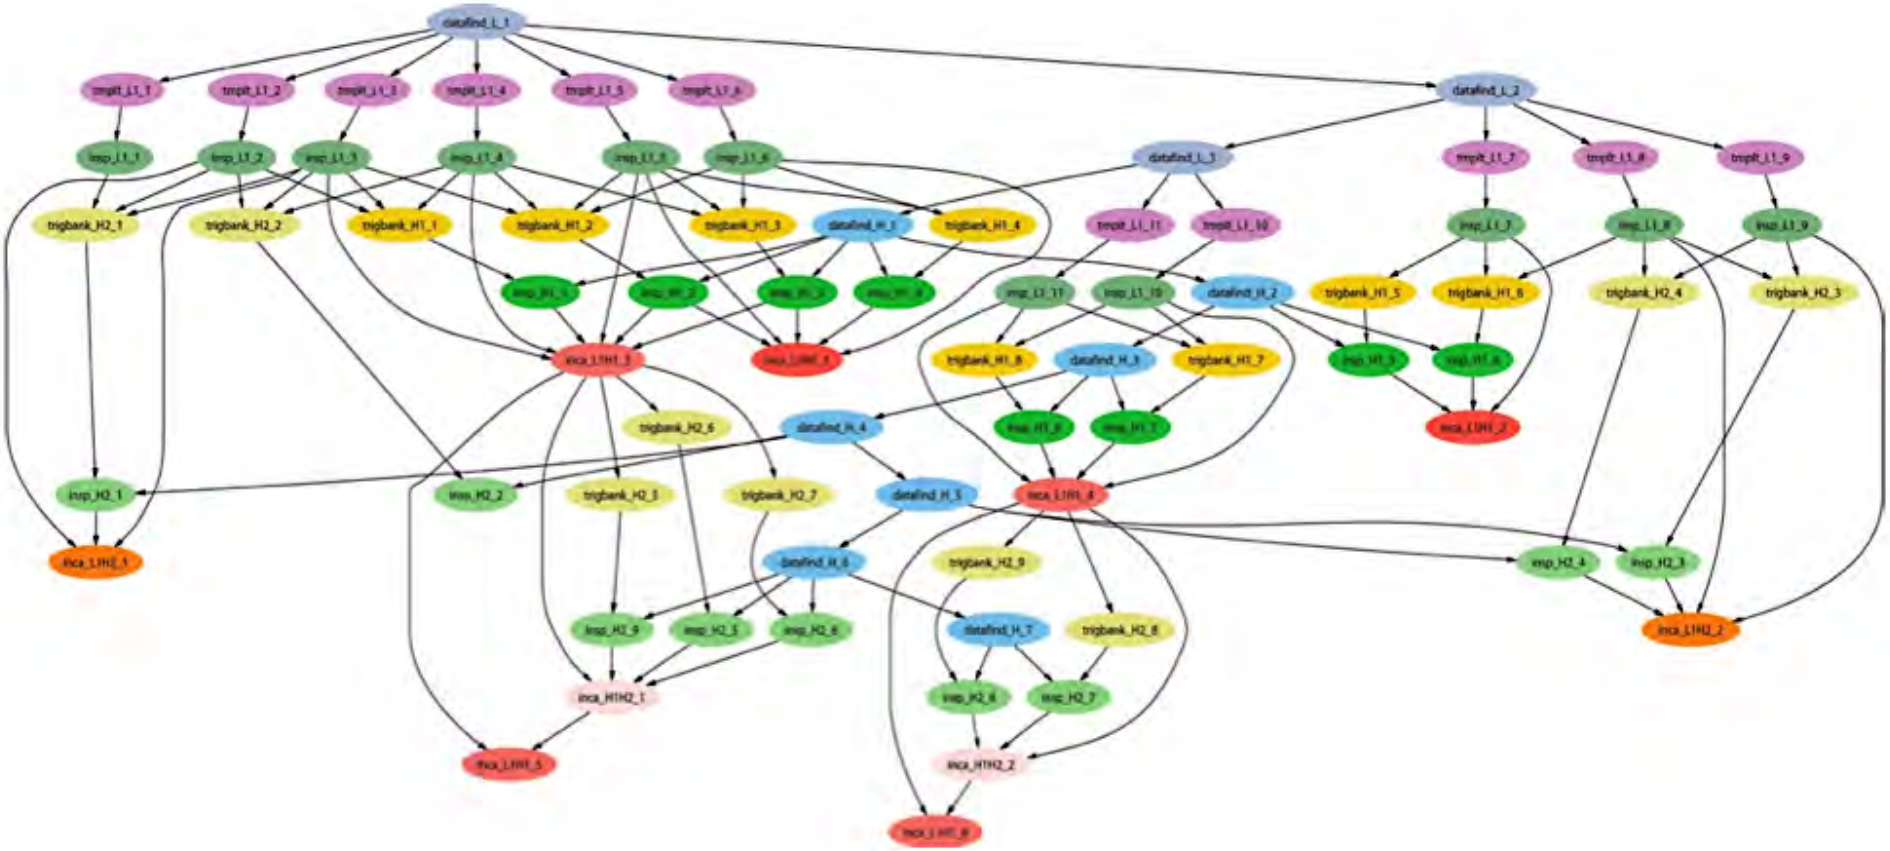

One month of data: 5 analysis DAGs containing ~45,000 jobs and few tens of Plotting DAGs each of 50 jobs.
For a year's worth of data, we run more than 500K+ nodes.

Why do we need the IndIGO Data Centre

- Maximum scientific exploitation requires data analysis to proceed at same rate as data acquisition
- Low latency analysis is needed if we want opportunity to provide alerts to astronomical community in the future
- Computers required for LIGO flagship searches
  - Stochastic = 1 unit (3 GHz workstation day per day of data)
  - Bursts = 50
  - Compact binary inspiral = 600 (BNS), 300 (BBH), 6,000 (PBH) .....
  - All sky pulsars = 1,000,000,000 (but can tolerate lower latency & .....)

Scientific pay-off is bounded by the ability to perform computations on the data.

Data Centre

Users and Usage

The current LIGO Data Grid (LDG) supports ~600 LSC scientists

Demand for resources is growing rapidly as experience increases and more data become available

The IndIGO data centre is expected to be setup on a similar footing

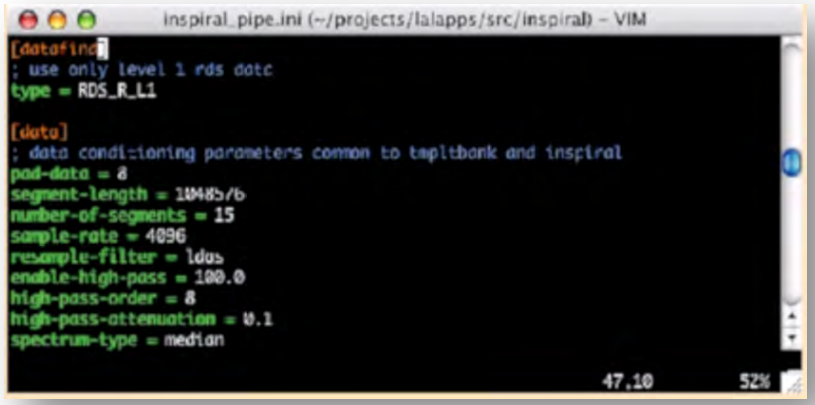

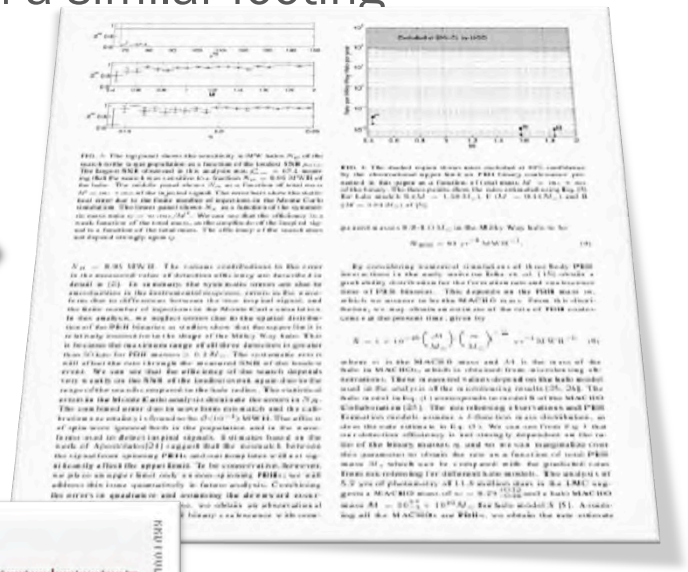

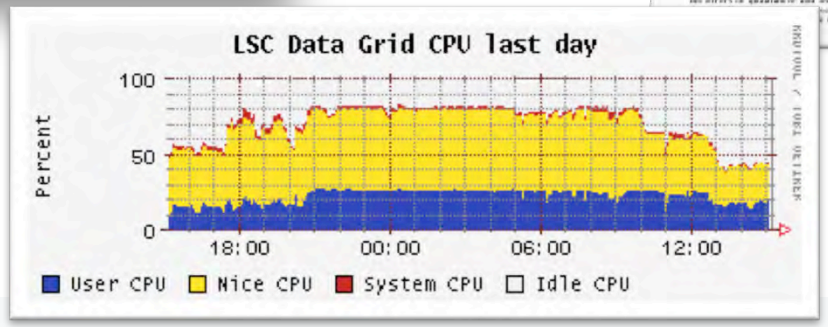

LSC computing resources

Distribution of LSC CPU cores

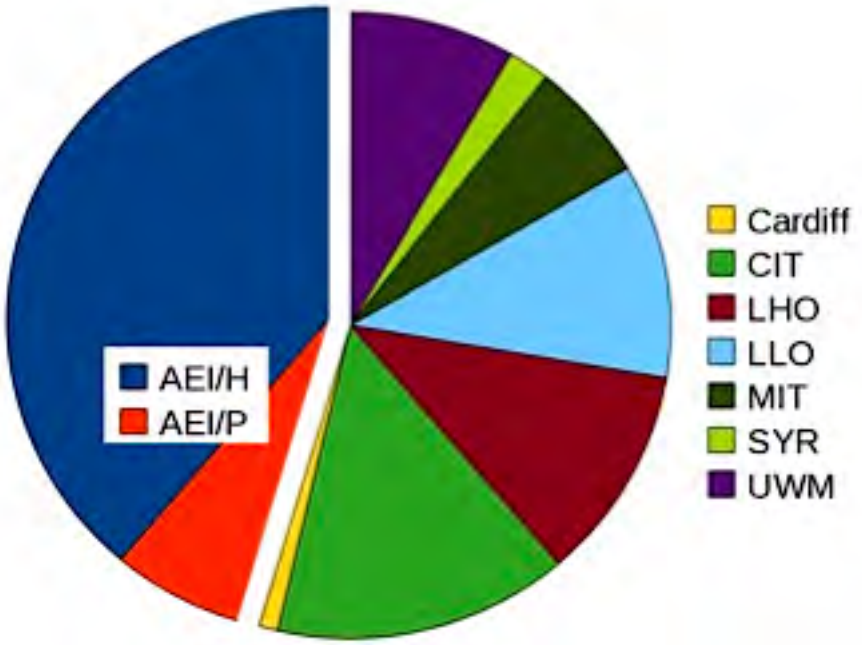

- LSC institutions and LIGO lab operate several large computing clusters for a total of 16,900 CPU cores.
- Used for searches and large scale simulations
  - Background estimates / assessment of significance
  - Pipeline parameter tuning
  - Sensitivity estimates, upper limits
- Analysis code-base: millions of lines of code
- Grid-enabled tools for data distribution

National Knowledge Network

- IndIGO data centre will need a high bandwidth backbone connection for data replication from Tier-1 centres as well as for users to use the facility from their parent institutions.
  - NKN can potentially provide this facility between IndIGO member institutions.
  - Outstanding issues: International connections, EU-India Grid
- The philosophy of NKN is to build a scalable network, which can expand both in the reach (spread in the country) and Speed.
- Setting up a common network backbone like national highway, wherein different categories of users shall be supported.

NKN TOPOLOGY

The objective of the National Knowledge Network is to bring together all the stakeholders in Science, Technology, Higher Education, Research and Development, GRID Computing, e-governance with speeds up to the order of 10s of gigabits per second coupled with extremely low latencies.

The major PoPs of ERNET are already a part of NKN – VECC, RRCAT, IIT (Chennai, Kanpur, Guwahati), IUCAA, University of Rajasthan.

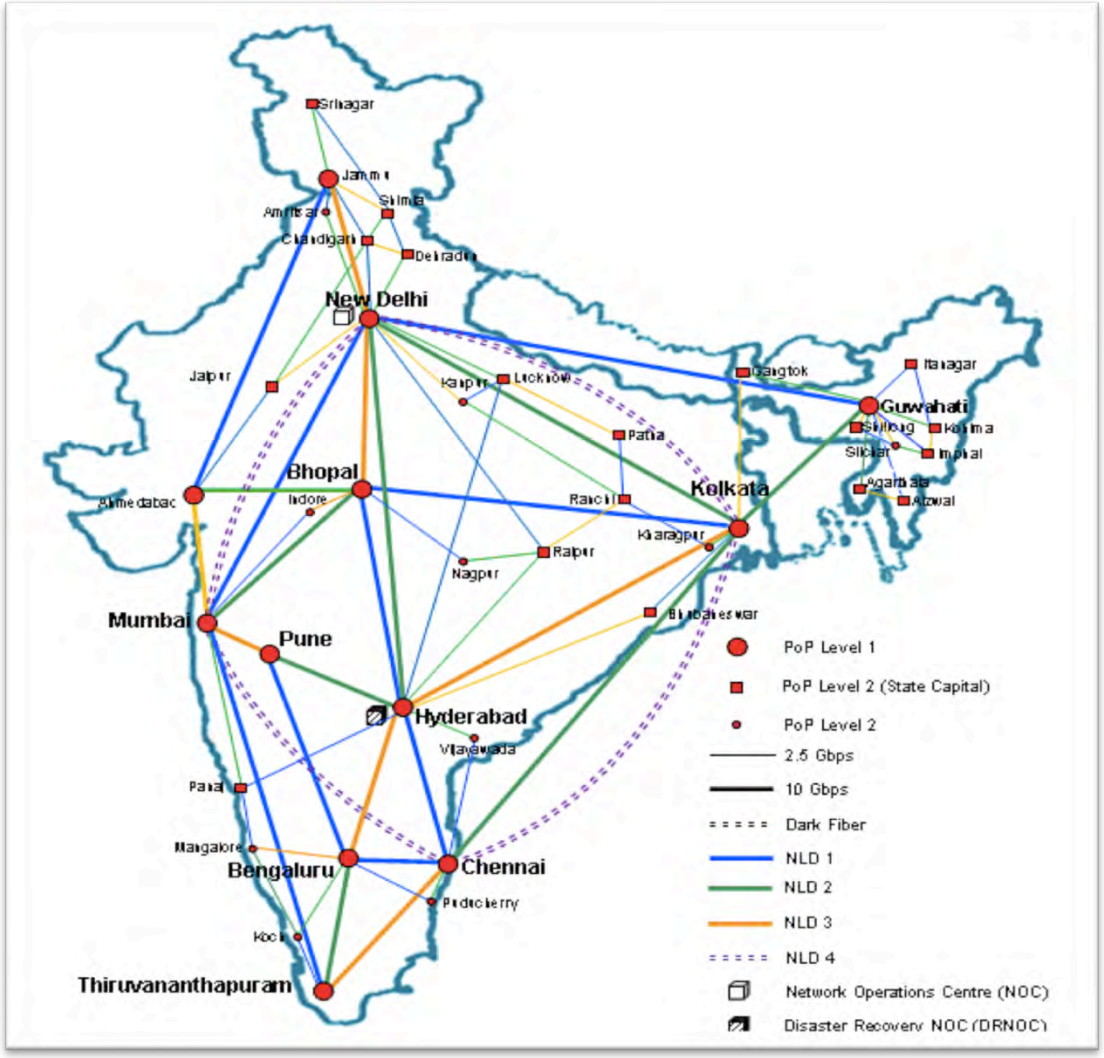

Collective wisdom

- Site selection / Bandwidth
  - IUCAA, Pune. Host to several large computational facilities. Delhi University?
  - External Gigabit Ethernet is probably sufficient although 10 Gig would be better.
- Storage, Cooling, AC
  - Typical: 1Pbytes on disk at Tier-1 centre. High throughput file system. RDS will require only a fraction of this at Tier-2 centre. Anticipate ~100TB.
  - At a rough estimate, 1000 cores = 35kW. Design Data Centre to hold 2/3 generation of equipment. Project 5-10 years in future. Need power to run the cooling itself, and power for disk storage and servers.
- Hardware / Cabling
  - Commodity off –the shelf computers, standard equipment racks. High density configurations. Co-exist with other user communities if need be. Typically top of the rack GigE switch to the machines in the racks and 10GigE to a central switch
- Middleware/Software/Security
  - Globus, VDT, Condor
  - Job management system
  - GSI for user authentication across LSC + IndIGO Consortium

Proposal Roadmap

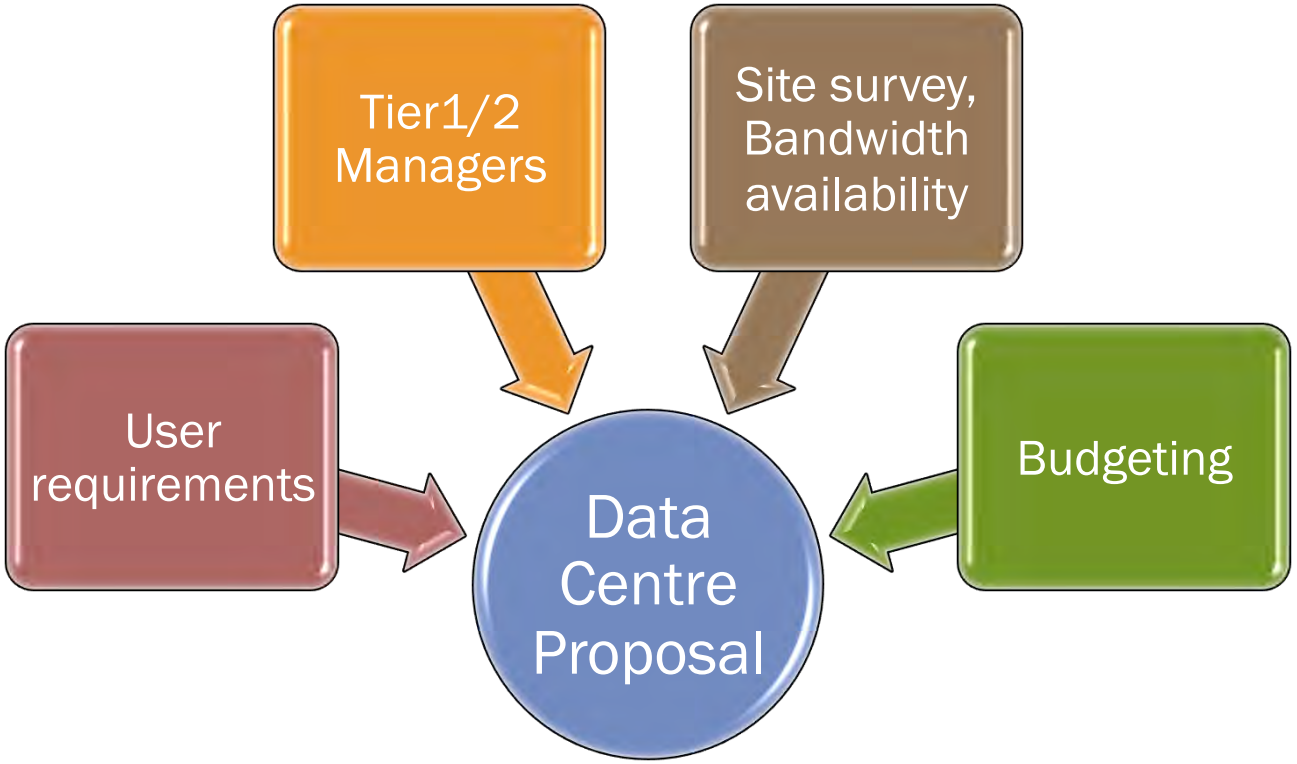

Proposal readiness by 15 May, 2011

Challenges

- Working with LDR and VDT involves a steep learning curve. Many new concepts. BUT, have a large user base and expert help.
  - Training system administrators and maintenance manpower
- Lot of uncertainties – bandwidth provider, site host, storage and node requirements etc. Ideas getting more concrete as we move along and start talking with LSC compute facility maintainers and experts from science and industry.
  - Very useful to visit a LSC cluster site (AEI Hannover e.g) and talking to the people involved in those centres.
- We should keep open the option of proposing this centre in conjunction with other (different kind, MPI based) scientific users. This would pose a host of challenges
  - Hardware, middleware and software requirements are different, hence some common ground has to be reached between groups.
  - Condor has a MPI environment – so MPI based codes are not a problem
  - Need to have this tested. Volunteers are needed.
  - Need to work out projections for next 5 years and gear up for Adv. LIGO and LIGO Australia.

Conclusions

- Need for a IndIGO data centre
  - Large Tier-2 data/compute centre for archival of g-wave data and analysis
  - Bring together data-analysts within the Indian gravity wave community.
  - Puts IndIGO in the global map for international collaboration
    - LSC wide facility would be useful for LSC participation
- Functions of the IndIGO data centre
  - Data archival: Tier-2 data centre for archival of LIGO data. This would include data from LIGO-Australia. LIGO Data-Grid Tools for replication.
  - Provide Computation Power: Pitch for about 1000 cores
    - Compare with AEI (~5000 cores), LIGO-Caltech (~1400 cores), Syracuse cluster (~2500 cores).
- Main considerations for data centre design
  - Network: gigabit backbone, National Knowledge Network. Indian grid!
  - Dedicated storage network: SAN, disk space
  - Electrical power, cooling, Air-Conditioning: requirements and design
  - Layout of rack, cabling
  - Hardware (blades, GPUs etc.), middleware (Condor, Globus), software (Data Monitoring Tools, LALApps, Matlab)
- Consultations with industry and experienced colleagues from Indian scientific community.