

Questions and Answers from the ETF Workshop held on September 20, 2002

Question:

Where have the meeting slides be posted?

Answer:

The slides presented at the September 20th meeting can be found on the TeraGrid Web site at <http://www.TeraGrid.org> . They are on the ETF Workshop agenda page.

Question:

I may have missed it, but what's the proposed schedule for the release of the TeraGrid specification document?

Answer:

The ETF specifications will be released as a series of documents on the ETF web site at <http://www.TeraGrid.org> , rather than as a single, monolithic document. The ETF sites expect to complete the first of these documents in a few weeks, with the complete set finished by the end of the year or early in 2003.

Question:

What types of proposals do you envision? Primarily connectivity with outlying computational resources, experimental equipment, or what? Will there be matching requirements? Are you interested in teams grouping resources together that extend the ETF?

Answer:

We are looking to enhance the diversity of resources available on the ETF, as well as add to its overall capability. We will not be paying for any new, or enhancements to existing, resources to be connected. NSF will only pay for networking costs, the hardware to connect to ETF, and the technical personnel who are needed to deploy the required hardware and software necessary to integrate the resource with ETF. No match is required. If a group proposal can enhance the capability of the ETF, then it will be welcomed.

Question:

Intel Pentium/Xeon and AMD processors are the most popular choices for clusters, especially in academia, and are projected to be for the next several years for price/performance reasons. Furthermore, Itanium marketplace acceptance is still unknown and is probably less certain than it was when the Itanium-oriented DTF/ETF proposals were formulated, as evidenced for example by Dell's decision not to build Itanium2-based systems. Thus, will very large compute resources based on commodity Intel Pentium4/Xeon and AMD processors be viable compute resources for ETF sites?

Answer:

We expect that future sites connected to the ETF will represent a diverse set of resources (computing,

data and storage, visualization, and instruments). There is no expectation that the processors of future sites be constrained to be Itanium family processors. The original ETF sites chose a homogeneous Itanium2 cluster configuration to simplify software deployment and integration of the initial ETF deployment. Appropriate configurations of future sites will be determined by the unique value they add to the ETF for national use.

Question:

If I understand correctly, FY03 will pay for the networking to make new ETF connections, but you will not pay for ANY personnel (e.g., to resolve the management and/or social issues of becoming part of the ETF). Is this true?

Answer:

This is not true. In addition to the hardware and connection costs involved in connecting to ETF, NSF will fund support of technical personnel who are needed to deploy the required hardware and software necessary to integrate the resource with ETF.

Question:

Where, exactly, will the funding necessary for research for the software itself (i.e., Grid software) for ETF come from? Surely there will be unique software requirements for the ETF, both from the middleware perspective and from the perspective of the higher-level user tools needed to facilitate collaboration. NMI is extremely valuable in this process, but it is not directly tasked to solve this problem. I'm concerned that NSF is perhaps too skewed toward viewing ETF as "merely a hardware problem", and NOT a software problem as well. Is there an "official position" from NSF regarding the sources for ETF-specific software research?

Answer:

NSF understands the multidimensionality of the ETF activity, as does the cyberinfrastructure advisory committee, which has suggested that software infrastructure, data infrastructure, and enabling research will be as important as hardware and networking systems. We expect some of the enabling research (in software and other areas) to continue to be supported by the ITR priority area, the NSF middleware initiative, and other programs as well as the terascale project itself. There are a number of FTEs funded through the ETF award that are directly associated with software development and deployment.

Question:

How will the proposed additions to the ETF be evaluated and selected? Is there a desire to have a distribution of new resources--computing, data, etc.--or will each proposal be evaluated on its own independent of other proposals (such that all new resources could be computing, or all be data, etc.)?

Answer:

All proposals submitted will be judged on their merits using the two usual NSF review criteria, intellectual merit and broader impact, as well as the other more specific criteria that are described in the solicitation, [NSF 03-553](#).

The overriding consideration will be the mutual benefit to the ETF and the partnering facility in the interest of enhancing science and engineering research and education opportunities. Unique capabilities or unique site expertise will also be strongly encouraged.

Question:

Is the review panel for FY2003 going to be the same as the "original" ETF review panel? If NOT, then how are you going to deal with the potentially inconsistent vision between the two panels?

Answer:

As we have done in the past, we will build knowledge of "program history" into the review panel. There has been considerable overlap on the TCS, DTF and ETF review panels, but for DTF added expertise was needed in certain areas, and reviewers were found to cover them. The same will be done for the review this year. We will attempt to assemble a panel with broad expertise.

Question:

How does NSF suggest incorporating other Federal Agency involvement in proposed extension to ETF? For example, many Federal agencies maintain important repositories that could be of great value to the scientific community. How does NSF expect funding policies to affect these potential partnership proposals with [local] universities, etc?

Answer:

There is a cross-agency organization called the Interagency Working Group on IT R&D (IWG), which is chaired by Peter Freeman, the Assistant Director of the CISE Directorate at NSF. Within the IWG, there are subcommittees on High End Computing and Large Scale Networking (and others) that meet monthly with representatives from each Federal Agency. Discussions are on going about Grid computing in these committees.

The ETF project will be as successful as the resources that it integrates. This clearly suggests resources supported by other agencies, universities, and other entities as well (hopefully international, too). Argonne National Lab (DoE) is already a member of the ETF community, so the project already has an "interagency" component. Other Federal agencies have also expressed interest in the ETF, and we are optimistic that other agency involvement will increase. The FY2003 solicitation is open to both academic institutions and FFRDCs.

Question:

Does the funding for ETF have to cover the cost for the current DTF locations and pay for a share of the ETF backbone costs? Is the funding for the current DTF already allocated via the DTF award?

Answer:

Funding for all ETF connections have already been covered either through the DTF and ETF awards or through institutional matching contributions to the DTF award. The Qwest-provided 40 Gb/s ETF backplane between the Chicago and Los Angeles hubs was funded via the original DTF award through March 31, 2006. Connections from the original four DTF sites to the ETF backplane hubs were funded

outside of the DTF award. Funding for the connection of TCS to one of the extensible hubs is included in the ETF award.

Question:

The DTF backbone costs are covered until March 2006. Should sites suggesting putting resources on the ETF discuss the costs for that time period, or just initial implementation costs?

Answer:

If a service provider is used, then the initial lease should extend at least through the period of the Cooperative Agreement. NSF anticipates making 5-year awards in FY2005 for extended management and operation of ETF through September 2009. These awards will include management and operations for all components of this facility. If it becomes necessary to extend leases for commercial service provider leases at this time, the additional costs will be built into the awards.

If a site chooses to buy dark fiber, then the NSF award could be used for this purchase. In this case, NSF would negotiate operations and management awards with the relevant ETF partners in FY2005 for continued management and operations through the end of FY 2009. The market is in a state of flux and many opportunities for favorable negotiations exist.

Question:

The network connections required to connect to the DTF backbone (lambda services) are typically very costly. Can you provide an estimate of the number of awards that are likely and the expected average dollar amount of the awards?

Answer:

Making predictions of networking costs has become more difficult recently, so the following are simply educated estimates. There are some fixed costs that will be covered by the FY2003 Terascale Extensions funding: the hubs and border routers. Our estimates are that these should cost about \$1.25M. Beyond that, distance from the hubs comes into play. Depending on how distant a site is, and how good an arrangement with a service provider can be worked out, we estimate that the total cost per site will be in the \$2.5M - \$5M range. So 3-4 awards may be all that can be expected.

Question:

Can you clarify the router structure one more time as far as the ETF router requirements and the site edge router requirements?

Answer:

The ETF backplane consists of two sets of routers- hub routers and border routers. All backplane routers are considered to be part of an integrated backplane. Border routers, which are located at the site that is connected to ETF are managed by the site; however, the site border routers are also considered to be part of the backplane as opposed to being part of the site network.

The border routers (and hub routers) are dedicated to the ETF project, and are not shared resources.

The resources that are being connected to ETF at a given site are connected “directly” to the backplane border router. There are no intermediate firewalls, routers, or other devices between the local ETF resource and the backplane border router.

Because backplane border routers (located at all ETF sites) must be integrated with the existing backplane routers, and directly connected to the hub routers, the selection of backplane routers located at new ETF sites must be done carefully, and in collaboration with the ETF networking team.

More details are available in the Primer at <http://www.TeraGrid.org> .

Question:

The ETF seems to focus on providing funds for network links like a terascale connections program. How would a site that connects to I-WIRE ideally participate in the ETF, leveraging the fiber/connectivity already in place?

Answer:

It is expected that, in general, networking costs will be high for sites and facilities integrating into the ETF. In the case of a site that is already connected to I-WIRE, there will still be equipment costs involved with local border routers, and the connection at Chicago to the hub routers. In this case I-WIRE may provide the fiber connection, but not the routers necessary for the ETF integration.

Question:

How does the ETF network relate to the National Light Rail initiative?

Answer:

The ETF backbone between Chicago and Los Angeles is provided through a partnership between the original four DTF sites (ANL, Caltech, NCSA, and SDSC) and Qwest. The partnership involves a collaborative design and deployment of 40 Gb/s between the backplane hubs in Chicago and LA and “end-to-end” monitoring and operational support for the backplane between the four sites. The four original DTF sites have all provided their own connections to the hubs in Chicago and Los Angeles taking advantage of available fiber or laying new fiber. PSC is still engaged in negotiations to provide a link between Pittsburgh and Chicago, and several options are under consideration.

There is a loose consortium of institutions that are exploring the idea of a “customer-owned” dark fiber network, and this consortium is called “National Light Rail” or “NLR.” Some institutions that are involved in the NLR activities are interested in using this approach to connect to the ETF backplane. Thus NLR represents one of many options that a site may consider to connect to the backplane.

Questions:

Will ETF job schedulers accept 3-week Gigaflop jobs that would take only a half hour on a teraflop system, or perhaps even 1,000-fold larger jobs that would take only a half hour on a potential petaflop

system, especially jobs that would require compute resources to be used in connection with the ANL visualization resource?

Answer:

The ETF, and more generally all of the NSF high-end computer systems target large-scale user applications that may not be feasible to consider running on academic institutional computer facilities. Hence, the goal is to support access to very large computing, storage, and visualization resources -- resources generally beyond the capabilities of single universities. We expect the ETF batch scheduler to support such large jobs. Similarly, the ETF will support submission and execution of large jobs that require concurrent access to multiple ETF resources (e.g., computing, storage and visualization) that may not be co-located.

Question:

I'd like my archived data collections hosted by ETF resources to be visible to the outside grid world. How can collections, which may be archived by ETF, be registered with replica catalog services belonging to my VO (virtual organization). Will ETF create its own RC (replica catalog) service and metadata catalog service?

Answer:

The ETF will provide data grid infrastructure for creating replica catalogs and digital library infrastructure for managing metadata.

The ETF resources will be used to host collections for multiple research projects. We currently see three types of access:

1. Data sharing within a project. Data Grid technology is used to create a logical name space that can be used to create a global identifier for sharing data with team members.
2. Publication of data. Digital library technology is used to organize a collection that can support discovery of published data.
3. Preservation of data. Persistent archive technology is used to support replication of data into archives.

Through the SDSC Storage Resource Broker (SRB) technology, one can:

1. Register existing archived data collection into a SRB logical name space, by creating logical names for each digital entity.
2. Replicate registered digital entities onto ETF resources, and actually make a copy of the data.
3. The data collections would then be accessible through any of the APIs that are provided by the SDSC SRB, including Web browsers, Windows browsers, Unix shell commands, C library calls, etc.

We would need to know more about the replica catalog services that are provided by your Virtual Organization to understand what is required for the registration of digital entities stored on ETF

resources. We anticipate the Open Grid Services Architecture to specify standard WSDL services for the registration of digital entities into replica catalogs.

Question

The implementation of a national 'visualization' resource is a new concept. What software and services will ETF offer to enable the national user community to take advantage of the large parallel rendering system at ANL?

I assume packages such as WireGL/Chromium will be offered, but this is a solution only for OpenGL codes. What else will be developed and deployed and what kinds of support will be offered?

Answer:

Visualization resources have been offered via the Internet in various forms such as rendering farms. The ETF visualization services will build on this concept both in terms of "batch" capabilities as well as streaming capabilities.

The current ETF management organization has a visualization services working group that is in the process of testing and evaluating several dozen tools and libraries for potential inclusion in the ETF visualization services. This working group expects to finalize an initial set of tools and libraries in early 2003.

Question:

I would like to set up a data staging and access service for external VO's such as iVDGL and ATLAS (high energy physics experiment) at the University of Chicago. This would provide a point of service for large-scale data staging to and from ETF from external networks peered in Chicago: Abilene, Esnet, and the dedicated CERN and Amsterdam (Surfnet) links.

We have an I-WIRE termination in the Geological Sciences building. What additional costs would be associated with using the fiber optic link to the Chicago Starlight hub?

Answer:

It is expected that sites will connect using a minimum of one 10 Gb/s channel. The connecting site is responsible for the bandwidth and all equipment required to connect to the ETF hub router and to the site ETF resources. For a single 10 Gb/s channel this means (a) a 10 Gb/s interface to the hub router (specifically a Juniper T640 in the current ETF architecture), (b) a backplane border router located at the connecting site, (c) a 10 Gb/s WAN interface for the backplane border router, and (d) a 10 Gb/s LAN or multiple 1 Gb/s LAN interfaces for the backplane border router.

The backplane border router must be close enough to the resource being connected that these 10 or multiple 1 Gb/s LAN interfaces can be directly connected without intermediate IP networks, firewalls, or other devices.

The 10 Gb/s bandwidth between these two routers requires optical fiber end-to-end. In some cases this

can be provided by a commercial service provider. In other cases the long distance portion may be provided by a commercial service provider and the local connection at the site may be provided by the site.