

Using the Grid for Satellite Imagery with UNOSAT

Sean Moran
University of Cambridge
St Catharine's College
Cambridge CB2 1RL, UK
sean.moran@cantab.net

Patricia Méndez Lorenzo
European Laboratory for Particle Physics
1211 Geneva 23
Switzerland
patricia.mendez@cern.ch

Abstract

This paper describes work which investigates how the Large Hadron Collider (LHC) Computing Grid (LCG) can be used to upload, store, compress and download satellite image data files on behalf of the United Nations Organization UNOSAT. Mechanisms for satellite image upload, storage and download are outlined. LCG computational resources are exercised by a JPEG2000 compression framework that transforms UNOSAT's original GeoTiff images to JPEG2000 ECW files. It is shown how the availability of these functions will aid humanitarian relief efforts by enabling UNOSAT to vastly increase its portfolio of satellite images and to supply images to fieldworkers much more rapidly through pre-compressed files. Importantly, these benefits arise through zero financial cost to UNOSAT.

1 Introduction and Related Work

1.1 UNOSAT

UNOSAT is a United Nations initiative to provide the humanitarian community with access to satellite imagery and Geographic Information System (GIS) services. UNOSAT is implemented by the UN Institute for Training and Research (UNITAR) and managed by the UN Office for Project Services (UNOPS). In addition, partners from public and private organizations constitute the UNOSAT consortium.

The goal of UNOSAT is to make satellite imagery and geographic information easily accessible to the humanitarian community and to experts worldwide working to reduce disasters and plan sustainable development. To do this UNOSAT acquires satellite images from all commercial providers.

UNOSAT is active on a wide range of projects and supports UN agencies and NGOs in many fields, from emergency relief, to refugee and IDPs to disaster prevention and longer-term

development. Some current and past projects of UNOSAT include Indian Ocean Tsunami Relief and Development (ongoing) (see Figure 1), Risk Assessment after the Matagalpa floods (finalized December 2002) and post-conflict assessment maps of Afghanistan (finalized December 2002).

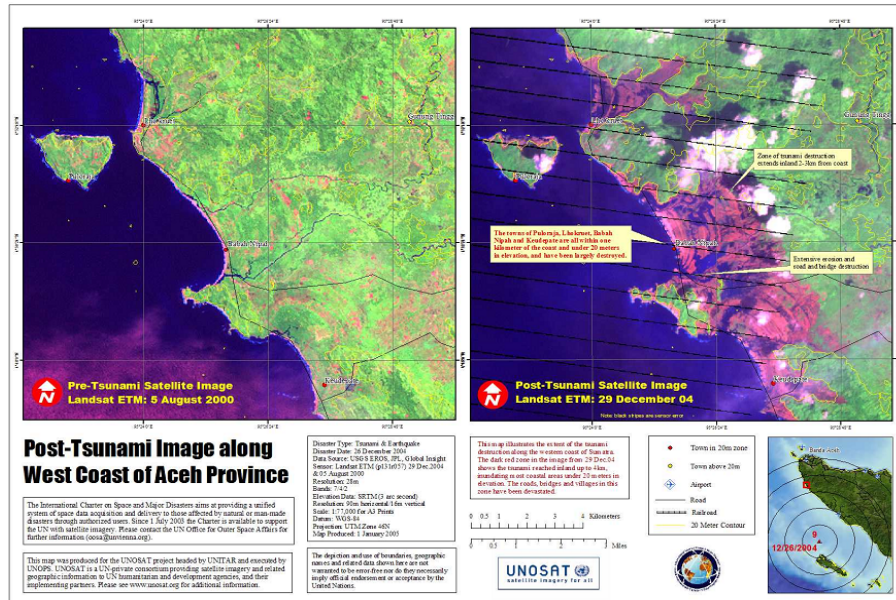


Figure 1: A map created by UNOSAT in response to the Indian Ocean Tsunami disaster. **Left:** Western coast of Sumatra before the Tsunami. **Right:** After the Tsunami. Dark red zone shows the Tsunami reached inland 4km inundating most coastal areas under 20m of elevation. The roads, bridges and villages in this zone were completely devastated. Source: UNOSAT

The UNOSAT service chain is represented in Figure 2. UNOSAT would typically respond to a disaster, such as the Matagalpa flood, by requesting detailed satellite maps of the area. Satellites travelling in orbit between 650 and 850 km above the Earth respond to this request and map out the region of using sophisticated radar and optical techniques. The optical images show the topography of the countryside, while radar reflects very differently from water and land, and thus shows up the flooded areas clearly.

The satellites then send the raw data (telemetry) back to Earth. When the Ground Station receives the raw data from the satellite, it performs some advance data processing on the data which yields several satellite images, each roughly 400 MB in size. These image files are then relayed by the Data Providers to UNOSAT. UNOSAT carries out some additional processing on the images in order to tailor them to the exact requirements of Humanitarian Workers in the field. The images are then placed on UNOSAT's website ready for end users to download over the Internet.




Figure 2: *The UNOSAT service chain, from image procurement to delivery to Humanitarian workers in the field.* Source: UNOSAT

1.2 UNOSAT and CERN

UNOSAT has been at CERN since 2002 in order to benefit from CERN's excellent IT infrastructure and network connectivity. An example of the benefit UNOSAT has found in this collaboration is given by the Indian Ocean Tsunami disaster. When the Tsunami struck on the 26th December 2004, UNOSAT immediately created maps of those areas thought to be most heavily affected by the wave. These maps were made available through the UNOSAT website a few days later on the 29th December.

Humanitarian organizations as well as the international media immediately started to make heavy use of this information to plan their activities. This put the Central Web Services at CERN under considerable strain, resulting in memory capacity problems. A fix was quickly found by the Internet Services Group, which effectively isolated the UNOSAT website from other web services. As a result, and as CERN's computer centre is manned around the clock, the UNOSAT data remained available despite the growing number of requests.

The UNOSAT director, Alain Retière, is confident that collaborating with CERN on advanced IT solutions can help facilitate the job of humanitarian organizations in responding to future natural disasters and complex crisis situations. Mr Retière considers that Grid computing, which is being pioneered by the particle physics community, could prove particularly useful in handling the large amounts of data and significant data processing that UNOSAT's work involves.

1.3 The LHC Computing Grid

The LHC Computing Grid (LCG) [1] is an emerging computing infrastructure that aims to provide scientists with seamless access to processing power and data storage distributed around the Globe¹. The LCG was initiated in order to provide the computing infrastructure for the simulation, processing and analysis of the data of the Large Hadron Collider (LHC) experiments.




Figure 3: The four experiments of the CERN LHC project. It is expected that the experiments will require a raw recording rate 0.1 - 1 GBytes/sec accumulating at 15 PetaBytes/year and processing equivalent to 100,000 of today's fastest PCs. Source: CERN

The LHC, currently being built at CERN near Geneva, is the largest scientific instrument in the world. When it begins operations in 2007, the various experiments (ALICE, ATLAS, CMS, LHCb) that use the collider will generate enormous amounts of data², which thousands of scientists around the world will access and analyse.

The LCG computing service is implemented as a geographically distributed computational data grid. The service uses computing and storage resources installed at a large number of computing sites in many different countries all interconnected by fast networks. Figure 4 provides an overview of the geographical distribution of LCG resources. The sites participating in the project are primarily universities and research laboratories.

The data from the LHC experiments will be distributed, according to a four tiered model (see Figure 5). A primary backup will be recorded on tape at CERN, the "Tier-0" centre of LCG. After initial processing, this data will be distributed to a series of Tier-1 centres, large computer centres with sufficient storage capacity for a large fraction of the data, and with round-the-clock support for the Grid.

The Tier-1 centres will make data available to Tier-2 centres, each consisting of one or sev-

¹The terms LCG and Grid will be used interchangeably in this paper.

²Estimates put this at roughly 15 Petabytes (15 million Gigabytes) of data annually.

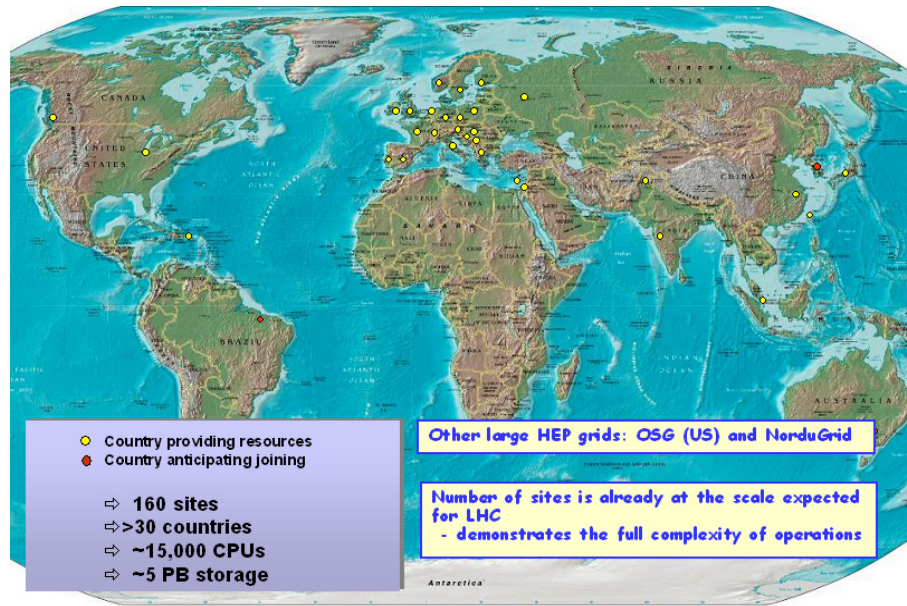


Figure 4: *Geographical distribution of LCG resources. The LCG currently commands resources from over 160 sites, located in over 30 countries worldwide which provide roughly 15,000 CPUs and 9Pb of data storage. Source: CERN*

eral collaborating computing facilities, which can store sufficient data and provide adequate computing power for specific analysis tasks. Individual scientists will access these facilities through Tier-3 computing resources, which can consist of local clusters in a University Department or even individual PCs, and which may be allocated to LCG on a regular basis.

The LCG middleware [2] hides much of this complexity from the user, giving the impression that all these resources are available in a coherent virtual computer centre.

1.4 LCG Operation

Figure 6 provides a high-level overview of the LCG in operation. The process begins with an individual user logging into a Grid User Interface (UI)³. A UI is a machine where LCG users have a personal account and where the user security certificate is installed. This is essentially the gateway to Grid services.

The user describes her job using a JDL (job description language) file which identifies the executable that will run on the Grid, the input files required for the executable and the names of the standard output and standard error files that will be produced when the job is executed. An example of a typical LCG JDL file is:

```
Executable = "upload.pl";
```

³In CERN this is the LXPLUS system running on Scientific Linux.

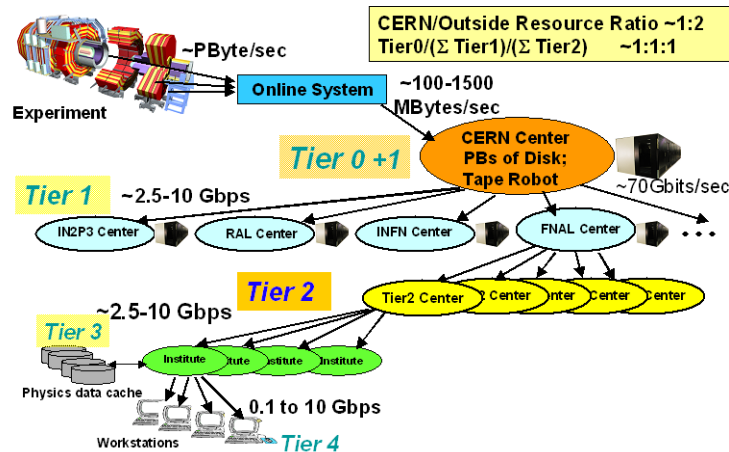


Figure 5: The LCG Service Hierarchy. The data in the LCG is distributed according to a 4-tiered model. Source: CERN

```
StdOutput = "message.txt";
StdError = "stderr";
InputSandbox = {"metadata_latlon.txt"};
OutputSandbox = {"message.txt", "stderr"};
```

The job will arrive at a Resource Broker (RB). A set of services running on the RB machine contribute to match job requirements to the available resources, schedule the job for execution to an appropriate Computing Element (CE), track the job status, and allow users to retrieve job output. Each CE is the front-end to a local batch system managing a pool of Worker Nodes (WN) where the job is eventually executed. The Logging and Bookkeeping (LB) service keeps the state information of a job and allows the user to query its status.

The RB uses an Information System (IS) to determine the most appropriate CE matching the requirements of the job as specified in the JDL file. The IS provides information about the Grid resources and their status, and this information is published following a specific schema referred to as Glue [3].

The output of the user's job is stored on a Grid Storage Element (SE). The output may be stored in any SE on the Grid and not just simply the SE at the site where the job is executed. The Grid Data Management Service allows users to move files in and out of the Grid⁴, to replicate them among different Storage Elements (SE), and to locate the files. A central file Catalog, the LFC (LCG File Catalog) keeps information about file location and associated metadata.

⁴A number of available and supported protocols (Globus, GridFTP) are used for data transfer.

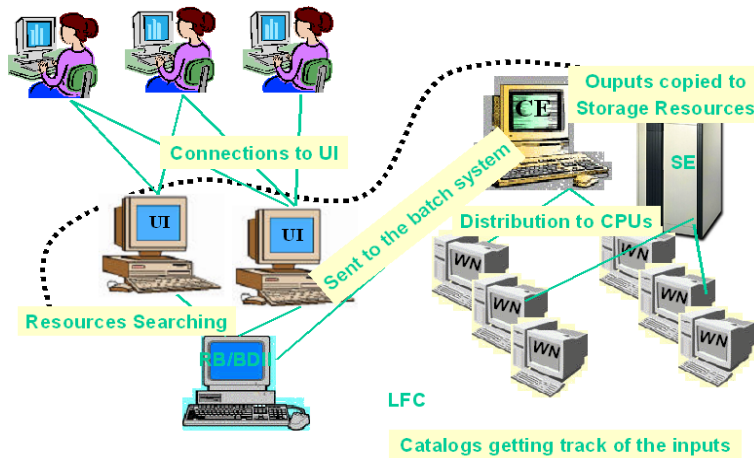


Figure 6: *The LCG in Operation.* A user logs into a User Interface (UI) which allows her to submit a job to be executed on Grid Resources. A Resource Broker (RB) uses the data provided by an Information System (IS) to send the job for execution on a particular Computing Element. The outputs of the job are stored on a Storage Element (SE) and the files are tracked on the Grid by using a File Catalog (LFC).

1.5 Satellite Imagery and the Grid

Despite the excellent support from CERN, the operations of UNOSAT are currently hampered by a low budget, modest human resources⁵ and limited computing power in terms of processing and storage.

The current UNOSAT computing infrastructure (see Figure 7) is supported by CERN and consists of Cernunosat01, a server with capacity of 1.5 Tb and Cernunosat02, a server with capacity of 782 Gb. Both servers are hosted by CERN and run on Intel Xeon processors. UNOSAT also commands a directory totalling 334Gb on the CERN Distributed File System (DFS).

This computing infrastructure is increasingly becoming a bottleneck in the UNOSAT operation due to limited processing and storage capacity. The total amount of data UNOSAT is storing at CERN is relatively small at present, about 2 terabytes (Tb), but this is growing rapidly. NASA⁶, the European Space Agency, the Centre National d'Etudes Spatiales and the Indian Space Research Organisation among others are giving UNOSAT access to archives containing terabytes of satellite images, and these need to be both processed and stored by UNOSAT, whilst giving the humanitarian easy access to the files.

Furthermore, there are plans to launch multiple high resolution satellites into orbit, which could eventually provide UNOSAT with some very large high resolution images. Commercial companies are also planning to release surveillance drones equipped with arrays of digital

⁵12 people in their CERN Headquarters including UN fieldworkers, satellite imagery experts, geographers, database programmers and internet communication specialists.

⁶NASA has recently donated over 20Tb of satellite images to UNOSAT.

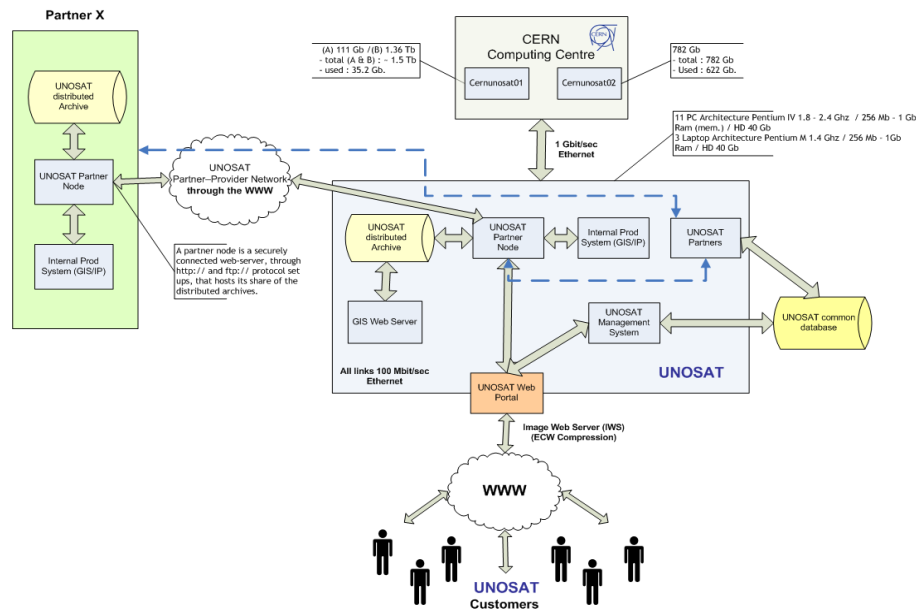


Figure 7: *UNOSAT Computing Infrastructure as of July 2005. CERN supports UNOSAT by providing two servers cernunosat01 and cernunosat02 that provide around 2Tb of storage space and data processing by Intel Xeon processors.*

cameras⁷. These data sources could one day produce petabytes of data, a large proportion of which could be processed and stored by UNOSAT.

The LCG may present a solution to this problem. For example, Satellite images could be uploaded to the Grid where they could be stored and processed, freeing up the UNOSAT local computing resources. The stored files could also be compressed on the Grid. These compressed images could then be downloaded, when required, by humanitarian workers in the field. As Bjorgo explains [4], timeliness is critical for the application of satellite imagery in disaster recovery and management. Pre-compressed satellite images will therefore be of great importance in enabling fieldworkers to rapidly download those images over slow internet connections.

The work described in this paper has the aim of assessing whether the Grid can meet the needs of UNOSAT by providing the capability to upload, store, compress and download satellite images. A skeleton framework implementing these various functions was created and executed on Grid resources in order to achieve this objective. The rest of this paper describes this framework and the results obtained.

⁷See BBC News Online: "Eternal planes' to watch over us", Jo Twist, August 2005

1.6 Related Work

The Grid has been applied to many different problems in a wide range of fields. Most directly related to UNOSAT is the research mentioned in the paper [5], that investigates how the Grid can be used for distributed Geo-rectification of Satellite Images. The computational tasks described in the paper include, wrapping satellite positional data to compensate the earth curvature, and consist of several steps such as image re-sampling, resolution conversion and image matching.

Research has also been conducted into how the Grid can be used for Medical applications (imaging, diagnosis and treatment) [6][7], Bioinformatics (study of the human genome and proteome to understand genetic diseases) [8][9], Nanotechnology (design of new materials from the molecular scale) [10], Engineering (design optimization, simulation, failure analysis and remote Instrument access and control) [11] and Natural Resources and the Environment (weather forecasting, Modelling and prediction of complex systems) [5].

2 UNOSAT Grid Infrastructure

Prior to developing software to upload and process UNOSAT's satellite images, it was firstly necessary to carry out a number of pre-development steps that ensured the necessary infrastructure was available to run UNOSAT's jobs on the Grid.

A Virtual Organization (VO) was created for UNOSAT on the Grid. A VO is an entity, which corresponds typically to a particular organization or group of people in the real world. The membership of a VO grants specific privileges to the user. For example, a user belonging to the UNOSAT VO will be able to read the UNOSAT files or to exploit resources reserved for UNOSAT.

Storage capacity of 3.5Tb was negotiated for UNOSAT on the CASTOR⁸ storage system. The CASTOR system now handles the bulk of the physics data of CERN with a total volume stored on tape of more than 1 Petabyte. As well as storage space on CASTOR, sufficient Computing Elements and Resource Brokers on the Grid have been reserved for the use of UNOSAT.

The authors collaborated with the CERN ARDA⁹ group in order to create a program to convert UNOSAT's metadata¹⁰ to the physical location of the files on the Grid. This facility was currently lacking on the Grid, and was required for UNOSAT to download their satellite images by entering the coordinates of the desired images (see Section 2.4).

It was also necessary to create UNOSAT an AFS area of 5GB. This is used as a temporary storage space for a selection of those satellite images that UNOSAT wishes to upload (see Section 2.1). This was an initial, less than ideal, solution to the problem of how to transfer files from the UNOSAT server to a Grid UI. In the future, this mechanism will be replaced by direct access from the UNOSAT server to the LCG.

⁸CASTOR CERN Advanced STORage manager.

⁹ARDA, which stands for "A Realization of Distributed Analysis of LHC", develops prototype systems that integrate Grid middleware developed by the EGEE (Enabling Grids for E-science in Europe) project with common analysis tools in high-energy physics and the analysis software of the LHC experiments.

¹⁰Specifically this was the geographic coordinates of the satellite imagery.

2.1 Image Upload to the Grid

A Perl script, `upload.pl`, was created in order to automate the process of uploading UNOSAT's satellite images to the Grid resources. To upload satellite images it is firstly necessary for UNOSAT to place the desired satellite images into the temporary CERN AFS directory (5GB) created for this purpose.

UNOSAT will also need to provide a metadata file containing details about each satellite image they wish to upload. Specifically these were the *file name*, *directory path of the files on the UNOSAT server* and the *geographical coordinates*. The Metadata file is required to be in Comma Separated Variable (CSV) format in order to simplify the parsing of the file ¹¹. The Perl script is then executed on a UI and carries out the following four steps:

1. Parse Metadata file to extract File Names, Coordinates, Path.
2. Locate the files (using the extracted file names) to be uploaded in the temporary AFS directory.
3. Copy and Register the files (using file names and paths) to Grid Resources.

This program is currently ready to interoperate with the ARDA Application (see Section 2.4). Testing has been performed using several satellite images from UNOSAT. At the time of writing plans are underway to scale up the testing and upload over 400 satellite images to the 3.5Tb CASTOR storage space.

2.2 Satellite Image Compression

As mentioned earlier, having pre-compressed satellite images is crucial in order to enable humanitarian workers to download large satellite images over low bandwidth links in developing countries.

To compress UNOSAT's satellite images, the authors applied the open source GDAL Utilities¹², which is a translator library for raster geospatial data formats that is released under an X/MIT style Open Source license. To compress the images the tool, *gdal_translate*, was used which took as input a GeoTiff¹³ and compressed this image to a JPEG2000 ECW file¹⁴.

GDAL utilities supports reading and writing JPEG2000 files using the ECW SDK from ERMapper. Coordinate system and georeferencing transformations are read, and some degree of support is included for GeoJP2 (GeoTIFF-in-JPEG2000), ERMapper GML-in-JPEG2000, and the new GML-in-JPEG2000 specification developed at OGC.

¹¹The original UNOSAT metadata file is a Microsoft Excel file. To convert this to the required CSV format, UNOSAT simply has to save the file in CSV format from the Excel File menu.

¹²These utilities are part of the FW Tools package which is set of Open Source GIS binaries for Windows (win32) and Linux (x86) systems. <http://fwtools.maptools.org/>

¹³This is a standard TIFF file with geographic information hidden in the file. It is utilized by a variety of industries for digital handling of geographic data.

¹⁴This compressed file *cannot* be read with ERMapper 6.4 installed in UNOSAT. However, ERMapper 7.0 is capable of reading the compressed files, and there are now plans in UNOSAT to update their version of ERMapper to achieve compatibility. In ERMapper 7.0, in order to read the compressed files, the file extension `.jp2ecw` simply needs to be renamed to `.ecw`

The performance of this compression was measured for a sample selection of UNOSAT's satellite images. For example, a 200 Mb GeoTiff was losslessly compressed into a 50Mb file in 56 seconds and a 60 Mb GeoTiff was losslessly compressed into 5Mb in 43 seconds.

2.3 Job Submission, Monitoring and Output

A Perl Script, `sw_unosat.pl` was created in order to send the `compression.pl` script to the Grid Computational Resources for execution. This script is submitted to the Grid, and carried out the following steps:

1. Locate the images, which have previously copied and registered to the Grid.
2. Download the images and unzip the files.
3. Compress the files using `gdal_translate`.
4. Re-zip the compressed images and associated files and copy and register back to Grid.

A Perl script, `submitter_general.pl` is used to submit `sw_unosat.pl` back to the Grid. `Submitter_general.pl` is an automatic toolkit delivered to submit large bunches of jobs to the Grid. It was developed for the Geant4 collaboration¹⁵ and now is used by any community arriving to the Grid and willing to run in a short time. This script has been specially adapted by the authors to conform to the UNOSAT requirements.

Another Perl script `get_output.pl` keeps track of the status of the jobs submitted previously with the `submitter_general.pl` script. A typical output of this script is:

```
The job run in grid01.phy.ncu.edu.tw:2119/jobmanager-torque-dteam is
in status: running
```

This output demonstrates that the job is present on the `grid01.phy.ncu.edu.tw` Computing Element and is currently in the running status.

At the present time the compression code required to execute the job is installed as and when needed on whatever Computational Site the job may be located. `Sw_unosat.pl` is effectively transferring the satellite files to the site and setting the environment for its execution, site per site.

2.4 Image Download from the Grid

The CERN ARDA group developed a program to map the metadata of UNOSAT's satellite images to the Logical File Names of the Satellite Imagery. A Logical File Name (LFN) is the unique name of a file in an abstract "grid space", for example,

```
/grid/dteam/users/graeme/dm_test/foobar
```

This name is independent of any Storage Url (SURL), with the LCG File Catalog (LFC) mapping these LFNs to the Physical Location of the files (or SURL's). In LCG each VO will have a separate file catalog, so their LFN spaces will be independent of one another.




Figure 8: *Downloading images from the Grid. To download images a user enters the geographical coordinates of the desired images. The ARDA application converts this metadata into the logical file names (LFNs) of the satellite images. The LFNs are sent the LCG File Catalog (LCG) which provides the Physical File Names (PFNs). These PFNs are then used by the CASTOR system to retrieve the desired files from the Grid.*

With this program it was then possible to create a Perl script to download images (see Figure 8) from the Grid based on geographical coordinates entered by a user.

The user provides the geographical coordinates of the images that they wish to retrieve and these coordinates are input into the ARDA application. The corresponding LFN's are determined and sent to the LCG File Catalog (LFC) which maps them to the actual physical location of the files on the the CASTOR storage system. These Physical File Names (PFN's) are sent to the CASTOR system which locates the required files and transfers them across to UNOSAT.

3 System Summary

Figure 9 provides an overview of the UNOSAT Grid system. To illustrate the system from the perspective of the user we will consider the example of a member of the UNOSAT team wishing to upload and compress satellite images.

The user will log into a UI and transfer the satellite images to the temporary AFS area. Having provided the JDL and metadata file the user would run the `upload.pl` script on the UI. The satellite images are then copied and registered onto the Grid resources.

To compress the images that were previously uploaded, the user would execute the Job Submitter script on the UI which sends the compression script to the Computational Resources

¹⁵Geant4 is a toolkit for the simulation of the passage of particles through matter. www.cern.ch/geant4

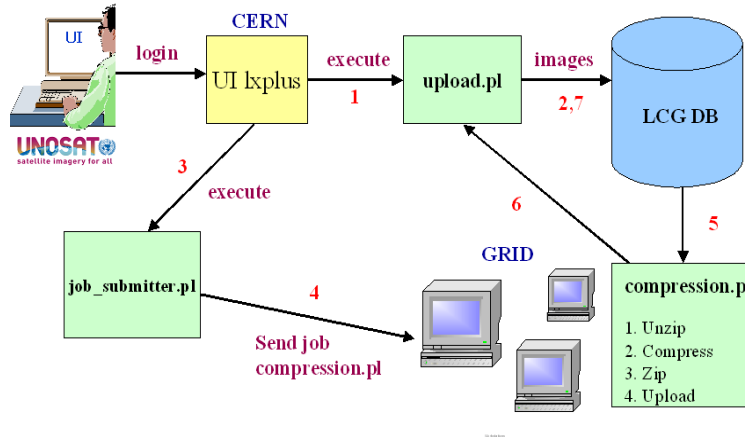


Figure 9: The UNOSAT Grid system. The user can upload files to the Grid, submit a compression job to compress previously uploaded images, and download satellite images from the Grid. Numbers on arrows indicate the sequence of action.

of the Grid. This script downloads the satellite images, unzips the files, compresses the images, re-zips and uploads (using `upload.pl`) the files back to LCG resources

4 Conclusions and Future Work

Our work has demonstrated that the grid paradigm was relevant to the needs of UNOSAT. It was shown how it is possible to upload, store, compress and download UNOSAT's satellite images using the functionality of the LHC Computing Grid being developed at CERN. Over a period of two months the authors have achieved the following results:

- An entire Grid Infrastructure was created for UNOSAT.
- A skeleton framework to process and store satellite images on the Grid.
- An application was developed in conjunction with the CERN ARDA group to allow UNOSAT to download images based on Geographical Coordinates.
- All programs were successfully tested on Grid Resources.

However, there is a significant amount of work remaining to be carried out in this area. For example, in the future it would be necessary to build on the skeleton framework described in this paper to eventually create a general UNOSAT Grid system. It would also be very useful to interoperate the Grid services with the UNOSAT website in order to provide widescale access to the Grid functionality to the humanitarian community. Finally it would be necessary to continue negotiations with owners of resources in order to secure sufficient processing power and storage for UNOSAT's current and future requirements.

The work described in this paper provides UNOSAT with preliminary research that they can use to make an informed decision on their future strategy regarding Grid technology. Ultimately the success (or otherwise) of using the Grid for satellite imagery with UNOSAT will be judged by how the added functionality and resources directly aids the humanitarian community working in the field.

References

- [1] LHC Computing Grid Project. <http://lcg.web.cern.ch/LCG>
- [2] EGEE: Enabling Grids for E-science in Europe, <http://public.eu-egee.org>
- [3] S. Andreozzi, M. Sgaravatto, C. Vistoli. Sharing a conceptual model of Grid resources and services. Conference on Computing in High Energy and Nuclear Physics (CHEP 2003), March 2003, La Jolla CA (USA).
- [4] E. Bjorgo. Current and Potential Uses of Satellite Imagery in UN Humanitarian Organizations. United States Institute of Peace, Virtual Diplomacy Series No. 12, March 2002.
- [5] Y. Teo, S. Tay, and J. Gozali. Distributed Geo-rectification of Satellite Images using Grid Computing.
- [6] J. Montagnat, V. Breton, I.E. Magnin. Using grid technologies to face medical image analysis challenges.
- [7] J.T. Moscicki, H.C. Lee, S. Guatelli, S.C. Lin, M.G. Pia. Biomedical Applications on the GRID: Efficient Management of Parallel Jobs.
- [8] M. Kitajima¹, T. Kosaka, K. Yamazaki, R. Teramoto, G. Kawamura, S. Date, S. Shimojo, H. Matsuda. Bridging life science databases using grid technology for genome-based drug discovery.
- [9] Y. Saito, M. Gomi, H. Matsuda, N. Goto, K. Kurokawa, T. Yasunaga. Real-Time Mapping System of cDNAs and Genomes Using Grid Computing.
- [10] K. Teranishi, P. Raghavan, Zi-Kui Liu. Towards A Grid Enabled System for Multicomponent Materials Design.
- [11] J.W. Park, S.H. Park, I.S. Hwang, J.J. Moon, Y. Yoon, S.J. Kim. Optimal blade system design of a new concept VTOL vehicle using the Departmental Computing Grid system.