

## Data Organization in the SDSS Data Release 1

A.R. Thakar, A.S. Szalay, and J.V. vandenBerg

*Johns Hopkins University, Baltimore, MD 21218*

Jim Gray

*Microsoft Research*

Chris Stoughton

*FermiLab, Batavia, IL 60510*

**Abstract.** The first official public data release from the Sloan Digital Sky Survey ([www.sdss.org](http://www.sdss.org)) is scheduled for Spring 2003. Due to the unprecedented size and complexity of the data, we face unique challenges in organizing and distributing the data to a large user community. We discuss the data organization, the archive loading and backup strategy, and the data mining tools available to the public and the astronomical community, in the overall context of large databases and the VO.

### 1. Introduction

The SDSS Data Release 1 (DR1) is the first officially scheduled public data release of the SDSS data. It is the successor to the Early Data Release (EDR) released in June 2001 ([archive.stsci.edu/sdss](http://archive.stsci.edu/sdss)). DR1 is scheduled for release in Spring 2003, and covers more than 20% of the total survey area (>2k square degrees). The raw data size is about 5 times that of the EDR, i.e., several Terabytes. The catalog data will be about the same size because there will be 3 datasets with several versions of each dataset.

This is the first single release of such a large dataset to the public, and naturally it presents unprecedented challenges. Simply distributing the data and making it available 24/7/365 will be quite an undertaking for the SDSS collaboration. Providing competent data mining tools on this multi-TB dataset, especially within the context and evolving framework of the Virtual Observatory, will be an even more daunting challenge. The SDSS database loading software and data mining tools are being developed at JHU ([www.sdss.jhu.edu](http://www.sdss.jhu.edu)).

### 2. Data Distribution

The master copy of the raw data (FITS files) will be stored at FermiLab. In addition to the master archive at FermiLab, there will be several mirror sites for the DR1 data hosted by SDSS and other institutions. Replication and syn-

chronization of the mirrors will therefore be required. We describe below the configuration of the master archive site. Mirror sites will probably be scaled-down replicas of the master site.

### 2.1. Data Products

There will be three separate datasets made available to the public - two versions of the imaging data and one version of the spectra:

- **Target dataset** - this is the calibration of the raw data from which spectral **targets** were chosen;
- **Best dataset** - this is the latest, greatest calibration and represents the **best** processing of the data from a science perspective;
- **Spectro dataset** - these are the spectra of the target objects chosen from the target dataset.

Within each dataset, the raw imaging data will consist of the Atlas Images, Corrected Frames, Binned Images, Reconstructed Frames and the Image Cutouts in addition to the **Imaging Catalogs** for the Target and Best versions. The spectroscopic data consists of the Raw spectra along with the **Spectro Catalog** and the **Tiling Catalog**.

### 2.2. Data Volume

Table 1 shows the total expected size for a single instance of the DR1 archive - about 1TB. In practice, however, the overall size of the catalog data at a given archive site will be several TB, i.e., comparable to the size of the raw data, since more than one copy of the data will be required for performance and redundancy.

Table 1. Data sizes of the DR1 datasets.

BEST		TARGET		SPECTRO		Indices	Misc. Catalogs	TOTAL
Catalog	Jpegs	Catalog	Jpegs	Spectra	Tiling			
400 Gb	50 Gb	300 Gb	50 Gb	10 Gb	10 Gb	150 Gb	20-30 Gb	1 TB

## 3. Archive Operations

### 3.1. Archive Redundancy, Backups and Loading

It will be necessary to have several copies of the archive at least at the master site, to ensure high data availability and adequate data mining performance. Figure 1 shows the physical organization of the archive data and the loading data flow. Backups will be kept in a deep store tape facility, and legacy datasets will be maintained so that all versions of the data ever published will be available for science if needed. The loading process will be completely automated using a combination of VB and DTS scripts and SQL stored procedures, and an admin web interface will be provided to the **Load Monitor** which controls the entire loading process. Data will be first converted from FITS to CSV (comma-separated values) before being transferred from Linux to Windows.

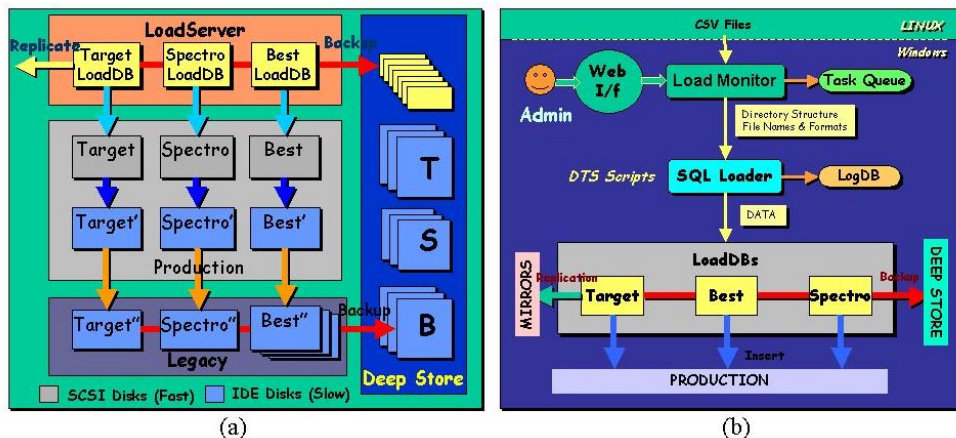


Figure 1. (a) Production archive components, (b) loading data flow.

### 3.2. Current Hardware Plan

The proposed hardware plan for the master DR1 site at FermiLab reflects the function that each copy of the archive must provide, but it also makes the most effective use of the existing SDSS hardware resources at FermiLab. Table 2 shows the plan for the various DR1 components.

Table 2. Hardware for DR1 Archive.

	<i>Load Servers</i>	<i>Production Servers</i>	<i>Legacy</i>	<i>Deep Store</i>
<i>Priority</i>	Hi Perf/Lo Capacity	Hi Perf/Hi Capacity	Lo Perf/Hi Capacity	Hi Capacity
<i>Vendor</i>	Dell PowerEdge 4600	Intel E7500	Dell PowerEdge 4400	Enstore
<i>CPU</i>	Dual 2.6 GHz Xeon	Dual 2.4 GHz Xeon	Dual 1 GHz PIII Xeon	
<i>Memory</i>	2-4 Gb	4-8-12 Gb	2 Gb	
<i>Disks</i>	2 × 120MB/s SCSI software RAID	2 × 8 × 160GB 3ware 7500 ATA RAID	14 × 73GB SCSI across 4 ultra 160 channels	Tape Silo
<i>Options</i>	Parallelize with multiple servers	Cluster of these for redundancy (warm sp.)	May use slower IDE Disks	

## 4. Databases

In January 2002, the SDSS collaboration made the decision to migrate to Microsoft SQL Server as the DB engine based on our dissatisfaction with Objectivity/DB's features and performance (Thakar et al. 2002). SQL Server meets our performance needs much better and offers the full power of SQL to the database users. SQL Server is also known for its self-optimizing capabilities, and provides a rich set of optimization and indexing options. We have further significantly augmented the power of SQL Server by adding the HTM spatial index (Kunzst et al. 2001) to it along with a pre-computed neighbors table that enables fast spatial lookups and proximity searches of the data. Additional features like

built-in aggregate functions, extensive stored procedures and functions, and indexed and partitioned views of the data make SQL Server a much better choice for data mining.

As the size of the SDSS data grows with future releases, we will be experimenting with more advanced SQL Server performance enhancements, such as horizontal partitioning and distributed partition views (DPVs). We are also developing a plan to provide load-sharing with a cluster of DR1 copies rather than a single copy. This kills two birds with one stone - it also removes the need to have warm spares of the databases, since each copy can serve as a warm spare.

## 5. Data Mining Tools

There will be a single web access point to all DR1 data. Our data mining tools will be integrated into a VO-ready framework of hierarchical Web Services (Szalay et al. 2002).

### 5.1. Catalog Access

Access to catalog data will be via a variety of tools for different levels of users.

1. The **SkyServer** is a web front end that provides search, navigate and explore tools, and is aimed at the public and casual astronomy users.
2. The **sdssQA** is a portable Java client that sends HTTP SOAP requests to the database, and is meant for serious users with complex queries.
3. An **Emacs** interface (.el file) to submit SQL directly to the databases.
4. **SkyCL** is a python command-line interface for submitting SQL queries.
5. **SkyQuery** is a distributed query and cross-matching service implemented via hierarchical Web Services (see Budavari et al. 2003).

### 5.2. Raw data

1. The **Data Archive Server (DAS)** will be a no-frills web page for downloading raw data files (FITS) for the various raw data products.
2. A Web Form or Web Service interface to upload results of SQL queries to the DAS and retrieve the corresponding raw images and spectra.
3. An **Image Cutout Service** (jpeg and FITS/VOTable) which will be implemented as a Web Service.

## References

- Budavari, T., et al. 2003, this volume, 31
- Kunszt, P. Z., Szalay, A. S., and Thakar, A. 2001, Mining the Sky: Proc. of the MPA/ESO/MPE workshop, Garching, A.J.Banday, S. Zaroubi, M. Bartelmann (ed.), (Springer-Verlag Berlin Heidelberg), 631.
- Szalay, A. S., et al. 2002, Proceedings of SPIE "Astronomical Telescopes and Instrumentation", 4846, in press.
- Thakar, A. R., et al. 2002, in ASP Conf. Ser., Vol. 281, Astronomical Data Analysis Software and Systems XI, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: ASP), 112