

LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY

- LIGO -

**California Institute of Technology
Massachusetts Institute of Technology**

Document Type: LIGO-T02135-00-E Sept 19th, 2002

**Path to Super Computing 2002:
LIGO-GriPhyN Demo**

**Kent Blackburn, Anne Carasik, Ewa Deelman, Scott Koranda, Albert Lazzarini, Karan Vahi,
Roy Williams**

Distribution of this document:

LIGO Laboratory - USC-ISI
LIGO/GriPhyN Applications Team

This is a publication
of the LIGO GriPhyN Applications Project.

California Institute of Technology

LIGO Project – MS 18-34

1200 E. California Blvd.

Pasadena, CA 91125

Phone (626) 395-2129

Fax (626) 304-9834

E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology

LIGO Project – NW17-161

175 Albany St

Cambridge, MA 02139

Phone (617) 253-4824

Fax (617) 253-7014

E-mail: info@ligo.mit.edu

<http://www.ligo.caltech.edu/>

PATH TO SC2002

**Kent Blackburn, Anne Carasik, Ewa Deelman,
Scott Koranda , Albert Lazzarini, Gaurang Mehta , Karan Vahi, Roy Williams**

Wednesday, September 11, 2002

The following draft aims to provide a roadmap is for the LIGO/GriPhyN Supercomputing 2002 demo, in November. This document is based on the discussions conducted during the meeting held at ISI on August 1st, 2002, attending were:

- 1) Ewa Deelman (ISI)
- 2) Gaurang Mehta (ISI)
- 3) Karan Vahi (ISI)
- 4) Kent Blackburn (Caltech)
- 5) Anne Carasik (Caltech)
- 6) Scott Koranda (UW)

The focus of this year's demo is LIGO's pulsar search, which is an analysis important to the LIGO' scientific community. The pulsar search involves the detection of binary pulsars in the gravitational wave signals collected by the interferometers at Hanford and Livingston.

From the point of view of computer science, we will showcase the newly developed Chimera Virtual Data System, which is the result of development conducted at ISI and University of Chicago.

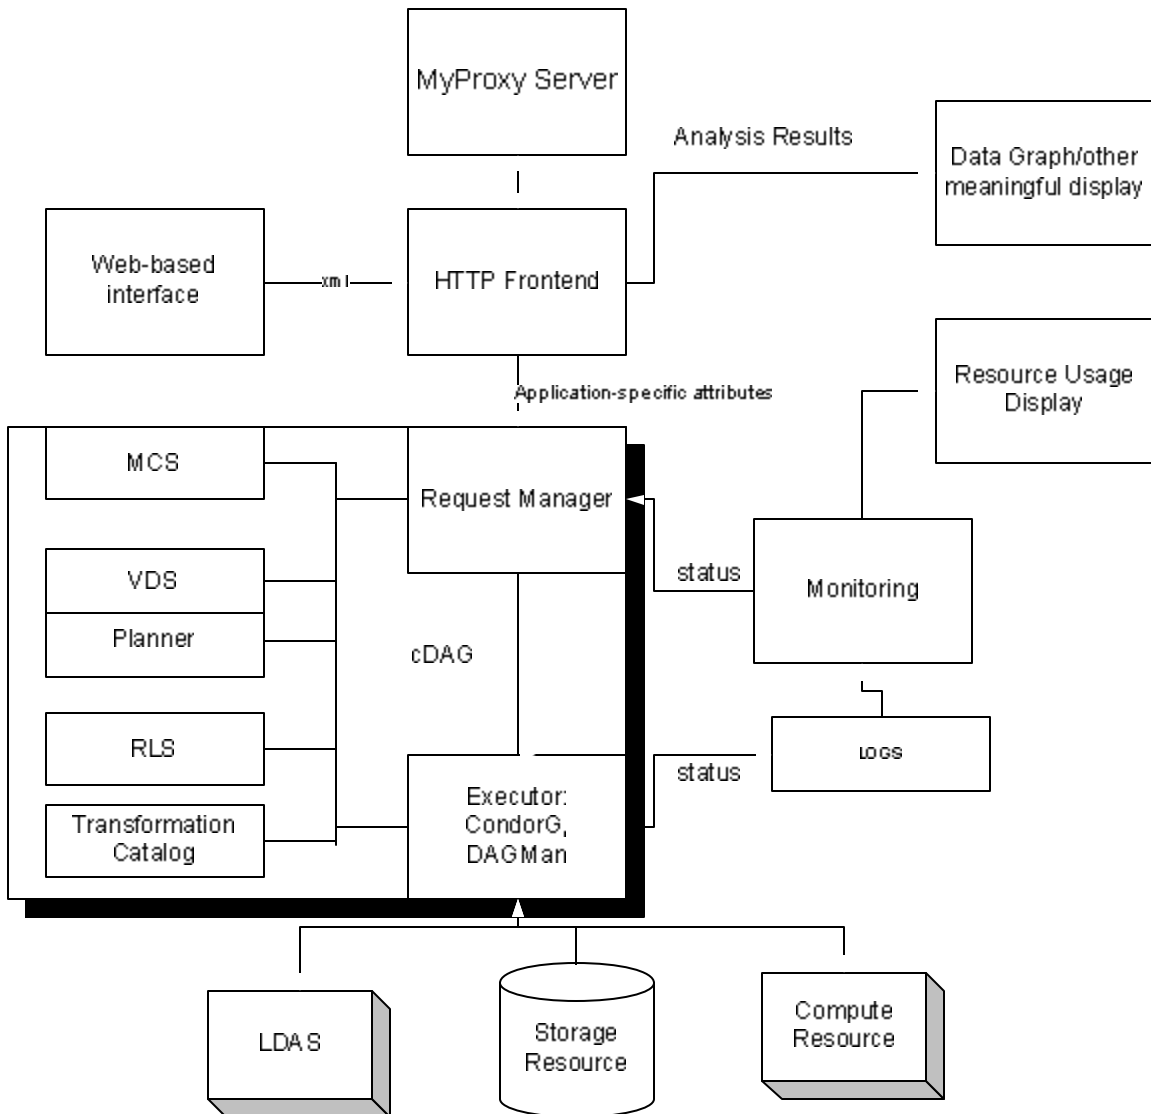
In order to conduct an all-sky pulsar search, a tremendously large amount of computational power and data are needed, more than the current LIGO resources are capable of supporting. Hence, in this proposed work, we will use both the computational resources of LIGO's LDAS systems as well as other grid resources to demonstrate the path towards supporting all-sky pulsar searches.

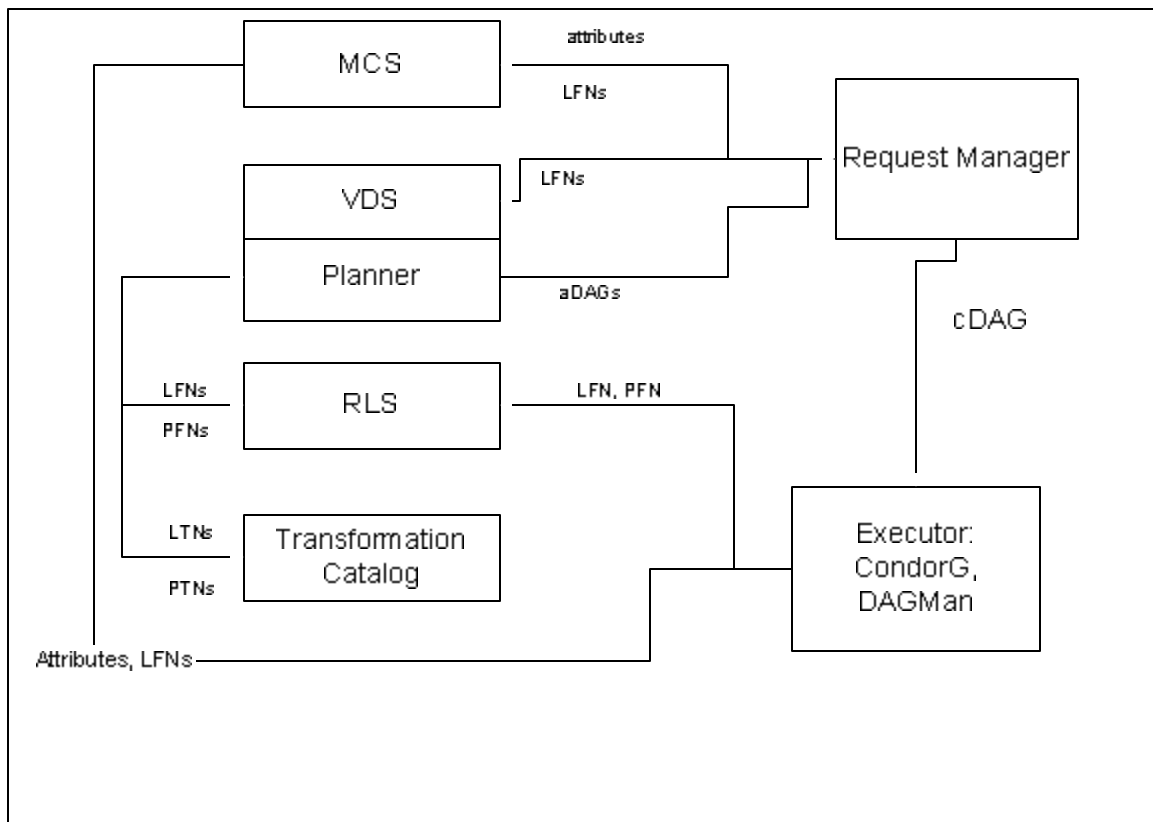
To enable LIGO analysis to be performed on grid resources, LAL shared objects wrapped in the LDAS wrapperAPI need to be ported to run in standalone mode on grid resources. Effort needs to be directed towards adapting this code (wrapperAPI/LAL/LALwrapper) to run in an MPI environment on a generic Grid compute resource with standard Globus GRAM jobmanagers, such as a jobmanager-pbs. In particular we are targeting the NCSA Platinum cluster.

Another approach, rather than submitting directly to a Globus jobmanager-pbs, is to submit the wrapperAPI jobs to a Condor pool in the MPI universe. We plan to investigate this approach both on the UWM Medusa cluster and the NCSA Platinum cluster.

A second approach to conducting the pulsar search does not use the LDAS wrapperAPI but instead is implemented as a stand-alone executable program. This hierarchical Hough transform code does not use MPI and is well-suited to run under Condor. We will investigate submitting such search jobs to the Condor pool at UWM via a Globus jobmanager-condor, and expanding the Condor pool via Condor glideins into the NCSA platinum cluster.

The architecture for the demo builds upon our system that we demonstrated during SC 2001 and is shown in the figure below.





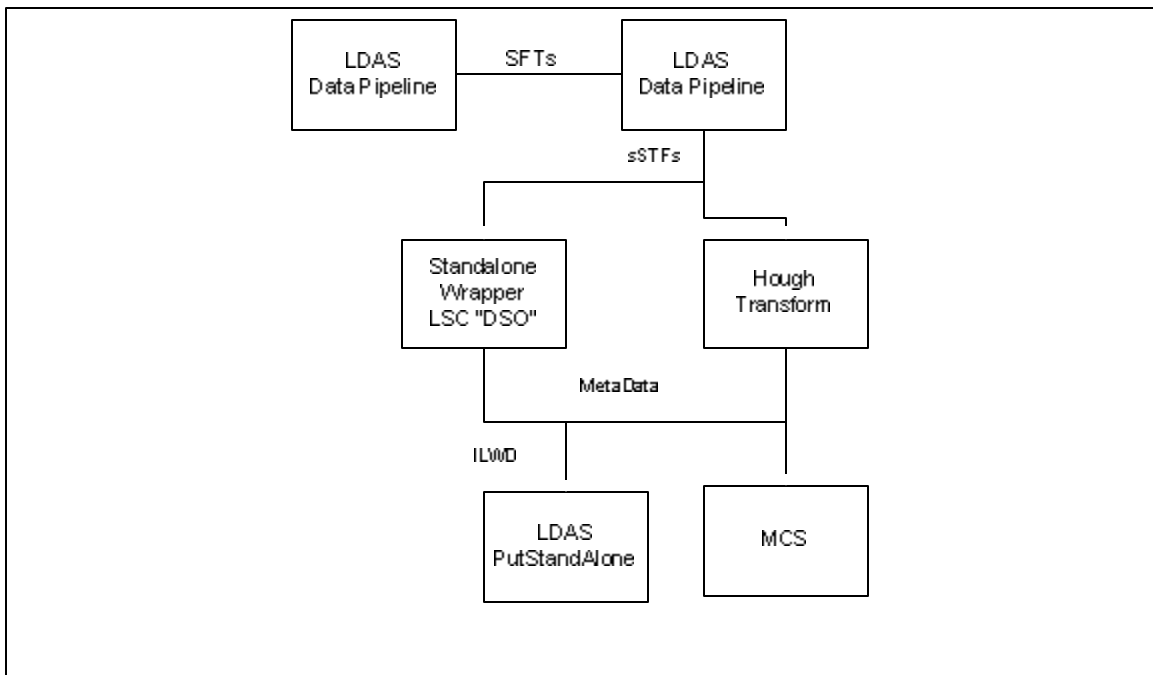
The components of the architecture behave in the following way:

- 1) The user interacts with the system via a web-based interface. As a first step the user authenticates to the system, and the system retrieves the user's credentials from the MyProxy server.
- 2) The user sends the application-specific attributes of the desired data product to the request manager. In case of a pulsar search the attributes specified would include location (celestial coordinates), time and a frequency range in which the pulsar search needs to be conducted.
- 3) The request manager queries the Metadata Catalog Service based on the attributes of the requested data. The MCS returns the logical filenames of one or more data items corresponding to those attributes.
- 4) These logical names are then sent to the Chimera Virtual Data System, which includes the Virtual Data Catalog (VDC). On the basis of the derivations recorded in the VDC, Chimera generates an abstract DAG that specifies the jobs that need to be executed to obtain the final logical data file. This abstract DAG only refers to logical entities.
- 5) The logical to the physical mapping of the data product and the transformations is done by the second part of the Chimera—the Planner—by looking up the Replica Catalog and the Transformation catalog. The Chimera planner queries the Replica

Catalog to determine the locations of all the logical files that are referred to in the abstract DAG. If any intermediate data products are found then the DAG is correspondingly reduced, for e.g., if the Replica Catalog returns the locations of the SFT files required for the run then the execution would bypass the generation of these SFTs.

- 6) The transformation catalog is queried to determine the paths of the executables on the execution sites. At present the pool where the jobs are to be executed needs to be specified. Later on, the system itself would determine the best execution sites based on the state of the network and the characteristics and availability of the resources.
- 7) The planner returns a concrete DAG (cDAG), where the physical paths of the transformations/executables and the replica locations are specified..
- 8) The Request Manager submits the cDAG to the executor: DAGMan.
- 9) DAGMan produces logs that are monitored by the Request Manager and a visualization system, which will show the activity: computation and data movement.
- 10) When new virtual data products are produced, the MCS and Replica Catalog are updated. If this new virtual data product is classified as an event, it is entered in the LDAS metadata catalog.

Some of the stages of the processing within a DAG are depicted in Figure 3.



Assuming that no intermediate data product for the run is already materialized, we need to execute the entire DAG. The first step is the generation of the SFT data products using the LDAS dataPipeline command on an LDAS system. The submit file to generate the SFTs are submitted to the managerAPI running on the LDAS system.

The SFTs are generated at the LDAS system using the dataPipeline command. Each SFT generation requires about 200 seconds to complete on an LDAS system. The pulsar search is planned to be run on up to 10 days of data. Each SFT contains approx 36 mins of data. Hence for 10 days of data we require 400 SFTs ($10 \times 24 \times 60 / 36 \approx 400$). These SFTs contain data from the frequency range of 0-2048Hz. The frequency corresponding to the desired range is extracted from each of the 400 SFTs. We call these data products sSFTs (sub SFTs). The sub SFTs corresponding to a particular frequency are then concatenated together. This concatenated product can then be shipped to grid resources to do the pulsar search. In the general case, the extraction and the concatenation of the SFTs can be done on another LDAS system, or on the same LDAS system where the SFT was generated.

These sSFTs are then stored in the LDAS system's storage area, or some storage system that will be accessible via gridFTP. (At present there are no GridFTP components installed on the LDAS's storage areas. But plans include adding these.)

The SFTs, after being transferred to the storage locations, are then registered with both the Replica Catalog and the MCS. The whole SFTs are registered, however, it is possible to move the stripped down sSFT data products (ILWDs to a temporary storage if necessary), and registered them in the RC and the MCS. The current assumption is that the whole SFT (the jumbo ones) will need to be registered as they are being stored on the LDAS site and each SFT generation requires 200 seconds of LDAS time.

Each sub-SFT data object (data corresponding to a particular frequency, a particular location and duration of up to 10 days) will then be sent to grid enabled resources, where the pulsar search standalone wrapperAPIs can be executed. This pulsar search code looks for excess power in the frequency corrected power spectrum over the length of the sub-SFT data set. If there is excess power at a particular frequency for the extended duration of time, the standalone wrapperAPI's search code identifies the metadata corresponding to the excess power as a candidate pulsar. This metadata is written to a file in the ILWD format that can next be used to record the candidate pulsar event in the LIGO database. This requires using GridFTP to move the metadata ILWD files to an LDAS system and then issuing a putStandAlone command.

Output

We would like to have a graphical representation of the output both in physics terms—celestial location for the pulsar search--as well as in CS terms---what is happening within the system, the resources involved, what data transfers are happening.

Code Development

Important components that are in development and necessary for the demo are:

- 1) The Standalone WrapperAPI: This wrapperAPI will allow the request to run LDAS wrapperAPI on a grid resource outside of LDAS, and would be used to do the scientific analysis required after the generation of the sub-SFTs. The SFTs will be generated by a LDAS system and kept in a grid accessible storage system.
- 2) MCS: This is the metadata catalog service which will be accessed by the Chimera VDS to determine the mapping of a users request consisting of domain-specific attributes to a logical filename or a set of file names.
- 3) The transformation catalog build on top of MCS (if time allows).
- 4) A more sophisticated planner

The DAG that we hope to execute for the Demo will look as follows

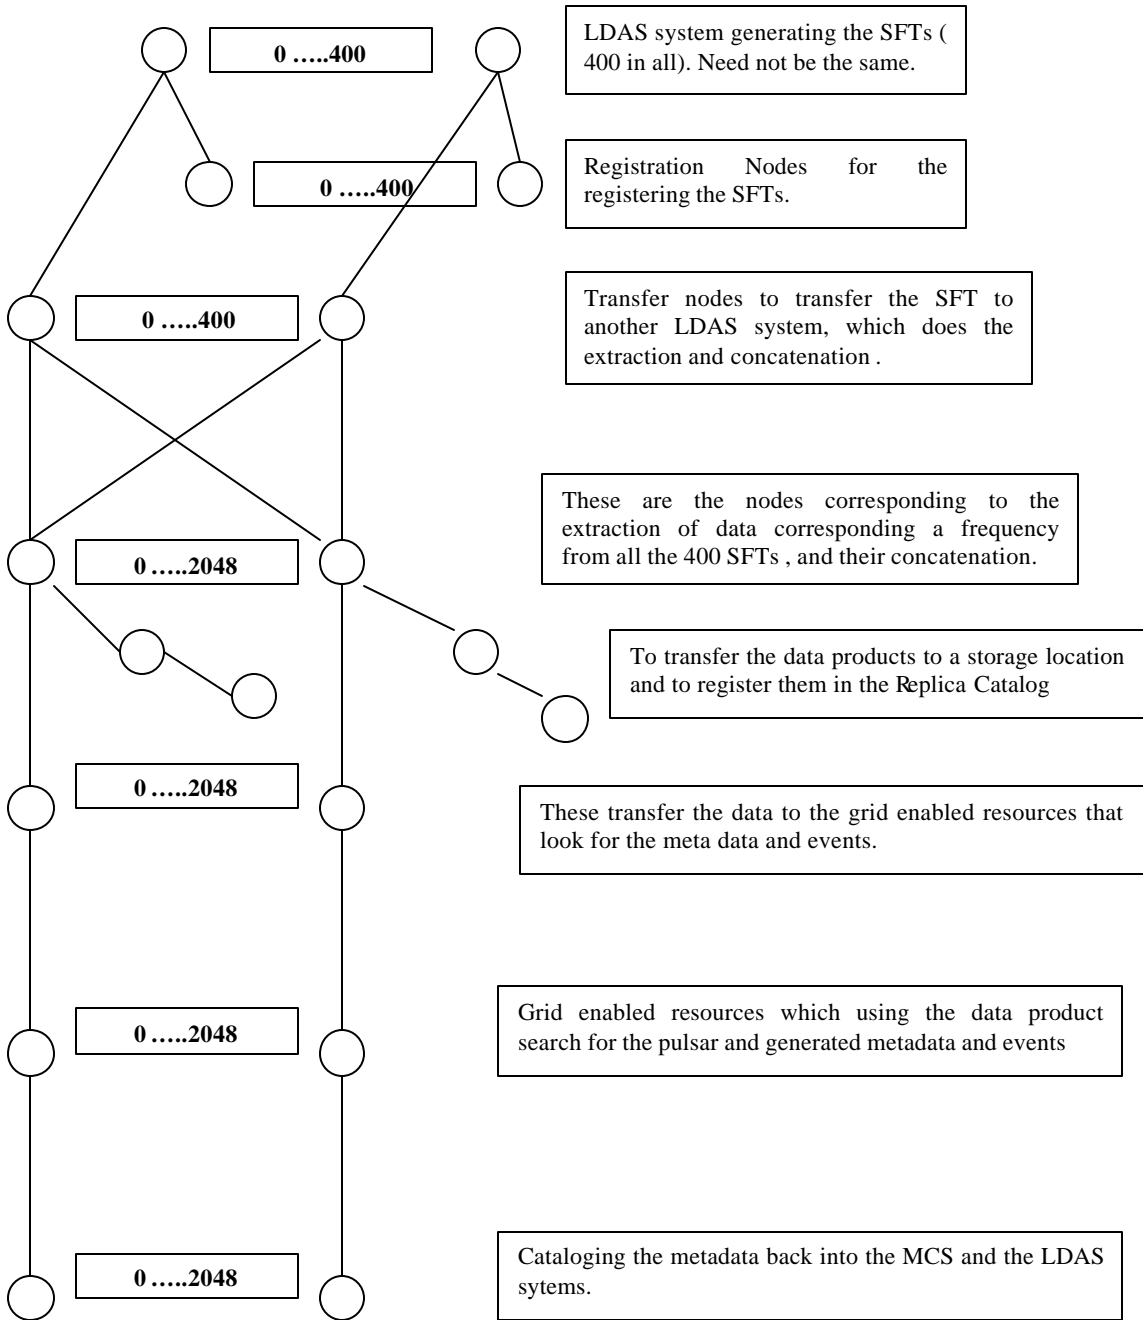


Figure 4: Ligo Monster DAG

Currently once all the standalone wrappers are up, Scott Koranda will provide a Concrete DAG corresponding to the pulsar search. From this DAG ISI will populate the VDC by writing the VDLt or VDLx. Time permitting one could write a generator, which generates out VDLx corresponding to LIGO to be inserted in the VDC.

TestBed

To run the pulsar search during the demo we hope to have access to resources at USC/ISI, Caltech and University Of Wisconsin. The detailed listing of resources is as follows:

1) USC/ISI:

- ISI storage plus a Condor Pool (consisting of 12 processors). The pool will have a gridftp server, jobmanager-fork and jobmanager-condor. *If we have time we could also setup a jobmanager-pbs. Would need Scott's help on that.*
- A Replica Catalog hosted at dc-user.isi.edu. We will upgrade to the new RLS code.
- Transformation Catalog (depending on whether the MCS is up by that time or not, otherwise a simple text file will be used).
- At USC we will have access to the HPC cluster. This is a 320 node cluster each with dual processors all running Linux 7.2. The cluster has grid-ftp, and PBS job manager.
- Also we will have access to the Almaak cluster, which is made of 64 processors all running Solaris. This cluster also is a part of the Condor pool at USC to which we will have access. It has jobmanager-condor and jobmanager-pbs.

2) Caltech

- At least one LDAS system.
- A linux beowulf cluster of 8-16 machines depending on the configuration at time of tests.
- 1+ TB of RAID storage for local data caches, a DB2 server and associated DB
- A number of SMP processors - linux (Xeon intel) and Sun (enterprise class servers). [CONDOR will not be available on LIGO/CALTECH systems]

3) UWM:

- One LDAS system

- a cluster which runs a 6.4.1 Condor pool and has a jobmanager-condor and a GridFTP server

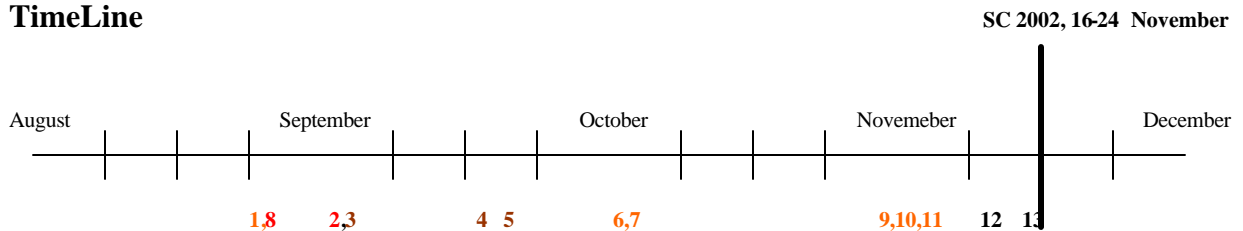
4) NCSA ?

- Probably access to < 64 nodes via a jobmanager-pbs, and probably a GridFTP server (we need to ask)

Important Dates and Responsibilities

	Task	Responsibility	Date
1	Interpool operability and migration of Chimera to globus-url-copy from the gsincftp.	Ewa, Karan , Gaurang , at ISI	3 rd week of August
2	Meta Data specifications of LIGO	Greg Mendell	Mid September
3	Standalone wrapper for generic Grid resource. Also the Huf transformations	Scott Koranda	Late august
4	Grid Stuff on the LDAS systems	Scott Koranda, Kent Blackburn	mid or late September
5	Concrete Dag by which to populate the Virtual Data Catalog	Scott Koranda	20 th September
6	Testbed at USC	Gaurang	Late September / early October
7	Population of the MCS	Ewa, Gurmeet, Karan	1 st October
8	TestMetadata	Kent Blackburn	Mid August
9	Transformation Catalog	Ewa, Karan, Gurmeet	Late October
10	AI Planning	Gaurang, Karan, Ewa	End October
11	UI for the Chimera. The map thing to pin point the location	Karan	End October – early November
12	Integration of all pieces, and the first test run	Everyone	Around 7 th November
13	SC 2002 (Crunch time)	Anyone and Everyone who can help	November 16th-20th

TimeLine



Key :

USC/ISI : Orange
Caltech : Red
Wisconsin : Brown

Other open issues:

Are we going to look into ways of visualizing dagman's output?