

Part of report to EMBRACE WP3 meeting at Hinxton on June 9, 2005./N-EE

Please note some ---remarks---

Extracted from original SweGrid proposal

---(as of October 31, 2003)---

Nils-Einar Eriksson, Uppsala Biomedical Centre (BMC), Computing Department and the Swedish EMBnet node, Linnaeus Centre for Bioinformatics (LCB), Uppsala

Summary:

To fully exploit the potential of GRID computing for bioinformatics, solutions will be needed that ensure the availability of updated bioinformatics database information at the GRID end nodes. These solutions have to be flexible and allow easy co-localisation of related information. In addition to semi-static fundamental databases like the EMBL nucleotide database, it should be possible to easily transport user generated database material to the end nodes. Mechanisms to split database information for parallel access on the vast number of GRID nodes should of course be provided.

We propose, for SweGrid, the re-implementation and further extension of those database distribution mechanisms that were developed at the Swedish EMBnet node more than 10 years ago. These mechanisms were tested and used in a full-scale pan-European production environment during some five years in the 1990's.

To enable us to create this, mostly infra structure, part of SweGrid, we need a rather limited allocation of storage and processor resources as well as cooperation from the systems managers.

Considering the time needed for some start-up work to be done locally, we estimate that the following SweGrid resources will be needed for the coming half-year period:

1. Up to 500Gbytes of disk-space in each of the six file servers.
2. Adequate CPU-time for database I/O- and formatting operations in the file servers. (Probably less than 100 hours in each file server, mostly independent of end node operations, i.e. concurrent with other general SweGrid use of the end-nodes).
3. Five Gbyte of allocated end-node disk space in each end-node. (An alternative, but more time consuming, would be to fetch prepared end-node files from the file server).
4. Up to 10 hours of end-node time per end-node for updates and testing.

Expected result: A database environment suitable for bioinformatics, available to all SweGrid users.

We thus apply for the above described resources.

Making data available at Grid-nodes, problem description and a suggested solution.
In bioinformatics a set of data is often compared to vast amounts of known data stored in databases.

How do we make these databases available to the Grid end nodes?

Even if we know which the most important databases are today, the method chosen should be of a general nature so as to accommodate the transport of new databases.

Obviously, we don't want to burden each Grid node with all database information. On the contrary, Grid parallel computing suggests that we only need to have very small parts of the database available at each node.

Some factors that influence the optimal degree of local database granularity are:

- The number of available Grid nodes.
- Need for redundancy. Some nodes will be down. Some regions might be unreachable due to data communications problems. The network topology obviously has to be considered.
- The size of a natural database granule. (This should be considered with respect to our knowledge of how, commonly used, application programs work. Grid applications for other disciplines might need less considerations of this type.)

Multiple complete sets of databases will not only contribute redundancy but will also give the corresponding number of possible parallel queues, subject to restrictions inherent in the supervising software.

In some applications we need frequent cross-references to related parts in other databases e.g. when comparing sequence information to structure information. This should be considered when allocating database granules to individual end-nodes.

To manage all of this, an elaborate service, will be needed to govern and to keep track of what data, is where and if it is up-to-date or not. A distribution mechanism with qualities mentioned above will be needed. In addition to computer science knowledge, bioinformatics domain knowledge will be essential if we want to get an optimal result. Resources should be allocated for the study of all related aspects and for the subsequent design of a solution. Prototyping, testing and some production processing could probably be performed at various grid test bed sites. In Sweden, we received funding to set up a test bed with 600 nodes divided into six groups with lots of storage capacity at each group centre, i.e. a structure similar to the overall structure of the planned European Grid. Swegrid will be operational this autumn. (A year-long successful Globus based test activity within the NORDUGRID project justifies this optimism. However, the test team has not dealt with our type of database distribution requirements.)

Proposed solution

BMC's computing department is a service organisation that housed the Swedish national EMBnet node from 1989 until recently, when the node was affiliated to the Linnaeus Centre for Bioinformatics (co-located at the BMC). During this time a rather unique experience as regards distributing bioinformatics database information to distant, domestic and international nodes, was gained.

We propose that this knowledge and associated software should be put into practical use within SweGrid, a grid test bed comprising some 600 computers organised similarly to the large European grid. We also offer this combined ---**experience**---to EMBRACE, a proposed EU Network of Excellence for Bioinformatics.

The software that was developed by Peter Gad at the Swedish EMBnet node for database distribution within EMBnet consists of a set of programs. Together, these programs constitute the realisation of a protocol, xNDT, that was assigned a 'Systems well-known port' by the Internet authorities.

xNDT was the major carrier of data within EMBnet for 3-5 years and has thus been tested thoroughly in a demanding, high volume, production situation.

--- (License conditions for possible use outside of SweGrid will have to be negotiated with Peter Gad as well – some kind of GPL-license should be attainable./N-EE 2005-06-07)---

As a second, equally important step, building on working routines for database distribution to the grid filesystems, we propose to establish mechanisms that will either bring the data to its final destination in the grid end-nodes or prepare the data to be fetched by each grid end-node from their respective filesystems. In both alternatives the data will be selected according to the optimisation considerations mentioned above. The extent and speed of these implementations will depend on the level of additional project financing. At least one of the primary bioinformatics databases will be implemented in the initial phase.

Annex

xNDT – brief history and possible openings for reuse in GRID computing.

In the early 1990-ies, a data distribution problem was solved at the Swedish EMBnet node at Uppsala Biomedical Centre, BMC. (EMBnet, the European Molecular Biology network, <http://www.embnet.org> , now consists of a world wide net of cooperating bioinformatics service centres.)

DNA sequence data from e.g. the Human Genome Project, was collected at an ever increasing pace at labs all over the world. This data was entered into, and copied between, databases at three different sites, Genbank in the USA, EMBL Data library in Europe, and DDBJ in Japan.

However, the search for and analysis of data was performed at decentralised national or local resources. In Europe, the database information was spread in a flat-file format to the national EMBnet nodes. Some of these files were quite large and the transmission of large sets of such files was often interrupted by technical problems (slow links, network congestion). In Sweden, 5-10 sub nodes got their data from the national node at the BMC. To simplify the process of keeping these sub nodes updated a set of client-server programs were written by Peter Gad who was one of the employees at the EMBnet node. Together, these programs constitute the realisation of a protocol that was named xNDT (extended Network Data Transfer, <http://www.embnet.se/ndt/>). It was adopted at almost all European EMBnet nodes (as shown by EMBnet's chairman Peter Rice at one of the 'Genes, Proteins and Computers' conferences at Chester. It was given a 'Systems well known port' by the Internet authorities - a rare achievement. Peter Gad wrote versions for a number of different systems including VAX/VMS, Silicon Graphics/Unix and SUN/Unix. The xNDT suite of programs thus was used in a demanding, large-scale multinational production environment during some 3-5 years. As the quality of data communications networks in Europe improved and when Peter Gad left the Swedish node for a job at a bioinformatics company, xNDT was replaced by script-governed FTP-transmissions.

However, xNDT still exists. It has recently been rewritten in Java and it can perhaps be put into use in a modified version to provide data for GRID-cluster end nodes. The central data-provider would, as in xNDT, ask the peripheral nodes what their status is, i.e. what files they have - including checksum validation.

If this answer is different from what was expected a 'data update push' would be initiated. Today, xNDT can feed information directly into a DBMS, e.g. MySQL and Oracle. Finally any reformatting of the data could be performed and the node would report a ready status to the central 'database- and job-submission facility'.

To minimise the need for long distance transmissions of data, xNDT was designed to enable a hierarchical structure. The intention was to have the EMBnet nodes take care of national redistribution. This was done in Sweden.