# Extending the *Fermi*-LAT Data Processing Pipeline to the Grid

S Zimmer<sup>1,2,6</sup>, L Arrabito<sup>4</sup>, T Glanzman<sup>3,6</sup>, T Johnson<sup>3,6</sup>, C Lavalley<sup>4</sup> and A Tsaregorodtsev<sup>5</sup>

<sup>1</sup> Department of Physics, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden

<sup>2</sup> The Oskar Klein Centre for Cosmoparticle Physics, AlbaNova, SE-10691 Stockholm, Sweden

<sup>3</sup> SLAC National Accelerator Laboratory, Stanford University, Stanford, CA 94305, USA

 $^5$ Centre de Physique des Particules de Marseille, 163 Av<br/> de Luminy Case 902 13288 Marseille, France

 $^{6}$  on behalf the  $Fermi\text{-}\mathrm{LAT}$  collaboration

E-mail: zimmer@fysik.su.se

**Abstract.** The Data Handling Pipeline ("Pipeline") has been developed for the *Fermi* Gamma-Ray Space Telescope (*Fermi*) Large Area Telescope (LAT) which launched in June 2008. Since then it has been in use to completely automate the production of data quality monitoring quantities, reconstruction and routine analysis of all data received from the satellite and to deliver science products to the collaboration and the Fermi Science Support Center. Aside from the reconstruction of raw data from the satellite (*Level 1*), data reprocessing and various event-level analyses are also reasonably heavy loads on the pipeline and computing resources. These other loads, unlike Level 1, can run continuously for weeks or months at a time. In addition it receives heavy use in performing production Monte Carlo tasks.

In daily use it receives a new data download every 3 hours and launches about 2000 jobs to process each download, typically completing the processing of the data before the next download arrives. The need for manual intervention has been reduced to less than 0.01% of submitted jobs.

The Pipeline software is written almost entirely in Java and comprises several modules. The software comprises web-services that allow online monitoring and provides charts summarizing work flow aspects and performance information. The server supports communication with several batch systems such as LSF and BQS and recently also Sun Grid Engine and Condor. This is accomplished through dedicated job control services that for *Fermi* are running at SLAC and the other computing site involved in this large scale framework, the Lyon computing center of IN2P3. While being different in the logic of a task, we evaluate a separate interface to the Dirac system in order to communicate with EGI sites to utilize Grid resources, using dedicated Grid optimized systems rather than developing our own. More recently the Pipeline and its associated data catalog have been generalized for use by other experiments, and are currently being used by the Enriched Xenon Observatory (EXO), Cryogenic Dark Matter Search (CDMS) experiments as well as for Monte Carlo simulations for the future Cherenkov Telescope Array (CTA).

<sup>&</sup>lt;sup>4</sup> Laboratoire Univers et Particules de Montpellier, Universite Montpellier 2, CNRS/IN2P3, Montpellier, France



**Figure 1.** Graph for a typical MC task. Each task contains three steps: *runMonteCarlo*, *transfer2SLAC* and executed from the Pipeline server *register-ds* which executes a Jython scriptlet registering the data proucts received in the Fermi Data Catalog. Since failures of streams for this task often require user intervention, we refrain from implementing an automatic roll-back if the first step fails.

## 1. The Fermi-LAT mission

The Fermi spacecraft supports two gamma-ray instruments; the Large Area Telescope (LAT) [1] and the Gamma-ray Burst Monitor (GBM) [8]. The LAT is a wide-field gamma-ray telescope (20 MeV - 300 GeV) that continuously scans the sky, providing all-sky coverage every two orbits. The GBM is an all-sky monitor (10 keV - 25 MeV) that detects transient events such as occultations and gamma-ray bursts (GRB). GBM detections of strong GRBs can result in an autonomous re-point of the observatory to allow the LAT to obtain afterglow observations. The satellite sends data to the ground every 3 hours. Data is transferred via relay satellites at 40 MB/s to the White Sands ground station. It then follows a leased line to the Mission Operations Center (MOC) at Goddard Space Flight Center where data is split into two parts and sent for science processing to both the GBM and the LAT teams, the latter being located at Stanford National Accelerator Laboaratory (SLAC).

# 2. Computing Requirements and Principle Task Types

The processing of downlinked satellite data is a time-critical operation. It is, therefore, necessary to automatically trigger the Pipeline for each newly arrived block of data, and to exploit parallel processing in a batch farm to achieve the required latency for the production of the various data products. This processing is complex and is abstracted in a process graph. An xml representation of this process graph is interpreted by the Pipeline to become a *task*. Once defined, a task is exercised by the creation of *streams*, each of which is one instance of the process graph and which consists of an arbitrary number of interconnected batch jobs (or scripts) known as *process instances* [7]. To that end, the software is designed to be able to handle and monitor thousands of streams being processed at separate sites with a daily average throughput of about 1/2 CPU-year of processing. To date peak usage has been 45,000 streams in a single day and 167 CPU-years of processing in a single month [4].

The Pipeline was designed with the L1 data processing task ("level 1 processing", the core data processing of raw data that comes from the satellite), automatic science processing (ASP) and Monte Carlo simulations as principle task types in mind and to ensure the tight connection to the Fermi Data Catalog. It is literally impossible to depict the whole L1 task scheme in one single figure as it contains many dozens of sub-stream creations, dependencies and automatic re-run mechanisms. Thus we refrain from including them in this paper. For details the reader is referred to [10].

Instead we show the simple layout of a Monte Carlo task (as our efforts initially are geared towards porting them to Grid sites) in Fig. [1] that does not rely on external dependencies such as local databases. This task simply consists of 3 steps, the generation of Monte Carlo data on a computing node, the transfer of the MC products to SLAC and their registration in the Data Catalog. The first 2 steps are batch operations where the registration step takes only split

seconds and is achieved through running a dedicated Jython [14] scriptlet.

## 3. Pipeline Components

We use a three tier architecture as shown in Fig. [2(a)] comprising of back-end components, middle-ware and front-end user interfaces. We describe them more in detail below.

#### 3.1. Back-End components

Core of the back-end components is the Oracle Database that stores all processing states. In addition we make extensive use of Oracle technologies such as the scheduler that is used to run periodic jobs for system monitoring including resource monitoring of the Oracle server itself. Most quantities are made available to the user through trending plots that allow quick judgment about the state of the system along with its resource usage. Another back-end component is the Pipeline Job Control Service. Using Remote Method Invocation (RMI [16]) they publish the uniform interface and communicate with the Pipeline server. We discuss some more details on this component in the next section.

To provide asynchronous persistent messaging from batch jobs to the Pipeline server we use email messaging. At the beginning of a job an email is sent detailing the host name and other worker node specific information. Another email is sent to indicate that a job has finished. This email may also contain additional commands that invoke new Pipeline commands, such as creation of sub-streams as the next step. In order to avoid overloading the SLAC email server, by relaying tens of thousands emails per day, we use a dedicated email server running the free Apache JAMES [12] software.

# 3.2. Middleware

The Pipeline server is the core of the system and contains two pools for threads, a worker and an admin pool. When a batch process is ready to run on one of the farms, a thread is allocated on the worker pool to perform the submission using the appropriate Job Control Service. Extensions to the pipeline (called *plugins*) can be used to add additional functionality, for example to provide access to experiment specific databases or to communicate with other middleware services without compromising the experiment independent design of the core Pipeline software. Plugins are written in Java and loaded dynamically when the Pipeline starts. The user can also provide Jython scriptlets to run within the Pipeline threads to perform simple calculations and to communicate with plugins. The Fermi data catalog is implemented as a plugin.

The Pipeline server API allows queries for processes, stream management as well as means to get or set environment variables. The admin thread pool is used to identify work to be delegated to the worker pool. This includes gathering processes which are ready to run as well as various database queries.

#### 3.3. Front-End Components

We provide a subset of the Pipeline API as Java Management Extension (JMX), that provides a call interface to various user-interface applications. These come both as Web interface and line command applications. The web interface provides password protected world-wide access to the Pipeline and its control interfaces and allow simultaneous monitoring of tens of thousands of jobs in various states. For detailed technical information on the Pipeline components the reader is referred to [7].

# 3.4. The Pipeline Job Control Service

Each Job Control Service (refer to Fig. [2(b)]) implements job control and status methods that are specific to the batch or Grid system. To that end each Job Control Service needs



(a) Pipeline Architecture and Key Technologies



(b) Implementation of Job Control Services and Data Catalog

**Figure 2.** Pipeline architecture and key technologies (a) and Connection of Pipeline Server with existing Job Control Services (b). To date the Pipeline supports interfaces to BQS, LSF, Condor, SGE and the DIRAC server that has been configured for the use with *Fermi*.

to provide the following commands: SubmitJob, GetJobStatus, KillJob. The code for the Job Control Service is written in Java and runs as a daemon on a dedicated service machine at each computing site. To date the Pipeline supports LSF [11], BQS and more recently also Condor as well as Sun Grid Engine (SGE [17]). At present *Fermi* uses LSF at SLAC and SGE at the Lyon computing center. The code can easily be adapted to any other desired batch system that follows the same job logic as the currently supported ones. To that end a java class *BatchJobControlService* and *BatchStatus* need to be implemented for the desired new batch system, where *Batch* denotes the system to be implemented.

For the use with DIRAC (Distributed Infrastructure with Remote Agent Control) the Job Control code wraps python scripts that provide the bridge in both format and language by using the dedicated DIRAC API. We describe further aspects of the DIRAC system and motivate our choice to use it in Section 4.2.

In the past a task was defined on top level to be handled by a specific Job Control Service. Recently the Pipeline has been enabled to support multi-site tasks that allow the Pipeline server to delegate the running of commands to our DIRAC Job Control Service while we leave the transfer step to the Lyon-based SGE Job Control Service.

## 4. Integrating the Pipeline with the Grid

As of this writing the LAT has been granted resources both at SLAC (dynamic allocation scheme) and Lyon (guaranteed 1200 cores allocation). At Lyon all resources are used for Monte Carlo productions while at SLAC the total allocation is shared between L1, Monte Carlo production and individual user jobs from the collaboration. As a recent challenge we have started reprocessing all of our data that was taken from the beginning of the mission up until now and are reprocessing it with our current state of knowledge about the experiment. The reprocessing requires a significant amount of our computing resources. As allocations are dynamic, users running science analysis may be directly impacted through less available slots on the batch farm. Another recent challenge was a massive Monte Carlo production of proton runs that occupied our resources both at SLAC and Lyon for several months. While not being particularly storage intensive, we do require significant amounts of CPU time. Since this production run was setup with the previous interation of the instrument response, dubbed "pass 6", it is likely that simulation requests of this kind may be repeated as our knowledge of the experiment grows. It is thus important to investigate possibilities to extend our resources that can be utilized within the current Pipeline framework.

## 4.1. The glast.org VO

Although we perform standard computing tasks at our two sites on local batch farms, there exists the virtual organization (VO) glast.org. This organization was founded in 2009 to provide access to gLite [9] resources granted by participating institutions in Italy and France. At present the VO includes 13 sites that are partially enabled for use by *Fermi*. Its use has however been limited to non-pipeline operations. Most notably, our existing Grid resources have been used for stand-alone pulsar blind searches and some large Monte Carlo simulations. These stand-alone tasks were unable to take advantage of the pipline's integration with the batch system and data catalog which made them more man-power intensive than their Pipeline counterparts.

#### 4.2. DIRAC as potential connection to EGI enabled sites

In order to optimize the resource usage provided by the EGI [5] sites supporting the glast.org VO, we are exploring the use of the DIRAC system as potential connection between the Pipeline and the Grid resources [18].

The DIRAC system, originally developed to support production activities of the LHCb experiment, is today a general solution to manage the distributed computing activities of several



**Figure 3.** Our proposed integration of the DIRAC system with the Pipeline. We establish an additional Job Control Daemon that delegates job requests to the DIRAC server. This server maintains its communication to the worker node through a dedicated notification service. Log files are stored on a DIRAC storage element and accessed via a web server.

communities. One of its main components is the Workload Management System (WMS). The key feature of the DIRAC WMS is the implementation of the pilot job mechanism, which is widely exploited by all LHC communities as a way to guarantee high success rate for user jobs (workloads). More in detail, workloads are submitted to the DIRAC WMS and inserted in the DIRAC *Central Task Queue*. The presence of workloads in the DIRAC *Central Task Queue* triggers pilot jobs submission. Pilot jobs are regular Grid jobs that are submitted through the standard gLite WMS. Once the pilot job gets on the worker node, it performs some environment checks and only in case these tests succeed, the workload is pulled from the DIRAC *Central Task Queue*. In case the environment checks fail or the pilot job gets aborted (for example because of a mis-configured site), only the pilot job is affected. The net result of this mechanism is a significant improvement on the workload success rate. Also, since the resources are pre-reserved by pilot jobs, the waiting time to get the workload execution started is reduced.

# 4.3. Establishing an Interface between Pipeline and DIRAC: Design considerations

While DIRAC has specifically been designed to tackle intrinsic Grid inefficiencies through its pilot job concept, it cannot solve some of the initial issues when being connected to the Pipeline system. In particular:

- The Grid uses personalized certificates. When submitting jobs through the Pipeline web interface, we submit them with a generic user ID, which is most feasible as the number of authorized Pipeline users is small.
- Some of our sites, in particular those from INFN, cannot directly send emails to the Pipeline server, thus rendering the standard communication of job status and task logic de-facto

unusable.

• In general large data on the Grid would be stored on Grid Storage Elements (Grid SE) and require user intervention to download it once finished. This makes the automatic copy to our central SLAC XROOTD space difficult to be used, in particular using the tight integration with the data catalog.

We decided to implement an interface to the DIRAC system for several reasons:

- Independence of Grid middleware: Currently the VO comprises only gLite sites but DIRAC is able to communicate with all common Grid middleware platforms, thus making it less difficult to connect to other Grid initiatives, such as the Open Science Grid in the US [15].
- Additional Monitoring functionality: DIRAC introduces a more detailed job monitor that in addition reports minor and major application statuses together with the overall job status, thus allowing easier debugging of tasks without the definite need to inspect log files manually. This task has always been among the most time consuming. We hope to achieve an improvement by using the new monitor, albeit as read-only system. The reason for that is that the Pipeline ID and the Grid ID are generally not the same and a re-submitted job on the Pipeline keeps its ID while on the Grid it is relaunched and assigned a new unique Grid ID.
- While *Fermi* is entering its 4th mission year, the number of developers for software has begun to dwindle. Thus it is important to keep the required manpower as low as possible while maintaining full performance and possibly improving the capabilities of the Pipeline. By using DIRAC we can build our interface on a system in active development with many possibilities to influence the development process to meet our needs.

Implementing the DIRAC interface comes in two steps, first the Job Control Daemon as described in the previous section that wraps the basic Pipeline commands through the DIRAC API and secondly the configuration of a dedicated DIRAC server. The implementation is shown in Fig. [3].

We make use of two of DIRACs core technologies: the pilot factory mechanism is used to renew proxies and authenticate the Job Control Daemon to the Grid allowing us to use one certificate retaining our previous user scheme. The DIRAC notification service provides means for the DIRAC server to communicate with Grid worker nodes. This service is both safe to inception and can be modified to relay the content of our status emails to the DIRAC server that itself implements the email communication with the Pipeline server.

One typical by-product of running our code are more or less extensive log files. Usually we have several hundreds of small log-files that each do not exceed a few MB. In the past other experiments had to artificially enlarge their log files to ensure the stability of the Grid SEs due to journalling of small files. We can effectively circumvent this by declaring a local storage element at the computing center in Lyon as a dedicated DIRAC SE. These storage elements are addressed like normal Grid elements but they do not need journaling to function. Since this is a local SE, we can view its content e.g. through a Web server providing access to log-files that can directly be linked from the Pipeline server.

In conclusion we believe that this solution provides an easy to implement and maintain interface to the Pipeline system used for *Fermi*-LAT.

# 5. Conclusions

The existing VO resources of glast.org suggest the possibility to establish a new connection to Grid services with our existing Pipeline architecture. We mitigate issues such as the transfer and subsequent registration of data products at the Fermi Data Catalog by using existing Pipeline technologies. Grid inefficiencies are handled by the DIRAC system that acts as a broker providing asynchronous communication with Grid worker nodes and a closed mechanism to automatically renew proxies to use for Grid operations. We leave handling of meta data to the existing Pipeline technologies.

As such, the Pipeline software itself was designed in a manner to not contain any Fermi specific functionality. Through the plug-in feature of the middleware it is successfully used for data handling for the Enriched Xenon Observatory or MC simulations for the Cryogenic Dark Matter Search experiment (CDMS [2]), the upcoming Cherenkov Telescope Array (CTA [3]) and the Large Synoptic Survey Telescope (LSST [13]), and is also being considered for use by other NASA missions.

#### Acknowledgments

We acknowledge the ongoing generous support of SLAC (US) and IN2P3 (France) as well as we thank them for increasing allocations at both sites. Furthermore we are thankful for Grid resources provided by INFN and IN2P3.

SZ likes to thank the organizers for a stimulating meeting and particularly helpful discussions with Stéphane Guillaume Poss.

The *Fermi*-LAT Collaboration acknowledges generous ongoing support from a number of agencies and institutes that have supported both the development and the operation of the LAT as well as scientific data analysis. These include the National Aeronautics and Space Administration and the Department of Energy in the United States, the Commissariat à l'Energie Atomique and the Centre National de la Recherche Scientifique / Institut National de Physique Nucléaire et de Physique des Particules in France, the Agenzia Spaziale Italiana and the Istituto Nazionale di Fisica Nucleare in Italy, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), High Energy Accelerator Research Organization (KEK) and Japan Aerospace Exploration Agency (JAXA) in Japan, and the K. A. Wallenberg Foundation, the Swedish Research Council and the Swedish National Space Board in Sweden.

Additional support for science analysis during the operations phase is gratefully acknowledged from the Istituto Nazionale di Astrofisica in Italy and the Centre National d'Études Spatiales in France.

#### References

- [1] Atwood et al. ApJ **697** (2009)
- [2] http://cdms.berkeley.edu/
- [3] http://www.cta-observatory.org/
- [4] Dubois, ASP Conf. Ser. 411, eds. D Bohlender, D Durand, & P Dowler (2009)
- [5] http://www.egi.eu/
- [6] http://www-project.slac.stanford.edu/exo/
- [7] Flath et al., ASP Conf. Ser. 411, eds. D Bohlender, D Durand, & P Dowler, SLAC-PUB-13549 (2009)
- [8] See http://gammaray.msfc.nasa.gov/gbm/ for details on the GBM
- [9] http://glite.cern.ch/
- [10] Focke, "Implementation and performance of the Fermi LAT level 1 pipeline", II. Fermi Symposium (2009)
- [11] http://www.platform.com/workload-management/high-performance-computing
- [12] http://james.apache.org
- [13] http://www.lsst.org/lsst/
- [14] http://www.jython.org/
- [15] https://www.opensciencegrid.org
- [16] http://java.sun.com/javase/technologies/core/basic/rmi/index.jsp
- [17] http://www.oracle.com/us/sun/index.htm
- [18] Tsaregorodtsev at al., J. Phys. Conf. Ser. 119 (2008)