

**Memorandum of Understanding**  
**between**  
**Fermi National Accelerator Laboratory**  
**and**  
**The Astrophysical Research Consortium**  
**concerning**  
**Archiving and Serving Data from the Sloan Digital Sky Survey**

October 22, 2008

## **I. Preamble**

Fermilab has been the primary host for processing, storing, and serving data from SDSS (2000 - 2005) and SDSS-II (2005-2008). The data have been released incrementally approximately once per year, starting with the Early Data Release (EDR), and extending through DR6 (currently). DR7 will be the final data release; it will be made public on October 31, 2008.

Since the scientific use of the SDSS data will continue indefinitely, ARC is seeking a way to secure the long-term future of the SDSS data and access tools. After the creation of DR7, the databases will be static and it is anticipated that the level of effort to maintain the databases will be dramatically lower than it has been during the operations phase of SDSS. It is envisioned that an SDSS data archive that is secure in the long term can be efficiently built on the foundations that are already in place.

Fermilab's involvement with creating and hosting the SDSS archive has been valuable to the astrophysics groups at Fermilab as well as to the user community at large. This value will hold well into the future.

The development and operations of the data-handling systems have included several universities in addition to Fermilab. A cooperative sense of mission has emerged for supporting the SDSS data products. As a first step to expand expertise and responsibility for storing and serving the SDSS data, the primary public access to the DR5 database is now at the University of Chicago and at the Johns Hopkins University, with Fermilab as a secondary site. For similar reasons, the Helpdesk responsibility was moved from Fermilab to the University of Chicago in Fall of 2007.

This MOU proposes a fixed-term agreement (see Section VII) for practical reasons, but this does not suggest that responsibilities for the curation of the data are expected to terminate in the foreseeable future. Rather, at the end of the agreement, ARC will reevaluate the status of the data archive and determine how best to construct a new agreement between the partners.

## **II. General Requirements**

There are fundamentally two requirements, each with corresponding tasks: archiving the data and serving the data.

The archiving task ensures that the data products described in the Appendix are preserved. This task implies a system for error checking and a system for replacing corrupted files from a back-up source. There is also an implication that the data be migrated as needed from a given hardware system or a given operating system to its replacement.

The serving task relates to the maintenance of publicly accessible interfaces to files and to a database system. ARC/SDSS will provide the corresponding stand-alone, working products to Fermilab as configured at the time of the release of DR7. No additional requirements will exist beyond the capabilities of the current configuration for DR7, e.g. speed with which queries are processed, speed of data transfer, and up-time (availability). Continuing use of DR1 (for example) suggests that researchers find value in earlier data releases. A goal (as opposed to a requirement) is to serve not just DR7, but also the earlier data releases.

In addition to preserving files and making files available, it is scientifically important to be able to reproduce, in principle, the values of derived parameters, given the input data and the processing configuration. ARC/SDSS will ensure that all of the information needed to retrace the processing steps is available within the set of files described in the Appendix.

### **III. SDSS Data Products**

The nature, products, and volume of the SDSS data are described in detail in the Appendix. This Section provides a brief overview to define terms.

- \* Data Archive Server (DAS): processed data in flat file format with tools for efficient navigation of the directory structure. Access to data is provided via rsync and http interfaces.
- \* Catalog Archive Server (CAS); SkyServer; CasJobs: photometric and spectroscopic parameters stored in a searchable database. The CAS database is built on SQLServer. It has a front-end web server interface for astronomers, and another (called SkyServer) that includes pages designed for education and public outreach. CasJobs is another front-end interface to CAS: it is a web service that allocates local compute resources (cpu and storage) that is designed for large jobs in batch mode.
- \* Survey Software: all software is currently checked in to the CVS code repository at Fermilab. See Appendix for an itemization of software by category. The contents of the CVS repository will be exported, and it is these exported files that will be available (via the DAS).
- \* Raw Data: pixel data as acquired from the instruments (imaging camera and spectrographs; Photometric Telescope frames, engineering log files). These data are distinct from the processed data that are publicly available in the CAS and DAS.

### **IV. Helpdesk**

A Helpdesk facility exists to assist users with the interfaces and the interpretation of the results of downloads. The service consists of a central hub (at The University of Chicago Library, which maintains the Helpdesk ticket tracking system called QuestionPoint) and a distributed

network of volunteer specialists who are available to answer a modest number of questions per month.

## V. Scope of Work at Fermilab

Fermilab will continue its current role as the primary center for storage of SDSS data and for serving the SDSS data to the public. After 1 October 2008, no new processing will occur and no new files will be loaded. The current SDSS machines used for CAS, DAS, and data processing will be reconfigured by 31 October 2008 to streamline operations and maintenance of the servers for CAS and DAS in the future. A copy of DR7 will be archived to the Enstore tape robot immediately after it is made public on 31 October. Once this work is accomplished, the routine level of effort for systems administration will be typical of systems this size that serve static sets of data. As Fermilab chooses to upgrade any of the hardware or other aspects of its overall operating environment, the SDSS data and interfaces will need to be appropriately migrated.

The roles and responsibilities for Fermilab will be:

- 1) Maintain and act as primary host for DAS, including both the web (http) and rsync servers. Maintain the server hardware and administer the systems, including for example configuring the apache web server. Maintain and administer the DAS interface CGI scripts and web pages.
- 2) Maintain and act as primary host for CAS (all DR's), SkyServer, and CasJobs. Maintain the service accounts necessary for hosting these systems.
- 3) Maintain an archive of static files of the exported contents of the CVS repository in which SDSS software is currently stored.
- 4) The *Dervish* and *astrotools* products that are currently in the CVS code repository will not be maintained as part of this MOU.
- 5) Maintain the integrity of the raw data. Fermilab will have one copy of the raw data in the DAS, and another distinct copy other than that in the DAS. Johns Hopkins University Libraries will have another copy of the raw data, such that each site backs the other up.
- 6) Provide recovery services to restore working systems in the event of hardware failure or file corruption.
- 7) In addition to the initial set-up, provide backup and technical support to the mirror sites at The University of Chicago and at Johns Hopkins University (see Section VIII), including access to data in the event that one of these sites suffers a loss.
- 8) ARC/SDSS will produce an inventory of all of the SDSS data files stored on the Enstore tape robot in addition to all of the files in the DAS. This inventory, which will be maintained by Fermilab, will identify the set of files on the Enstore tape robot, not replicated in the DAS, that would be required to restore the mirror sites.

9) Update as necessary the registration of the sdss.org domain name. Maintain the sdss.org site, or, alternatively, assist with a transfer of the sdss.org pages to another site.

10) Members of the Center for Particle Astrophysics will be expected to help answer Helpdesk questions concerning scientific issues, and others in the Computing Division will be expected to help answer questions concerning technical issues. Based on current experience, the average number of questions per month assigned to an individual should not exceed 2.5. These individuals will be designated on a list of experts available for this purpose (along with others from other institutions). All such designated individuals will be expected to use the system for addressing queries employed by The University of Chicago Library (currently QuestionPoint).

11) Key performance indicators and metrics, as described in Section IX, will be reported quarterly.

The Head of the Computing Division at Fermilab will be responsible for the conduct of this work and will serve as the point of contact with ARC.

## **VI. ARC Resources**

ARC will contribute at least \$60k per year to Fermilab for five years in partial support of the effort described in Section V (including for example covering the cost of licenses for software products). These funds will be applied by Fermilab according to a Statement of Work that is generated annually by the ARC Business Manager. The SoW will be reviewed and revised annually by both parties. It is expected that the first SoW, and the corresponding ARC resources made available to Fermilab, will reflect a larger scope of effort than in the following years, namely transition activities after 31 October.

## **VII. Duration**

This agreement extends for five years, specifically from the date of signing through 31 December 2013. If at the end of this interval Fermilab anticipates that there will be significant changes to the next agreement, adequate time shall be given to ARC so that it can properly devise a plan for the following years.

## **VIII. Interactions with Other Institutions**

The overall goal is to create a cooperative network of sites for load balancing and with mutual back-up services for data redundancy and recovery. Responsibilities of other institutions are sketched below; these will be augmented, refined and formalized in separate MOU's.

### ***University of Chicago***

contingent on securing funding, University of Chicago:

- 1) will act as secondary host (mirror) for CAS (DR7) and DAS
- 2) will implement a load-sharing arrangement with Fermilab for CAS (as is already being done for DR5) and provide technical support to maintain availability of CAS and DAS at the UofC Library, i.e., serve as effective back-up to Fermilab
- 3) will lead responsibility for running the Helpdesk
- 4) will identify, organize, document, and curate the physical and electronic SDSS Administrative Records (proposals, reports, governance policies, reviews, letters, email list archives, public information materials, etc.)
- 5) may act as secondary archive for raw data, contingent on discussions with Networking Services and Information Technology at The University of Chicago, as well as the securing of funding.

### ***Johns Hopkins University Department of Physics and Astronomy***

- 1) act as secondary host (mirror) for CAS, SkyServer and CasJobs, including the mydb's.
- 2) develop techniques for mining the archived mail lists
- 3) monitor the traffic at the CAS sites at JHU, Fermilab, and UofC, and make this information available to the other sites

### ***Johns Hopkins University Libraries***

- 1) JHU Libraries will archive a copy of the raw data, and the processed DAS datasets.
- 2) JHU Libraries has an interest in trying to develop the provenance chain from the processed data within the releases to the underlying raw data. This task may require cooperation with Fermilab to ensure that all necessary files, not otherwise accessible from the DAS or the files exported from CVS, are available.
- 3) In support of 2) above, the data streams from the telescope and instruments that document prevailing environmental conditions and parameters; the processing programs and their configurations; and the underlying operating systems need to be available. JHU Libraries will not focus on providing support for the processing programs, but rather archive them so that in the future investigators may examine the type of processing that occurred.

## IX. Key Performance Indicators and Metrics

Each institution serving as a host site for the SDSS data archive shall implement systems for measuring and reporting on Key Performance Indicators (KPIs) and metrics. At a minimum, each host institution will monitor and report on the following:

### *System Performance*

- DAS system uptime, per month
- CAS system uptime, by data release, per month

### *Helpdesk Support (or User Support)*

- # of Helpdesk tickets assigned to the organization, per month
- % of Helpdesk tickets closed within 2 business days
- % of Helpdesk tickets closed within 5 business days

### *Data Usage (or Access)*

- Volume of data transferred monthly from DAS via rsync interface
- Volume of data transferred monthly from DAS via wget interface
- # of web hits per month, by CAS data release
- # of SQL queries processed per month, by CAS data release
- # of returned rows, by CAS data release

### *Data Integrity*

- # of corrupt files found on the DAS, per month
- # of corrupt files replaced on the DAS, per month
- Volume of data transferred to archive mirror sites, per month

Each host institution shall provide ARC with a quarterly report that summarizes performance against performance indicators and metrics, and highlights operational issues or concerns (for example, Fermilab shall report activity related to checking the integrity of files on Enstore).

These quarterly reports, as well as the annual invoices, shall be sent to:

Business Manager  
Astrophysical Research Consortium  
c/o Department of Astronomy, Box 351580  
University of Washington  
Seattle, Washington 98195-1580

## X. Cessation

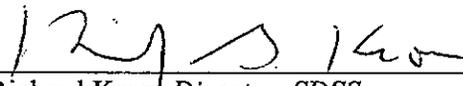
During the interval of agreement, 1 January 2009 through 31 December 2013, should either party to this MOU fail to obtain the necessary resources to support the scope of work described herein, it shall notify the other in writing, and all parties will seek alternative resources (possibly including transfer of responsibilities to another party) or arrange for an appropriate reduction of the project activities.

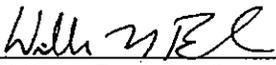
## XI. Amendments

This MOU establishes the basic guidelines for the agreement between ARC and Fermilab, but it is natural to expect that the relationship will undergo evolution, which may lead to changes to be recorded in occasional future amendments as mutually agreed upon in writing by the two parties.

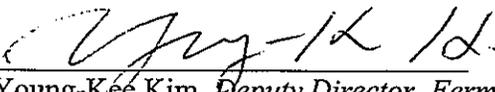
### Concurrence:

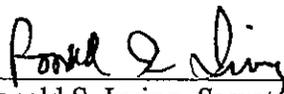
 (date) 10/22/08  
Victoria White, Head, Computing Division, Fermilab

 (date) 22 Oct 2008  
Richard Kron, Director, SDSS

 (date) 10-22-2008  
William Boroski, Project Manager, SDSS

### Approved:

 (date) 10-22-2008  
Young-Kee Kim, Deputy Director, Fermilab

 (date) 10-30-08  
Ronald S. Irving, Secretary/Treasurer, Astrophysical Research Consortium

## Appendix

### SDSS Long-Term Scientific Data Archive

---

This document summarizes the SDSS data products that will be maintained in the long-term scientific data archive. In particular, the long-term archive will include the following items:

- Final version of the Data Archive Server (DAS)
- All of the incremental Catalog Archive Server (CAS) releases (EDR through DR7)
- All of the software used to collect the data
- All of the software used to process and calibrate the data
- The interface code for the SkyServer interface
- The raw data in binary format.

**Data Archive Server** - The final version of the DAS will contain the full set of processed data in flat file format. Since the DAS contains the full set of flat files spanning all releases, there is only one DAS version of the DAS to be archived. The DAS will require approximately 65TB of usable disk space at survey completion.

**Catalog Archive Server** - A total of eighteen distinct CAS databases were released to the public during the course of the survey. All eighteen versions should be preserved in the long-term archive. Table 1 summarizes the contents and data volume of each CAS database. Storing the CAS databases in uncompressed form will require approximately 34 TB of disk space. It may be possible to reduce this by storing the CAS files in compressed format; however, compression tests have not been performed.

Table 1. Data volume of the CAS databases

Release	Product	Volume (TB)
EDR	EDR	0.065
DR1	BestDR1	0.64
	TargDR1	0.50
DR2	BestDR2	0.94
DR3	BestDR3	1.23
	TargDR3	1.35
DR4	BestDR4	1.74
	TargDR4	1.63
DR5	BestDR5	2.15
	TargDR5	1.77
DR6	BestDR6	2.85
	TargDR6	1.95
	SegueDR6	0.57
DR7	BestDR7	3.60
	TargDR7	1.95
	SegueTargDR7	1.10
	RunsDB	6.70
	Stripe 82 DB	7.00
<b>Total</b>		<b>37.74</b>

**Survey Software** - All of the software associated with production survey operations is currently contained in the CVS code repository at Fermilab. To avoid the need to maintain the CVS code repository, its contents will be exported to static files which will be available via the DAS. The set of files will include:

- All of the software used at Apache Point to operate the telescopes; instruments, and various support systems, including the data acquisition and observing software;
- All of the pipeline software used to process and calibrate the data;
- All of the software used to load the data into the CAS databases;
- All versions of CasJobs and the SkyServer interface;
- The html web pages associated with [www.sdss.org](http://www.sdss.org).

The estimated size of the SDSS code repository at survey completion is 6 GB compressed. The repository is currently using 20GB of disk space. Approximately 7.5GB of that is a subset of the astrolog data, which do not belong in CVS (they will be included in the DAS). A significant portion of the remaining space is consumed by items that also do not need to be saved in the repository. For example, a number of people have used the CVS repository to collaborate on science papers over the years, and the various associated plots and graphics take up significant space. The SDSS Collaboration is reviewing these files to identify non-production items before final archiving. After excluding the astrolog data and compressing the remainder, the total size drops to 5.4 GB. Removal of draft science papers and like items should drop the final size even further.

**Raw Data** -- All of the raw data are currently stored in binary format in the Enstore tape robot at Fermilab. Fermilab intends to continue hosting the raw data and periodically migrating the full data set to new media as necessary and appropriate. There is benefit in having the raw data set mirrored at geographically dispersed sites to protect against catastrophic disasters. The estimated volume of raw data at survey completion is 32 TB.

Table 2 summarizes raw data volumes by data type; the description notes which data are stored in compressed format.

Table 2. Raw data volume stored on tape in Enstore.

File Family	Data Type	Description	Size (TB)
apoLogs	Observing log files	Files include astrolog, mcptpm, md5sum, murmurlog. Prior to MJD 53616, these were backed up every full moon and not compressed or gzipped. From MJD 53616 forward, all of the apoLog files are gzipped when tarred for Enstore.	0.25
rawImagFields	Imaging data	Field and gang files. idR files are compressed (.Z); gang files are not compressed.	24.40
mtRaw	Photometric Telescope data	All files are gzipped when tarred for Enstore	5.56
spectroRaw	Spectroscopic data	All files are gzipped when tarred for Enstore	1.35
<b>TOTAL</b>			<b>31.55</b>

A recent test was performed to verify the gain due to compression. A night's worth of raw spectroscopic data was compressed and the size of the gzipped tar file was compared to the amount of disk space used by the raw data. The size of the gzipped tar file was 0.60 the size of the uncompressed raw data.