

Data to Science

Data Management: Archiving, Standards, and Accessibility

Bryan J. Field, Kathy Turner, and Christopher Jackson (detailee)

Office of High-Energy Physics, Cosmic Frontier

NSF Headquarters, September 18/19, 2023

bryan.field@science.doe.gov

kathy.turner@science.doe.gov

christopher.Jackson@science.doe.gov



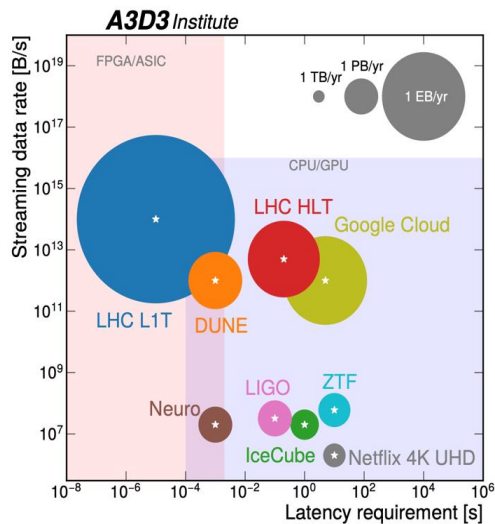
U.S. DEPARTMENT OF
ENERGY

Office of
Science




The Big Data Frontier: LHC Case Study

- In **2018** – during the final year of the 2nd season of LHC running (Run 2) – more than **115 petabytes of data in total**, including about 88 petabytes of LHC data, were recorded on tape, with a record peak of **15.8 petabytes** in November 2018
- The “Tier 0” center at CERN first reconstructs the full collision events and analysts start to look for patterns; but the **data has a long way to go**. Once CERN has made a primary backup of the data it is then sent to large “Tier 1” computer centers in seven locations around the world: in France, Germany, Italy, Spain, Taiwan, the UK and the US. Here events are reconstructed again, using information from the experiment to improve calculations using refined calibration constants. Then onto several “Tier 2” centers...



The Data

The Rate



Optimizing the use of Supercomputing Facilities

- High-Energy Physics Center for Computational Excellence (HEP-CCE) is a DOE High-Energy Physics (HEP) collaboration of 4 national labs, six HEP experiments, and 4 HPC centers
 - Argonne Lab, Brookhaven Lab, Fermilab, Lawrence Berkeley Lab, ATLAS, CMB-S4, CMS, DESI, DUNE, LSST-DESC, Argonne Leadership Computing Facility (ALCF), Energy Sciences Network (ESnet), Oak Ridge Leadership Computing Facility (OLCF), National Energy Research Scientific Computing Center (NERSC)
- HEP-CCE addresses the development and implementation of HEP scientific applications on next-generation computing, storage, and networking systems.
- A current focus of the HEP-CCE is to develop common strategies to efficiently run HEP software applications on pre-exascale and exascale high-performance computing systems that will shortly be deployed at ASCR computing facilities.
- The Department of Energy has a long and proud history of funding projects that have successfully created, analyzed, archived, and served some of the largest scientific datasets ever created.

Alpha Magnetic Spectrometer ^{1/2}




U.S. DEPARTMENT OF
ENERGY

Office of
Science



- AMS is an international collaboration involving 44 institutions from America, Europe and Asia sponsored by US **DOE and NASA**. DOE, through the AMS-MIT team, is responsible for detector operations, data processing, setting the physics analysis topics and publications.
- AMS was launched by the Space Shuttle Endeavour (STS-134) to the ISS on 16-May-2011 and installed on the upper Payload Attach Point (S3) on 19-May-2011. It is expected to continue through the lifetime of the ISS currently scheduled through 2024 and possibly to 2028.
- Members of the AMS collaboration have unrestricted access to the data.
- A raw event is the collection of the information read out from the AMS detector's electronics after the latter are triggered by particle(s) crossing the AMS sensitive volumes. Raw events are logically grouped in sequences called **runs**. One run typically includes all events taken during a quarter of ISS Low Earth Orbit movements, i.e., about **23 minutes** and contains on average **700,000 events**.



Alpha Magnetic Spectrometer ^{2/2}



U.S. DEPARTMENT OF
ENERGY

Office of
Science

- Data frames are transmitted via the TDRS satellites to the ground at White Sands, NM, and then relayed to NASA Marshall Space Flight Center (MSFC) in Alabama. At MSFC the framed data are archived (under NASA responsibility, for the lifetime of the ISS program) and simultaneously transferred to AMS relay computers at MSFC using the User Datagram Protocol (UDP). From the relay computers the data is copied into the POCC at CERN and made available via the Network File Service (NFS) to the AMS Science Operation Center (SOC), also located at CERN.
- Raw data are collected at **10 Mb/sec** or **40 TB/year**. Currently, the combined dataset (raw, reconstructed and simulated data) is approximately **22.5 PB**.
- All the finalized AMS data are published in Physical Review Letters and readily accessible from the journal web site (supplemental material).
- According to the NASA-DOE agreement, the complete data will be released to the public one year after the termination of the experiment.



AMS has collected

225,657,950,810

cosmic ray events

Last update: September 12, 2023, 3:55 PM

Fermi Gamma-ray Space Telescope/LAT ^{1/2}



- NASA is responsible for the FGST (Fermi) mission. DOE led the fabrication (managed by SLAC) of the **Large Area Telescope (LAT)**, the primary instrument on Fermi, in partnership with NASA and with contributions from international partners
- The LAT detects gamma-rays of energies from 20 MeV to greater than 300 GeV.
- After launch, DOE led the LAT Instrument Science Operations Center at SLAC; in the last few years, the ISOC is carrying out the critical roles that need to be done at SLAC (the rest has moved to GSFC).
- The ISOC supports LAT operations, **LAT data processing**, hosting and archiving, and LAT data analysis, in conjunction with NASA and members of the LAT Collaboration.
- The LAT **collects data continuously**; data are stored temporarily on the spacecraft between highspeed downlinks to the ground, occurring about **15 times per day**.





- The LAT sends about **15 GB of raw data to the ground each day**. Level 1 processing of the raw data at SLAC reduces the data volume by a factor of about 100. About **200 MB of processed photon data** are then delivered back to NASA every day. **SLAC currently holds about 5 PB of LAT data.**
- Data analysis tools are provided with data to members of the LAT Collaboration and to the public. Those tools can be installed on user computers, and data can be downloaded from servers for local analysis. LAT Collaboration members can also use installed tools and data on SLAC computers.
- All photon data are made public, as quickly as possible after collection on the LAT, generally **within a few hours of being taken**. The LAT Collaboration has access to all LAT data, including and beyond the photon data. Other data products made public are measurements of detected celestial transient sources, light curves of known bright variable sources, and detected source catalogs.
- Public data releases are continuous, released as quickly as possible after processing, or publication for source catalogs.

LuSEE-Night^{1/2}



- LuSEE-Night is a DOE-NASA collaboration led by DOE's Brookhaven National Lab (BNL) and the Space Science Laboratory (SSL) at UC Berkeley. BNL is the DOE lead institution, SSL is the overall lead institution.
- DOE is responsible for the main instrument and sponsors some equipment (panels, battery, communications subsystem, and battery management). NASA is responsible for integration, thermals, and overall systems engineering.
- LuSEE-Night will be delivered to the far side of the Moon on a future Commercial Lunar Payload Services (CLPS) flight. The vendor is Firefly Aerospace, and the launch is currently scheduled for late 2025.
- Data is limited by the available download volumes from the Moon and are very modest by modern standards.
- Raw data will be about **6-10 GB per lunar cycle**, so over 18 months, the total dataset will still be under **200 GB**. Collaboration will be about 20 people when data collection is complete. They will mostly use their own compute resources, there is no plan for a dedicated allocation given the modest data volumes.

- NASA requires data to be public 6 months after **data taking**. Collaboration plans regular 6-month data releases. Initially these will be limited to basic data products, at the end of mission, the collaboration will release everything, including simulations.
- During the proprietary period, data will be otherwise limited to Science Collaboration members.
- The modest data volume will be served via the web. Data will be available on the website for the foreseeable future due to its limited size.



CMB-S4 ^{1/2} (Project's draft plan)



Office of
Science

- CMB-S4 is anticipated to be a joint DOE/NSF project, with possible additional contributions from international and private partners. It is currently in Conceptual Design, with DOE CD-o approval and is awaiting entry to the MREFC pool at NSF.
- DOE is expected to be responsible for the Detectors, Readout, Modules, Small Aperture Telescopes, and Data Acquisition & Control
- NSF is expected to be responsible for the Large Aperture Telescopes, and the Chile and South Pole Sites.
- DOE and NSF are jointly responsible for the Data Management. **These comments are about the draft plan for data!**
- Data is expected to be produced at about **1.2Gbps** at the South Pole and **2.6Gbps** in Chile, for around 7 years, eventually totaling some **40PB**.
- The project will be responsible for reducing the raw data to well-characterized single frequency maps. This is expected to use a mixture of allocated resources at NERSC and individual members' own computing resources.
- During the proprietary period, all CMB-S4 collaboration members will have access to the raw data and to intermediate data products. CMB-S4 is in the process of determining buy-in from international partners.
- The details of the proprietary period have not been finalized yet.





- The main mode for the distribution of data to the collaboration will be via **the primary data center**, where all collaboration members will also be eligible to have user accounts and have access to significant compute resources. Given its **volume**, they typically expect users to bring their code to the time-domain data, and only to transfer map-domain data to their local computing resources (though these can obviously be processed at NERSC too).
- There have been no public data releases yet, though they do hope to release the various Data Challenge dataset once they have been analyzed by the collaboration. The first of these - Data Challenge 0 - **could be released in 2024**.
- The data is expected to be made available at NERSC, via the dedicated CMB-S4 data portal (<https://data.cmb-s4.org/>), and eventually potentially through the NASA LAMBDA archive.
- **Real-time processing for rapid transient detection** alerts imposes stringent demands, especially at the South Pole. Should the bandwidth to the South Pole increase sufficiently to stream data to the US then the project's processing model for those data would change significantly.

Dark Energy Survey (DES) ^{1/2}



U.S. DEPARTMENT OF
ENERGY

Office of
Science



- DES is an international collaboration led by NSF and DOE. During the fabrication phase, DOE funded the DECam (led by FNAL) and NSF funded the telescope upgrades (led by CTIO) and the data management system. DOE and NSF supported the DECam and telescope operations and both supported the data management during the operations phase.
- Data taking was from 2012 to 2019. The collaboration is in the midst of the final data processing and Key Project analyses.
- Data was processed at National Center for Supercomputing Applications (NCSA), UIUC Campus Cluster, Fermilab, NERSC using pipeline code developed by collaborators and computing professionals.
- All of the data, observations and images & catalogs made from them, value-added information from subsequent analysis pipelines, including (but not limited to) those that are necessary for Key Project cosmology results, are made available to the public. The timescale depends on the specific data product.

Dark Energy Survey (DES) ^{2/2}




U.S. DEPARTMENT OF
ENERGY

Office of
Science



- A full night of observations comprised approximately 400 images (including calibrations). Each image is about **1 GB** in size. There would be typically 20,000 images (including calibrations) per year. That totals about **120 TB** of raw data.
- This raw data is processed several times both one image at a time, and then combining the images. Cosmology data products are derived from these catalogs and images. Final data products are about **1.5 PB**.
- Public data releases DR1 and DR2, based on the Y3 and Y6 data, were released by DES on-schedule in January 2018 and January 2021, respectively. DES has released value-added cosmology data products upon the acceptance by the journals to which the results were submitted.
- At present NCSA serves DES data products of all kinds (including cosmology data) to both the public and to the DES Collaboration. This data is accessed via the web using standard database search and access techniques.
- It is expected that all data will be served at Fermilab ultimately.

Dark Energy Spectroscopic Instrument (DESI) ^{1/2}



- DOE leads the DESI project and operations, with contributions from NSF, and US and international institutions. DOE “leases” the Mayall telescope from NSF.
- DESI operations started in December 2020, with their survey validation (SV) phase, followed by the 5-year survey. The SV data was included in their early data release (EDR).
- Although the raw data is relatively small, **50-100 GB/night**, the combined raw + processed + mock data and scientific working group output adds up to approximately **1 PB/year** growth
- The data are processed at the National Energy Research Scientific Computing Center (NERSC) using a real-time queue for prompt processing of data as they arrive from the telescope, and the regular batch queue for yearly reprocessing of all data with the latest algorithms and calibrations.
- The majority of DESI analyses are performed at NERSC, using time granted through the annual allocation process, which covers both the needs for core data processing (10-20% of the allocation) and the scientific analyses and supporting DESI-specific mocks/simulations (80-90% of the allocation).



Dark Energy Spectroscopic Instrument (DESI) ^{2/2}



U.S. DEPARTMENT OF
ENERGY

Office of
Science

- The collaboration has immediate access to the data at NERSC. DESI plans public data releases every ~2 years, following a 1–2-year proprietary period.
- **Early Data Release (EDR)**, publicly released in June 2023 includes spectra and redshifts for 1.2 million galaxies and quasars, as well as observations of nearly half a million stars in our own Milky Way galaxy. The dataset is approximately **80 TB** and represented about **2% of the full DESI dataset** at the time (**4 PB**).
- The EDR in 2023 covered “Survey Validation” data prior to the DESI Main Survey; future data releases are planned for the year 1 dataset (public release in 2024), and the year 3 and year 5 (final) datasets.
- Data are primarily served through NERSC, including direct file access for DOE collaborators with a NERSC account; web download; and Globus for bulk transfers. A subset of the data are also available via the NOIRLab Astro Data Lab.
- After DESI operations have finished, they will continue to serve the data from NERSC through the Cosmology Data Repository which also hosts other cosmologically useful public datasets for joint analysis at NERSC.



Rubin Observatory ^{1/2}



Office of
Science

- The Rubin Observatory is an NSF/DOE partnership, with NSF as the lead agency, and with private and international contributions. DOE roles are led by SLAC and NSF by AURA. During construction/commissioning, DOE is responsible for the camera (LSSTCam) and NSF for the telescope facility and data management system. DOE and NSF are splitting the operations costs ~ 50/50, with DOE primarily responsible for LSSTCam and the US Data Facility (USDF) at SLAC
- The 10-year survey is projected to start in mid 2025. Commissioning has already begun.
- Expected to produce **20 TB/night** off the telescope, about **30 PB/year** after processing. After 10 years, we're looking at **300 PB**.
- The USDF at SLAC leads the data processing and management, with public alerts within 60 seconds and annual calibrated image catalogs. The USDF will orchestrate data processing in 3 centers (US, UK, France) and archive all the results using LHC tools for workflow (PanDA) and data (Rucio) management
- The plan to host 5-10k science users in the Google cloud, serving up data from the USDF via the Rubin Science Platform (RSP). Current cloud resources are not large enough for all analyses; additional resources will come from some 12 Independent Data Access Centers.
- Researchers in the US, Chile and international with data rights, gained via in-kind contributions, have immediate data access. After two years, prior annual release data sets will be made public worldwide.



- After the proprietary period, the plan to do annual releases where all the data taken to date is reprocessed each time. Public access to a manageable subset of the LSST data will be supported through partnership with external data centers (from the international in-kind program) and/or NOIRLab.
- Data will be served out of the USDF primarily to the cloud, by queries to a database of objects (stars and galaxies), cutouts around sources/areas of interest, and bulk image downloads.
- The USDF will provide these functions through end of survey and for 2 years thereafter. It is expected that NOIRLab, as lead lab, will take responsibility for long term archiving of the LSST data.



Legacy Survey of Space and Time/ Dark Energy Science Collaboration (LSST/DESC) ^{1/2}



U.S. DEPARTMENT OF
ENERGY

Office of
Science



- DESC is a primarily DOE-supported **collaboration** organized for dark energy science with the DOE + NSF-supported Rubin Observatory LSST survey. DESC also receives in-kind support from IN2P3 (France), STFC (UK), and individual institutions.
- DESC was organized in 2012 and started receiving DOE Operations funding in 2017.
- Rubin-LSST survey will begin in 2025 and the annual reductions will be done by Rubin. Rubin dataset will be **100s of Peta-bytes. DESC will reprocess subsets of the data.**
- NERSC is the primary computing resource for DESC (Perlmutter CPU and GPU). DESC will process its value-added catalog data and simulation data and reprocess a subset of raw Rubin data for systematics studies. DESC has additional computing resources available in France and UK.
- Members have a default NERSC allocation of computing time for studies and propose to the DESC Computing Coordinator for larger allocations.
- DESC use of LSST survey data will be governed by the Rubin Data Policy. DESC will release its value-added data products consistent with Rubin policy.

Legacy Survey of Space and Time/ Dark Energy Science Collaboration (LSST/DESC) ^{2/2}



U.S. DEPARTMENT OF
ENERGY

Office of
Science



- Data Challenge 2 (300 deg² simulation of the LSST survey) images and catalogs (a reprocessing by Rubin was released as Data Preview 0.2) have been released. Photometric LSST Astronomical Time series Classification Challenge (PLAsTiCC) and Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC) simulated time domain data are also currently available.
- DESC serves its public data via its NERSC data portal and Community File System (cfs) portal.
- NERSC serves as our primary data storage and archive facility. Long-term archiving is planned to be done via NERSC HPSS tape storage.



U.S. DEPARTMENT OF
ENERGY

Office of
Science