

# Life Sciences & The Dutch Grid: An Analysis from a Grid Supporter's perspective

Lammerts, E.<sup>1</sup>,

<sup>1</sup>e-Science Support Group, SARA Computing and Networking Services, Science Park 121, 1098 XG Amsterdam, The Netherlands

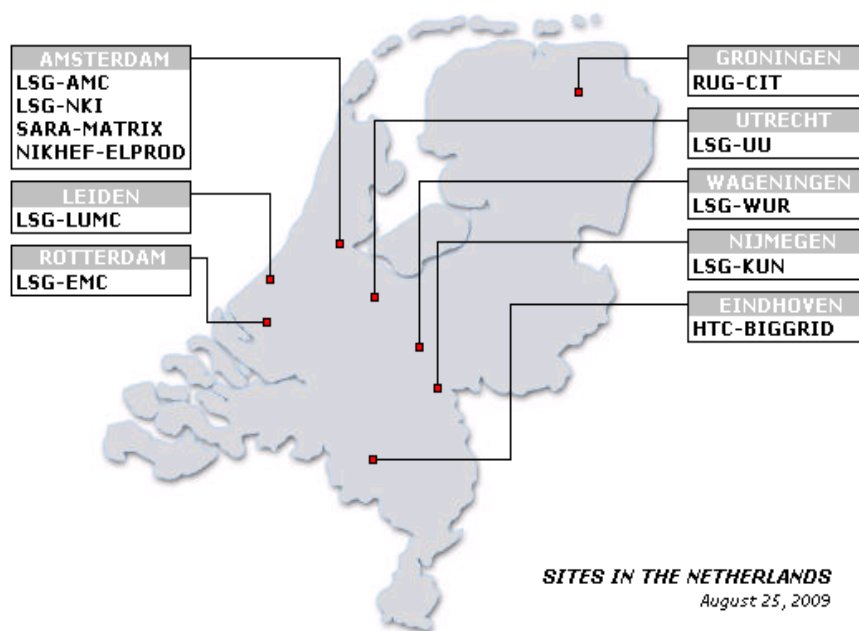
## ABSTRACT

**Motivation:** Over the past few years extensive effort has been undertaken to enable the Life Science community to implement the Grid in their research. Although much progress has been made we observe that Grid is still far from being utilized to its fullest extent. We do, however, believe that the concepts of Grid have great potential for this specific target group. Therefore it is important to summarize the problems on a conceptual level in order to provide a sense of direction to the ongoing international development efforts.

## 1 INTRODUCTION TO THE DUTCH LIFE SCIENCE GRID

All sites with the 'LSG' prefix have been set up in the context of the *Life Science Grid*<sup>1</sup> project. These sites, with the exception of LSG-EMC, have 16 (32 bit) cores and 1.5 Terra Byte of storage. The storage will be upgraded to 20 Terra Byte in the near future. LSG-EMC has 32 cores and 20 Terra Byte storage, and five more sites will be set up in the near future with the same capacity. In addition, all LSG sites will be upgraded to 64 bit architectures. All sites run the gLite 3.1.35 middleware with DPM 3.1.28-0 as Storage Resource Manager (SRM) protocol implementation for the Storage Elements (SE's).

The *Life Science Grid* project is an initiative of SARA Computing and Networking Services<sup>2</sup> and commissioned by the



**Illustration 1:** Sites in The Netherlands accessible for Dutch Life Science VO's

The part of the Dutch Grid accessible to Dutch Life Science Virtual Organizations (VO's) is divided over eleven different sites on eight different locations (see Illustration 1: Sites in The Netherlands accessible for Dutch Life Science VO's).

Netherlands Computer Facilities Foundation<sup>3</sup> (NCF) and the Netherlands Bioinformatics Center<sup>4</sup> (NBIC) in their capacity of

<sup>1</sup> [https://grid.sara.nl/wiki/index.php/Life\\_Science\\_Grid](https://grid.sara.nl/wiki/index.php/Life_Science_Grid)

<sup>2</sup> <http://www.sara.nl>

<sup>3</sup> <http://www.nwo.nl/ncf>

<sup>4</sup> <http://www.nbic.nl>

founding partners in the BigGrid project<sup>5</sup>. SARA, one of the core partners in BigGrid, is responsible for the placing- and administration of the clusters, as well as providing support for their users.

## 2 DIFFICULTIES EXPERIENCED BY THE DUTCH LIFE SCIENTISTS ON THE GRID

The difficulties that the Life Scientists experience when using the Grid are communicated back to us through incident-based notices. Most of these can be filed under one of these three categories:

- **Organizational:** *how does XXX work, where can I find information about it and through what channel can I request support?*
- **Technical:** *my job has successfully generated output but when I try to retrieve it from the SE it does not seem to exist, what is wrong?*
- **Naive:** *I run hundreds of jobs which stage output files of 1gb each in the output sandbox, why do almost all of my jobs fail?*

### 2.1 Naivety-based problems

Over the past two decades the modality of performing Life Science has been changing<sup>6</sup>. Research has started to shift away from hypothesis drivenness and focuses more on bulk-data-generation-and-interpretation drivenness. This shift is where HPC and HTC come into the picture.

However, the shift is not complete. Still today, many Life Scientists are struggling with the new perspective. Fundamental concepts of data-centric research have not yet settled in the field. Because more financial means become available to stimulate the use of HPC and HTC within the Life Sciences, the traditional- and new perspectives are intertwining; but not without the expected difficulties that cause the struggle of the Life Scientist.

A good example is Bioinformatics. Over the past decades the fields of Biology, Informatics and Mathematics, of which Statistics and Probability Theory in particular, have begun to overlap (see Illustration 2: The overlap between Biology, Informatics and Mathematics), and as a result the field of Bioinformatics was created. However, we notice a lack of 'pure' Bioinformaticians – not many know the details of an algorithm, how to implement it in software and interpret the results. Most scientists are located in the field of what we call **BioICT**: they know how to *use* software and interpret results, but have little knowledge of the ins and outs of the algorithm, let alone of its implementation. Hence, the requirements for data-centric Biological research are not completely met. Although BioICT is a logical first step for the scientist it still requires close cooperation with BioStatistics and Statistical Informatics.


While the Life Scientist tries to find his way in Computer Science he runs in to the problem of dealing with- and management of large scale distributed systems. For example, when running a software like BLAST<sup>7</sup> (**B**asic **L**ocal **A**lignment **S**earch **T**ool) on a desktop computer, the researcher specifies his input data, runs the program and when it finishes, his output data is

available directly. When utilizing the Grid this work flow is less intuitive (from his perspective, at least). “*Your job is not being submitted because you have not delegated your proxy*”, “*You should stage your output data on an SE because the WMS cannot handle that much output data in an output sandbox*”, “*Your output datasets are too many and too small, and cannot be handled by the SE*”, and other similar hints and tips keep recurring when supporting the scientist.

Since these problems have to do with the transition that is taking place in the field, they are hard to solve. But the answer does not end there. From our role as Computer Scientists and System Engineers we have a responsibility to ease the transition – if not from us, from who can a scientist learn to handle the computing facilities?

In practice these problems manifest in the following areas:

- Using the command line interface
- Job planning and management
- Data management



**Illustration 2:** The overlap between Biology, Informatics and Mathematics

### 2.2 Technology-based problems

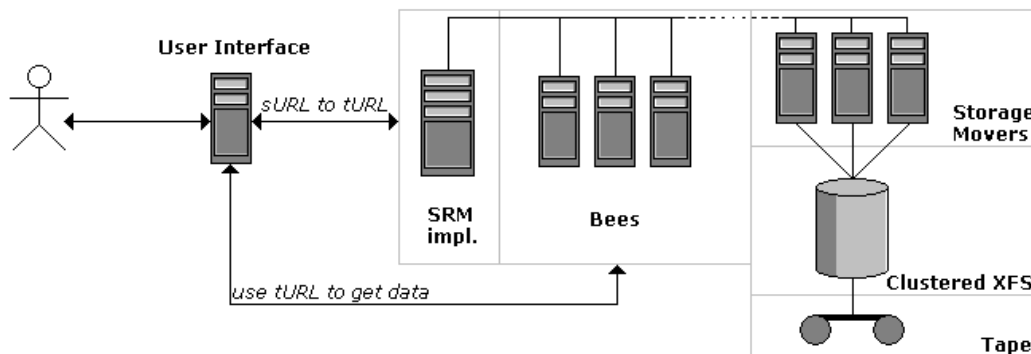
Other problems originate from the technology rather than the scientist. The Grid, although continuously maturing, is still not a production infrastructure. Further more, it is based on academic software, which is not always as stable as we would like.

One of the general issues is the trade-off that exists between the scale and the stability of the Grid. The growth of the Grid (in systems as well as people) is parallel to the growth of the surface on which errors can occur. We notice that the occurrence of errors as a result of the scale are not just conceptual; often when an update or upgrade of some component is done, an other component, an interface or a client breaks in one way or another.

<sup>5</sup> <http://www.biggrid.nl/>

<sup>6</sup> Gusfield, D. (2002) [http://webcast.ucdavis.edu/Engineering/2008/EC\\_S124\\_02/ECS124\\_4-1-02\\_L-1.aspx](http://webcast.ucdavis.edu/Engineering/2008/EC_S124_02/ECS124_4-1-02_L-1.aspx)

<sup>7</sup> <http://www.ncbi.nlm.nih.gov/BLAST/>



**Illustration 3:** The concept of Storage Elements

There is another dimension that requires attention. The major early adopter (if not initiator) of Grid technology was the High Energy Physics (HEP) community. As such it contributed much to the current state of the Grid; it defined its needs ever more clearly over the years and influenced development of several systems that are now common ground on the Grid.

A common example is data management. The concept of Grid data management in The Netherlands is based on the SRM protocol. An implementation of this protocol is at the heart of an SE (in The Netherlands this implementation is usually either *dCache*<sup>8</sup> or *Disk Pool Manager*<sup>9</sup> (DPM)). For an overview of a common Data Management solution see Illustration 3: The concept of Storage Elements.

This set-up is a fairly general one for distributed data storage and scales fairly well. However, the SRM protocol has three downsides from the perspective of the Life Scientist:

*First*, each transfer has major overhead because of the scope of the protocol. It is known to be very inefficient when transferring many small files, which is a common use-case within the Life Sciences, in contrast to HEP.

*Second*, the protocol is not common. It has no general purpose clients available and the storage is not locally mountable. Further more, it has only one official client interface, which is based on the command line.

*Third*, the implementations of the protocol are academic software. Experience learns that they are not as stable as we would like them to be. This is a problem for small communities – while the bigger communities like HEP are able to increase redundancy by replicating their data over multiple countries, the Life Science VO's are much more dependent on the local sites.

Without going into further detail by comparing the computing needs of HEP and the Life Sciences, it is important to realize that some of the current solutions do not fit the normal Life Science use-case. Awareness of this issue is still young and was raised mainly because the method of *technology* push, as applied so far, was not successful. The signals coming from the field strongly

indicate the need of the scientist for more influence on the state of the Grid, or in other words, the need for a *technology pull*.

Most technical problems we see are related to one of the following:

- Authentication
- Job management and failure rates
- Data management

### 2.3 Organization-based problems

After the two previous, external factors, we need to turn around and look at ourselves. One of the issues causing difficulties is based in the horizontal way that support and documentation is provided – no distinction is being made between the different types of scientists or VO's. A single type and version of documentation and education is provided for all. However, scientists can differ in many ways from each other, from expertise and frame of reference to substance of research and needs for computing. Although the situation in The Netherlands might differ in other countries, it remains crucial that support is provided in a way that fits the scientist and his community.

Issues with an organizational background are related to one of the following:

- Starting with the Grid
- The role of the community or VO

## 3 MAPPING THE LIFE SCIENCES ONTO THE GRID AND VISE VERSA

What steps can we take to map the Life Scientist onto the Grid and the Grid onto the Life Scientist? The remainder of this paper is an attempt to structure, and give context to a number of possible solutions.

### 3.1 Easing the transition between research modalities

Based on the shift of research modalities as mentioned in the previous chapter, we can define the technological factor of the shift as one of “desktop to large scale”. So an obvious (but not concrete) suggestion is to introduce the technology from the Desktop-based perspective of the Life Scientist – letting him use the Grid from his local Desktop computer.

Many attempts have been made to provide Desktop access to the Grid. The mission is, summarized, *to let a scientist use large scale*

<sup>8</sup> <http://www.dcache.org/>

<sup>9</sup> <https://twiki.cern.ch/twiki/bin/view/LCG/DataManagementDocumentation>

distributed computing and storage resources from a familiar environment. Examples of such initiatives include:

- *VBrowser*<sup>10</sup>, an attempt to provide access to storage services on the Grid through an interface that resembles a standard file-browser;
- *inQ*<sup>11</sup>, a browser tool for SRB;
- *jGridStart*<sup>12</sup>, a Java WebStart application to handle certificate requests and to import a certificate into your browser (currently in alpha);
- *GridApps*<sup>13</sup>, a REST (Representational State Transfer) (Fielding, R.T. 2000) based interface for applications running on the Grid;
- *Leiden Grid Infrastructure*<sup>14</sup> (LGI), similar to GridApps.

We believe that such low level interfaces offer real added value to Life Scientists. They allow the scientist to, instead of dealing with the Grid itself, do a familiar task from a familiar environment. Of course, such interfaces do have major dependencies, and an update or upgrade of one of the many components can (and will) break some of them.

Attempts of a different kind to provide Grid interfaces are workflow systems. These systems, such as *Taverna*<sup>15</sup> and *Moteur*<sup>16</sup>, allow scientists to define a workflow of Grid jobs. Although these applications are mostly Desktop based (or sometimes on-line, such as the *P-GRADE Portal*<sup>17</sup>), their principle is different from that of the preceding. They often replace the complexity of the Grid with a type of complexity that approaches large scale computing from a perspective that might be closer to the mindset of the scientist.

We believe that such platforms provide much added value for many different use-cases, but question its use for general application within the Life Sciences. We notice that Life Science research groups that utilize such software need to go through a learning curve of which the steepness is not proportional to the increase of productivity it delivers. Even though Life Scientists do have a need for *simple* workflows, these workflows are often the same or very similar across their domain. Therefore it should be possible to take a simpler, but maybe less generic, approach.

An example of such an approach is annotated application services. Providing access to applications running on the Grid in the way GridApps does is a good first step. However, after collecting data, the Life Scientist typically needs to do more than a single atomic step (as provided by GridApps) to get useful information from his data. It is common that the Life Scientist needs the Grid to pre-process, process and visualize or interpret his data. Using GridApps he can do all of these steps, providing that the applications he needs are available. It would be more useful, however, if he could specify all his actions at once. That way he would not have to deal with moving of data, monitoring of his jobs, etcetera.

<sup>10</sup> [http://staff.science.uva.nl/~ptdeboer/vlet/page\\_vbrowser.html](http://staff.science.uva.nl/~ptdeboer/vlet/page_vbrowser.html)

<sup>11</sup> <http://www.sdsc.edu/srb/index.php/InQ>

<sup>12</sup> <http://www.nikhef.nl/pub/projects/grid/gridwiki/index.php/JGridstart>

<sup>13</sup> <https://ws2.grid.sara.nl/apps/>

<sup>14</sup> <http://fwnc7003.leidenuniv.nl/LGI/>

<sup>15</sup> <http://taverna.sourceforge.net/>

<sup>16</sup> <http://modalis.polytech.unice.fr/software/moteur/start>

<sup>17</sup> <http://portal.p-grade.hu/>

Providing annotations for application services on the Grid would be a major step. (See Illustration 4: Tying Application Services.) These annotations should specify input parameters and the type of output generated. By checking them it would be possible to see which application can take the output of another application as its input. As an additional step, a client application could be developed that checks for available application services and allows the user to select which applications to run in which order.




Illustration 4: Tying Application Services

Another major issue is that of data management. From his old-modality frame of mind, the Life Scientist has no conception of the complexity of data management on the Grid. Concepts like replication or checksums do not come natural, since in his Desktop environment they are irrelevant.

An important issue regarding data management is that of its current unreliability when storing many smaller files at once. This is due to technical issues (for which we will make a suggestion in the following chapter), and the available workarounds are hard for the scientist to handle.

We believe that the concepts of data management are complicated, but knowledge of them is essential when working with the Grid. Therefore, as long as no alternative is available, our suggestion is to generate more teaching material on how data management works and how it should be used, specifically by the Life Scientists.

Another option is to develop an interface based on common and familiar protocols, such as *WebDAV*<sup>18</sup>, while maintaining the current technologies. The advantage of such protocols is that the remote file system can be mounted easily, because most operating systems provide native support. However, the implementations of these protocols on top of the existing infrastructure might prove problematic.

### 3.2 Adapting the Grid to suit the Life Sciences

Apart from providing interfaces that fit the Life Scientist, much work can be done to improve technical concepts of the Grid. Two major issues can be identified. Respectively:

- Data Management
- Job failure rates

The technical issues with Data Management on the Grid are well known and documented. Although the concepts of distributed storage as used on the Grid has been proven to be successful in other domains, on the Grid it is to a lesser degree, for reasons that have been discussed in the previous chapter.

Based on the identified problems with the SRM protocol, which in our opinion are the cause for most of the technical issues with

<sup>18</sup> <http://www.webdav.org/>

Data Management on the Grid, we propose to implement a different protocol. During selection the following criteria should be maintained:

*More efficient for smaller files.* Since SRM is not able to deal with many small files, the next protocol should be more efficient.

*Common for enabling access to distributed storage.* If we choose a protocol that is widespread among different domains, that means many clients and tools are available.

*Not academic.* By choosing a protocol that goes beyond the academic perspective, we can define a better framework for our services. In the end, we as supporters, should provide a service to our customers, the Life Scientist (among others). Since not-academic software is much easier to rely on, we are able to go as far as defining a Service Level Agreement (SLA).

The issue of job failure rates asks for a different approach. There is no widely available alternative with the promise of better performance.

Since most of the errors occur between the a job submission and the point it lands on a Worker Node, effort needs to be spent on more actively advocating the use of Pilot Jobs. These are Grid jobs that run on a *meta-level*; their main purpose is to get onto a Worker Node and check the environment. After this is done, it can fetch a task definition from an external source and start processing this task. If all goes well, it can store the results of the task, delete the task definition on the external source, and fetch a new task (also see Illustration 5: Pilot Jobs).

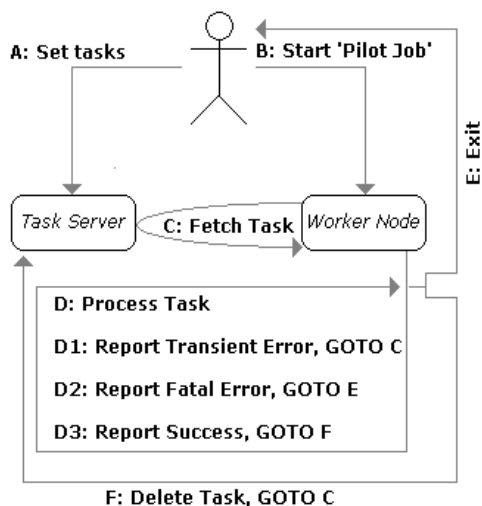


Illustration 5: Pilot Jobs

There are some frameworks available that support this concept.

- *Token Pool Server*<sup>19</sup> (ToPoS). An extremely efficient, open source, REST based interface to store task definitions;
- *Distributed Analysis Environment*<sup>20</sup> (DIANE). Provides automatic control and scheduling of computations, and is part of the EGEE respect suit;

<sup>19</sup> <http://topos.grid.sara.nl/4/>

<sup>20</sup> <http://it-proj-diane.web.cern.ch/it-proj-diane/>

- *Condor*<sup>21</sup> GlideinWMS (Sfiligoi, I. 2007). Provides a similar approach, but then from a Workload Management System (WMS) perspective.

### 3.3 Organizational suggestions

Enabling the the Life Science researcher to use the Dutch Grid infrastructure asks more than technical adjustments and additions. It is important that the dissemination of information is handled in a way that comes natural to the Life Scientist, to provide him with a platform that has information dedicated to his specific problems, and to provide clear channels for support. The details of such organization are expected to be more of an internal issue, and therefore out of the scope of this discussion. We do give an overview of a situation that seems a promising improvement.

We propose to introduce VO specific support and documentation, in which the VO manager plays an active and crucial role. Since a VO is build around a notion of similarity of its members, we can assume that the Grid is used in similar ways to achieve similar goals. Because these ways are directly related to the research, it makes sense to enforce the generation of documentation on tips and tools. Sharing is the keyword here.

But what means do we have to enforce this? Since currently all support is delivered from a central place, we can distribute responsibilities and introduce a new first-line support, dedicated to a specific VO. Since a VO does not have a dedicated person for support, this channel needs to be maintained by the VO members. The perfect tool for this is a 'many-to-many' mailing list as applied by many successful open-source projects. A search-able archive needs to be maintained, making up for documentation. Of course there should be a possibility to escalate a question. For this we suggest to keep on using a ticket system.

Further more, an on-line repository needs to be maintained, in which static documentation is stored. This repository needs to include all VO specific documentation as well as links to generic documentation. Another useful addition is a search-able database of public papers, created in the context of this VO.

## 4 DISCUSSION

The Grid has the potential to play an important role in data-centric Life Science research. There are, however, still many problems to be solved. These problems are on three distinct levels; organization, technological and naivety. Currently a solution is mostly being looked for on a technical level. However, it is our believe that the solution is only to be found by keeping the bigger picture in mind. We need to build a realization about the identity of the Life Scientist and his research, and try to find solutions that fit in his modality of research, mindset and capabilities.

Of course many technical improvements can be made as well. An important one, from the Life Scientist's perspective, is Data Management. The current solution has not proven suitable for the Life Sciences and effort needs to be put into finding a new solution. Again, it is important to match this solution to a number of criteria to make sure it is an actual workable solution.

On an organizational level it is important to build stronger Virtual Organizations that have a sense of self-preservation. This means that knowledge specific to a VO should be contained and

<sup>21</sup> <http://www.uscms.org/SoftwareComputing/Grid/WMS/glideinWMS/>

disseminated within the VO itself. To suit this set-up we need to reorganize our knowledge dissemination and support.

Further discussion and analysis is needed. It is important to be aware that not only the Life Scientist needs to change to suit the technology, we as Grid supporters also need to facilitate change to suit the Life Scientist. Lets move on to find the right balance between technology push and technology pull, to teach both ourselves and the Life Scientist and to find forms of organization that fit the natural form of the Life Science community.

## ACKNOWLEDGEMENTS

The observations and corresponding suggestions in this paper were obtained through, apart from own experiences, many discussions with colleagues in the e-Science Support Group of the department of High Performance Computing and Visualization. The paper has been written with support from, and in the context of, the International Workshop on Portals for the Life Sciences (IWPLS) '09.

## REFERENCES

- Fielding, R.T. (2000) Representational State Transfer (REST) *Architectural Styles and the Design of Network-based Software Architectures*, University of California, Irvine (CA), pp. 94-124
- Sfiligoi, I. (2007) *Journal of Physics: Conference Series glideinWMS – A generic pilot-based Workload Management System*, 119 2-5