




EOSDIS Handbook

July 2016



1. Introduction.....	3
1.1. Purpose.....	3
1.2. Scope.....	3
1.3. Document Organization.....	3
2. The Motivation for EOSDIS	4
2.1. Origin of EOSDIS.....	4
2.2. Standardization and Interoperability	4
2.3. Managing Data from Heterogeneous Earth Observing Missions.....	4
2.3.1. Description of EOSDIS Context.....	4
2.3.2. Interfaces within the EOSDIS System.....	6
2.3.3. User Community	7
3. Enterprise Viewpoint	8
3.1. Architecture	8
3.1.1. Distributed Active Archive Centers (DAACs)	8
3.1.2. Science Investigator-led Processing Systems (SIPS)	9
3.1.3. Networks.....	9
3.1.4. EOSDIS Common Services	12
3.1.5. Metrics and Monitoring	21
3.2. Standards and Conventions Organizations	21
3.2.1. ESDIS Standards Office (ESO)	21
3.2.2. ISO19115 Metadata Standards	21
3.2.3. Committee on Earth Observation Satellites (CEOS)	22
3.2.4. Group on Earth Observations (GEO)	22
3.3. Key System Functions / Business Use Cases.....	22
3.3.1. Ingest.....	22
3.3.2. Science Data Processing.....	23
3.3.3. Archive.....	25
3.3.4. Discovery	25
3.3.5. Data Access and Ordering	27
3.3.6. User Identification	27
3.3.7. User Services	28
3.3.8. Metrics	28
3.3.9. Near Real Time Support	29
3.4. Related Initiatives.....	29
3.4.1. Federation of Earth Science Information Partners (ESIP)	9
3.4.2. Earth Science Data Systems Working Group (ESDSWG)	30
3.4.3. Big Earth Data Initiative (BEDI)	30

4. Information Viewpoint.....	31
4.1. Earth Science Information Models.....	31
4.1.1. UMM and Relationship to ISO metadata model.....	31
4.1.2. Collection metadata.....	32
4.1.3. Granule Metadata.....	33
4.1.4. Browse model.....	33
4.1.5. Usage model.....	34
4.1.6. Processing Levels.....	34
4.1.7. Data Model.....	35
4.2. Data Format Standards.....	35
5. Service Viewpoint.....	37
5.1. Discovery Services.....	37
5.1.1. CMR.....	37
5.1.2. OpenSearch.....	38
5.1.3. CWIC.....	38
5.2. Online Data Access and Transformation Services.....	39
5.3. Identity Management Services.....	39
5.3.1. User Identification.....	40
6. Outlook and Conclusion.....	41
7. Appendices.....	42
7.1. Adding new data to the EOSDIS Collection.....	42
7.2. Acronyms.....	43

1. Introduction

1.1. Purpose

The purpose of this handbook is to document the overall architecture of the Earth Observing System Data and Information System (EOSDIS) from several viewpoints:

- Enterprise Viewpoint describing the context surrounding EOSDIS, an architectural overview, important EOSDIS standards, and finally key components.
- Information Viewpoint documenting key data models and representations used in EOSDIS.
- Service Viewpoint providing introductions to key technical interfaces and services.

Readers should be able to get a high level view of what EOSDIS does and how it accomplishes it.

1.2. Scope

This document covers the high-level organization of EOSDIS into constituent Distributed Active Archive Centers (DAACs), Science Investigator-led Processing Systems (SIPS) and EOSDIS common elements that apply across the system. Information about individual DAACs will be compiled in a second volume of this handbook. This handbook covers primarily the functional aspects of EOSDIS that are visible to external interfaces and users, underpinned with some discussion of internal architecture. However, EOSDIS's business processes are not covered.

1.3. Document Organization

This document is organized to provide a top down view of NASA's EOSDIS. Section 2 explains the origin of EOSDIS, the high level handling of Earth Observation data, and EOSDIS' role in the broader context of global science communities. Section 3 describes EOSDIS as an enterprise; documenting how data is archived, core EOSDIS services and systems, networks, and key capabilities. Section 4 covers the information viewpoint - how is data described, made available for discovery and use, and what data formats are used. Section 5 explains EOSDIS services from a technical perspective, providing a high level overview of key APIs into EOSDIS holdings, their role in the enterprise, and where to find more information on using them. Finally, Section 6 offers an outlook on the future of EOSDIS and plans in progress to make one of the world's largest holdings of Earth Observation data even more readily accessible.

2. The Motivation for EOSDIS

2.1. Origin of EOSDIS

EOSDIS began as a data system to support NASA's Earth Observing System (EOS), a set of Earth Observing satellites covering a wide range of Earth science disciplines. The original concept consisted of a constellation of DAACs, each with expertise in a particular discipline and tasked with supporting the data in that discipline. The DAACs were linked together through an interoperable Information Management System. While the number of DAACs has fluctuated slightly from time to time, the system remains largely intact today. This formulation produced a system that was adaptable to the evolving needs of a variety of user communities, yet still provided a largely consistent level of service across the system.

While much of the focus of EOSDIS has been on satellite data, the DAACs have also been repositories from the beginning of data from field campaigns, airborne data, and synthesized model data.

2.2. Standardization and Interoperability

From its beginning, EOSDIS was planned to support interoperability among its DAACs, demonstrating catalog interoperability at both data collection and file level as early as 1994. A common user interface supported data search and order across the DAACs, so that a user did not need to know which DAAC held a given dataset. Interoperability was further enhanced by adoption of the Hierarchical Data Format (HDF) as a recommended standard for science data products within EOSDIS. EOSDIS use of standards has evolved over the last two decades to embrace more community standards, which are explained in more detail later in this document.

2.3. Managing Data from Heterogeneous Earth Observing Missions

The Earth observation data supported by EOSDIS runs the gamut from solar irradiance to oceanographic to atmospheric to land surface and subsurface. The instruments producing these data include radiometers, lidars, radars, interferometers, sounders, active and passive microwave and gravimeters, among others, borne on a variety of satellites. In addition, EOSDIS supports related field data, including in situ stations and airborne instrument simulators. EOSDIS also contains some model output, particular assimilations of the satellite measurements. This heterogeneity is handled by a combination of agile, adaptive DAACs together with a core set of practices, standards and facilities (described later in this document) to abstract the complexity.

2.3.1. Description of EOSDIS Context

As shown in Figure 1, data are generally downlinked from the satellite by one of two paths. In the first path, used by the Terra satellite, data are relayed to a geosynchronous satellite called the Tracking and Data Relay Satellite (TDRS), which in turn downlinks the data to a facility at White Sands, New Mexico. From there, data are transmitted to the EOS Data and Operations System (EDOS) at the Goddard Space Flight Center in Greenbelt, MD. Many other Earth Science satellites downlink their data to Polar Ground Stations for forwarding directly to EDOS.

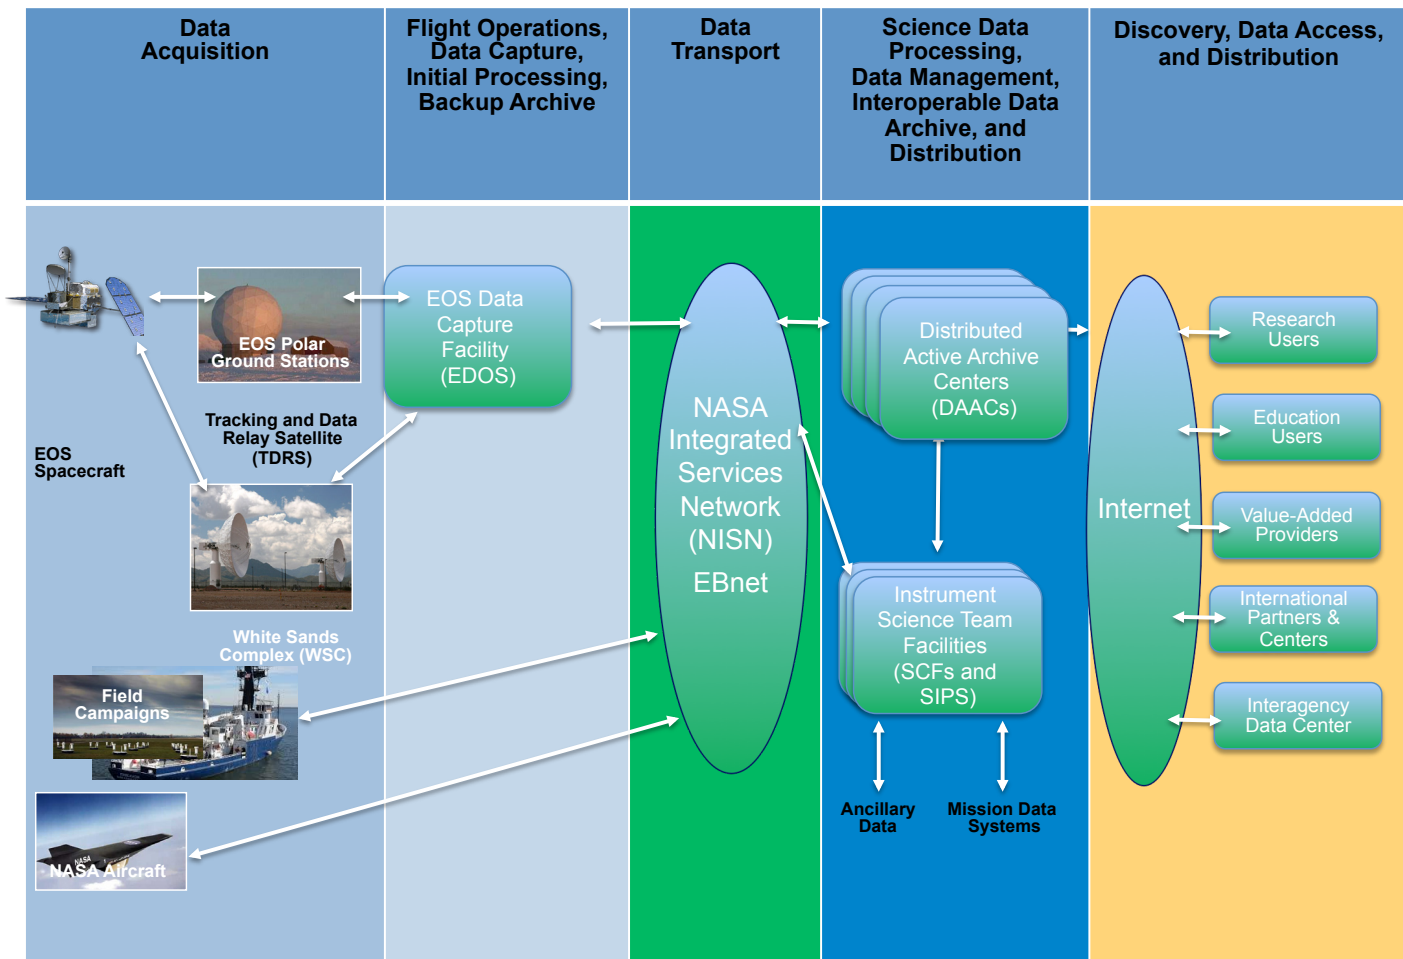


Figure 1: EOSDIS data lifecycle from acquisition to distribution

EDOS processes the rate buffered data, reordering packets and removing duplicates as necessary, into Level 0 Production Data Set (PDS) files, which are then sent to DAACs, who then relay the data to SIPS by subscription. The SIPS process the raw PDS data into geophysical variables and send them to the DAACs for archive and distribution. In some cases, the DAACs themselves execute the retrieval processing on behalf of the science team, using processing algorithm software submitted to them by the science team. The DAACs are then responsible for distributing the science data to research users and other communities.

In airborne missions, the instrument teams perform the tasks that the SIPSs do for satellite missions: collect Level 0 data and process these products into higher-level products. Data as collected from the on-board instrument is designated the level 0 product. Products produced at the instrument teams' facilities using investigator-provided systems and software are sent to appropriate DAACs for archiving and distribution to general users. Level 0 data products and ancillary data that begin the processing sequence are stored at the DAACs, as are higher-level products that are created from the Level 0 data. The instrument teams may reprocess standard data products as needed. Airborne mission interface agreements may be more informal than the Interface Control Documents used for satellite mission data, but at a minimum they establish what data is to be transferred, how data transfer is accomplished, and what information needs to accompany the data to ensure accurate receipt and accountability by the data center.

A field campaign may result in a very large number of diverse data products. The Large Scale Biosphere-Atmosphere Experiment in Amazonia campaign, as an example, has generated over 300 data products, ranging in size from a few granules that are each smaller than a megabyte to complex data products with thousands of granules, with differing formats and scientific content, totaling into the gigabyte range. For field campaigns, raw data from field-deployed instruments are captured by a variety of specialized hardware and software components, such as data loggers (e.g. Campbell Scientific products). Field campaign data also includes observations recorded by field personnel (possibly in paper notebooks) and may result from measurements made by laboratory instruments on samples collected in the field.

The processing of this raw data is the responsibility of the field project personnel, and the methods employed are highly variable, given the extremely diverse types of measurements required to address hypotheses of an overall NASA field campaign. The initial data processing done by these teams is highly diverse, meeting the needs of the individual science teams. Standards are not widely implemented since measurements and processes are campaign specific. The observational campaigns may range from measuring carbon dioxide flux using eddy covariance flux towers, to measuring soil respiration, to making biomass estimates by measuring tree diameter at chest height using a tape measure. Field data are often kept on individual computing resources (desktop, project servers) and not at a DAAC. The degree of centralization of storage is a function of how the field campaign project is organized, and an area of evolving technology and lessons learned through practice. The science processing is typically done by individual scientists or science teams as they begin analyzing the data to examine spatial and temporal trends. The initial processing and science processing are often done by the same individuals. In some cases, the raw data is inserted into Excel spreadsheets and spreadsheet formulas are used to generate Level 1 data in that same spreadsheet. This data is eventually moved to more standard formats, often by a DAAC, and stored.

EOSDIS is also home to a wide collection of socioeconomic data. Socioeconomic data are available in a variety of formats and must often be converted, checked carefully for errors, gaps, and inconsistencies, and filtered to obtain relevant, usable data. Often data must be geo-referenced, i.e., individual data items must be associated with spatially defined points, lines, or areas (e.g., centroids, borders, or administrative units). This may require use of a gazetteer to associate place names with geographic coordinates, or spatial processing to identify appropriate buffers around coordinates, taking into account coastlines, borders, or other geographic features. Geo-referenced data from multiple sources may also need to be reprojected to ensure spatial consistency and accuracy. The socioeconomic data is also archived and collected for use across the EOSDIS program at the Socioeconomic Data and Applications Center (SEDAC).

2.3.2. Interfaces within the EOSDIS System

Data exchange with the variety of missions supported by EOSDIS can be a challenge. Since the turn of the century, a standard exchange mechanism has been used with nearly all of the satellite missions with which EOSDIS interfaces. This interface is typically implemented by the receiving system polling a directory, either a remote one or a local one, for the arrival of an American Standard Code for Information Interchange (ASCII) Product Delivery Record (PDR) which details the data that are ready for transfer, including the data file location, size and checksum. When a delivery record from the sending system is detected, the data are fetched (if on a remote system) and ingested, that is, inventoried and made available to the rest of the system. Product Delivery Records are standardized (albeit with minor variations); the format is documented and controlled via Interface Control Documents (ICD) among organizations.

2.3.3. User Community

NASA’s Earth observation data are nearly all available for complete public access after the post-launch checkout, in keeping with the NASA Earth Science Data Policy . (Data from other agencies or nations may have restrictions negotiated with the provider, however.) The primary target community is researchers in Earth science and related fields. Although NASA is a research agency and not an operational agency, as shown in Table 1 a significant number of application users utilize NASA data in a variety of applications and decision support systems. NASA’s data and related services and tools also find use within the Education and Outreach communities, though this is not a primary focus of the data. Note that “users” may be either human or machine, in the sense of other software systems incorporating NASA data and services.

User Community Type	Human	Machine
Research (primary focus)	Remote sensing researchers Interdisciplinary researchers Modelers	Models
Applications	Policy makers Businesses	Decision Support Systems
Education and Outreach	Teachers Student Museums and other outreach	Teaching tools Outreach tools

Table 1: EOSDIS user community types and points of interaction

NASA cooperates extensively with other agencies, both here and abroad.

3. Enterprise Viewpoint

3.1. Architecture

EOSDIS is a true system-of-systems with a sophisticated architecture to enable the movement of data from acquisition to processing to archival and distribution. This section describes the overall organization of EOSDIS and the core systems that make this enterprise function.

3.1.1. Distributed Active Archive Centers (DAACs)

The actual archiving and distribution of data in EOSDIS is done by a distributed system of twelve DAACs. Each DAAC supports data and associated user communities from a small number of specific science disciplines and the associated user communities as shown in Figure 2. While seven of the DAACs are operated by NASA centers, two are operated by other federal agencies (Department of Energy and United States Geological Survey) and three by universities.

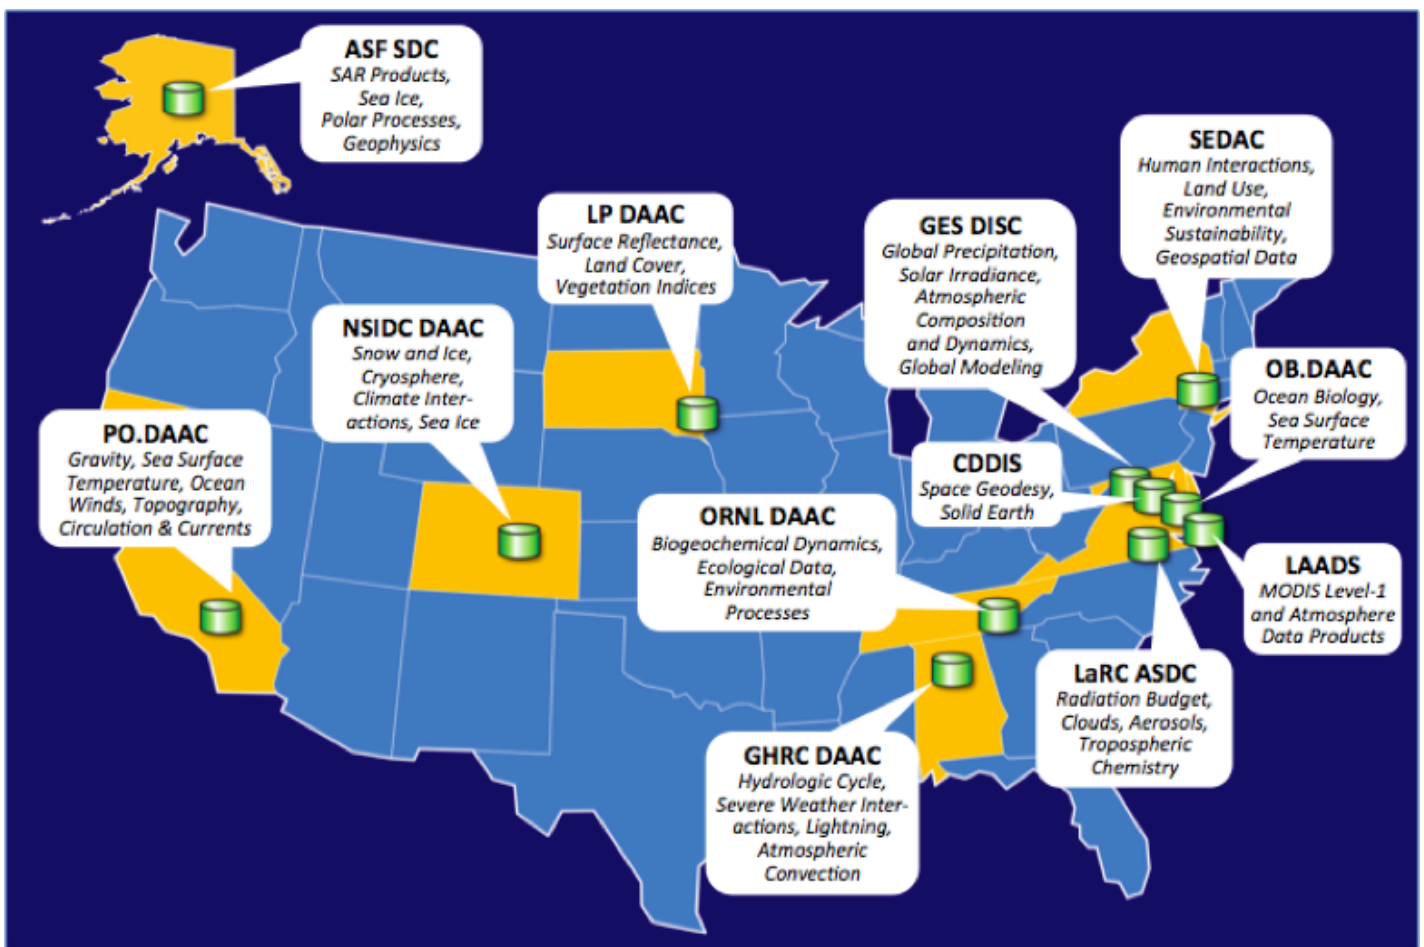


Figure 2: EOSDIS DAACs are geographically distributed and support diverse science disciplines

The DAACs share a common set of requirements with respect to archiving and distributing data, as well as interfacing to EOSDIS common services such as the metadata clearinghouse and metrics systems. However, the individual DAACs have evolved different architectures to best support the missions, data types and user communities assigned to them. EOSDIS is thus a loosely coupled system: the DAACs do not depend on the

architecture of other DAACs. Data and information are exchanged through a small number of well-defined and controlled interfaces, and changes in one component rarely force changes in other components, except in cases where evolution of standards and compliance thereto are concerned.

3.1.2. Science Investigator-led Processing Systems (SIPS)

Most of the EOS standard products are produced at facilities under the direct control of the instrument Principal Investigators/Team Leaders or their designees. These facilities are referred to as Science Investigator-led Processing Systems (SIPS). The SIPS are geographically distributed across the United States and are generally, but not necessarily, collocated with the PIs/TLS' Scientific Computing Facilities. Products produced at the SIPS using investigator-provided systems and software are sent to appropriate DAACs for archival and distribution. Level 0 Data Products and Ancillary Data that begin the processing sequence are stored at the DAACs and retrieved by the SIPS.

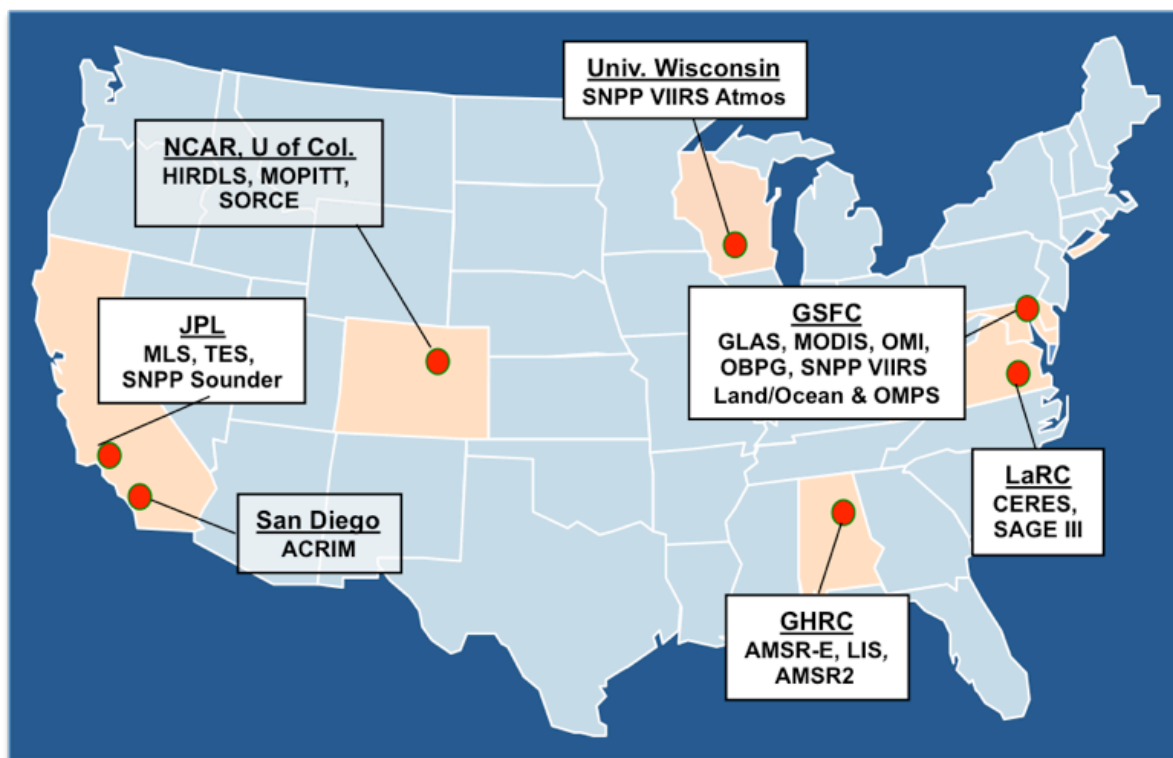


Figure 3. Map of Science-investigator Led Processing Systems (SIPS).

3.1.3. Networks

EOSDIS is a comprehensive data and information system designed to perform a wide variety of functions in support of a varied national and international user community.

Mission operations, managed by the Earth Science Mission Operations (ESMO) Project, coordinate communications through the Space and Ground Network facilities of the Tracking and Data Relay Satellite System (TDRSS) and the ground stations.

Science operations, managed by the Earth Science Data and Information System (ESDIS) Project, are performed within a distributed system of many interconnected nodes with specific science operations responsibilities for

production, archiving, and distribution of Earth science data products from Science Investigator-led Processing Systems (SIPS) and discipline specific EOSDIS data centers.

NASA network capabilities transport the data to the Mission Operations facilities and Science Operations facilities (Figure 4). Mission Operations are managed by the ESMO and consist of instrument/platform operations and data capture. Science Operations, managed by ESDIS, consists of science processing, archive and distribution to science teams and end users.

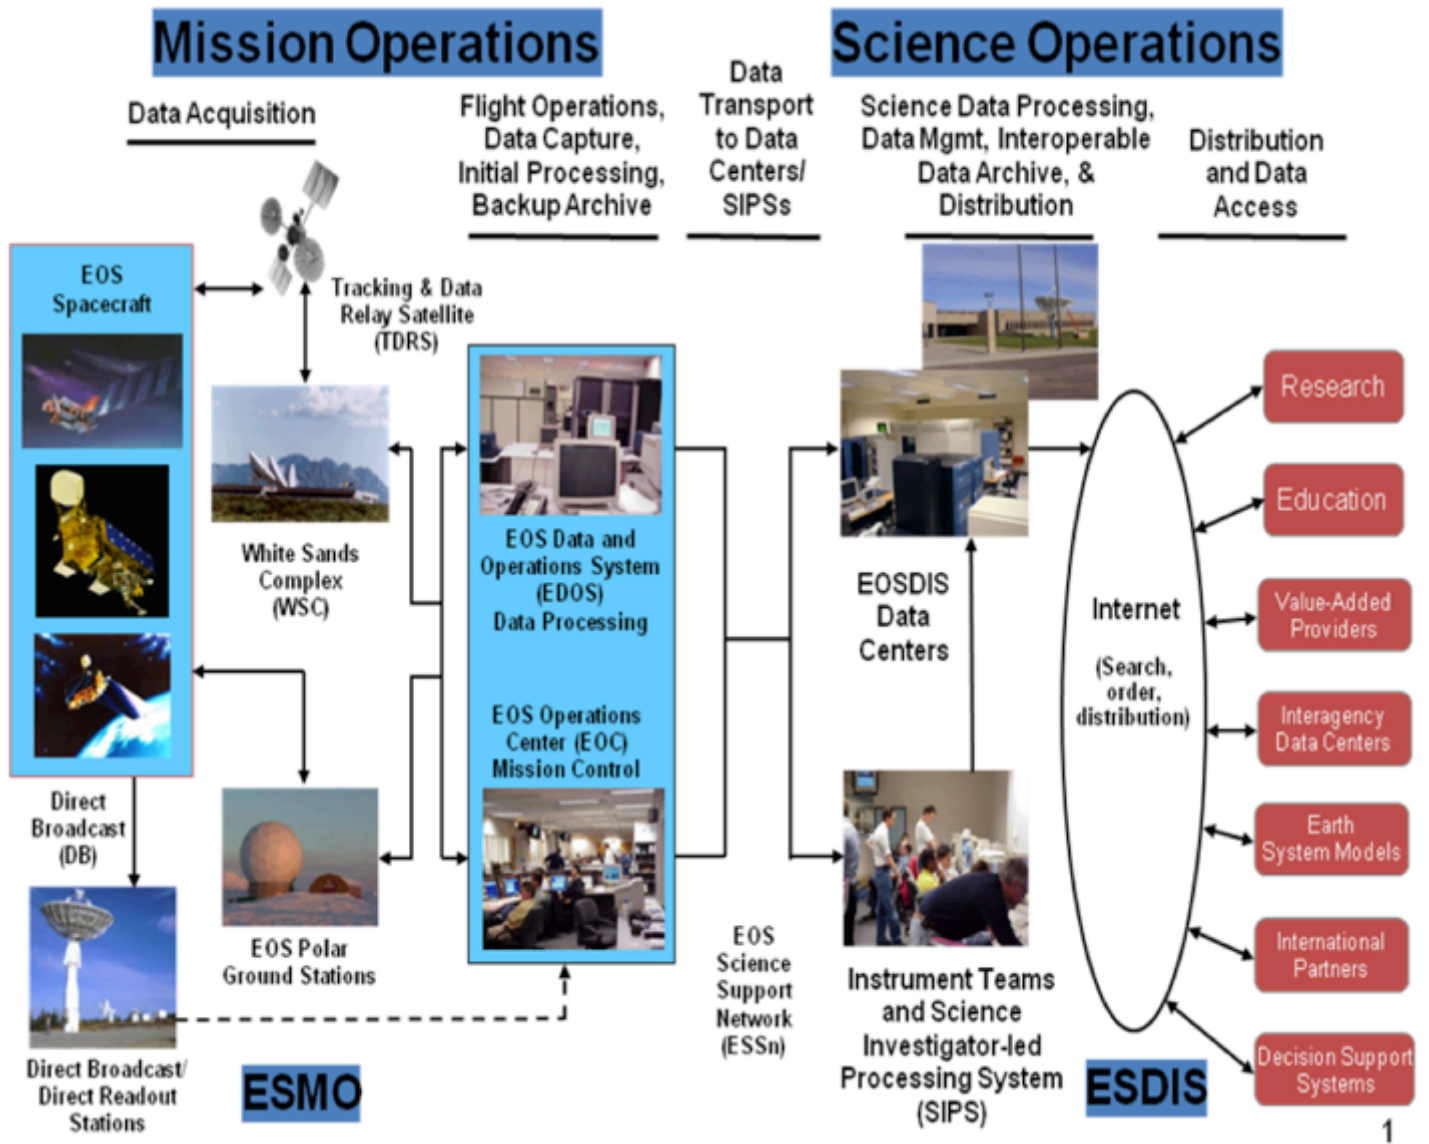


Figure 4: EOSDIS networks provide the backbone for moving data from acquisition to users. Mission Operations data flows deal with instrument/spacecraft operations and data capture. Science Operations transfers data to be processed, archived and eventually distributed to science teams and end users.

Thus, the ESMO network requirements are driven by mission operations and mission critical data transfers, while the ESDIS network requirements are driven by science data transfers and distributions. To meet these requirements, the ESMO and ESDIS projects employ a variety of network resources. The resources used for the mission operations and mission critical data transfers are grouped into a logical network called the “EOS Mission Support network” (EMSn). The resources used for the science data transfers and distributions are grouped into a logical network called the “EOS Science Support network” (ESSn). Note that these logical networks utilize some of the same physical resources.

These two logical networks consist of many segments of physical local area and wide area networks in the U.S. and abroad to satisfy the EOSDIS subsystems and users connectivity needs. These include NASA provided network services from NASA’s Communication Service Office (CSO) and network services provided by other entities.

The EMSn includes local-area and wide-area communication circuits and facilities between and among various EOS ground system subsystems to support EOS mission operations and mission-critical data transfers.

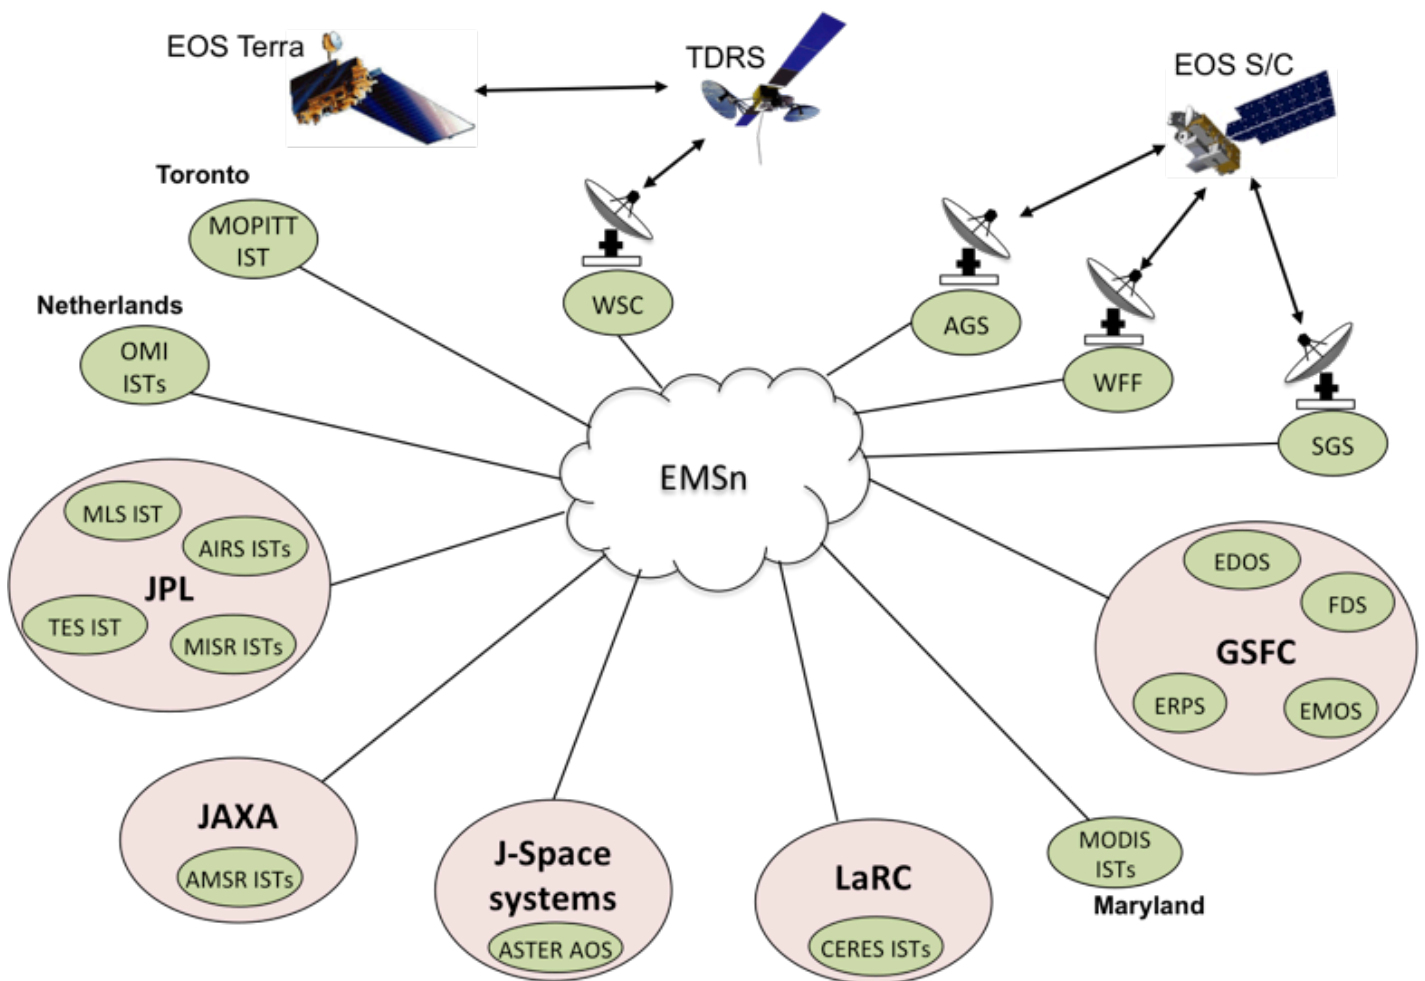


Figure 5: EMSn Logical Network. The EMSn provides mission operations and critical data transfer for EOSDIS

The ESSn includes shared globally connected logical or virtual science data communication networks comprised of several segments of shared IP-based internal and external physical networks such as the NASA Integrated Services Network (NISN) and the Internet2 IP backbone. The ESSn serves the needs of NASA’s diverse science and research community worldwide with high-performance links and gateways connecting to numerous research, educational, and public commercial networks via the Internet and national research networks in Europe, Asia, and other continents.

EOSDIS Networks Overview

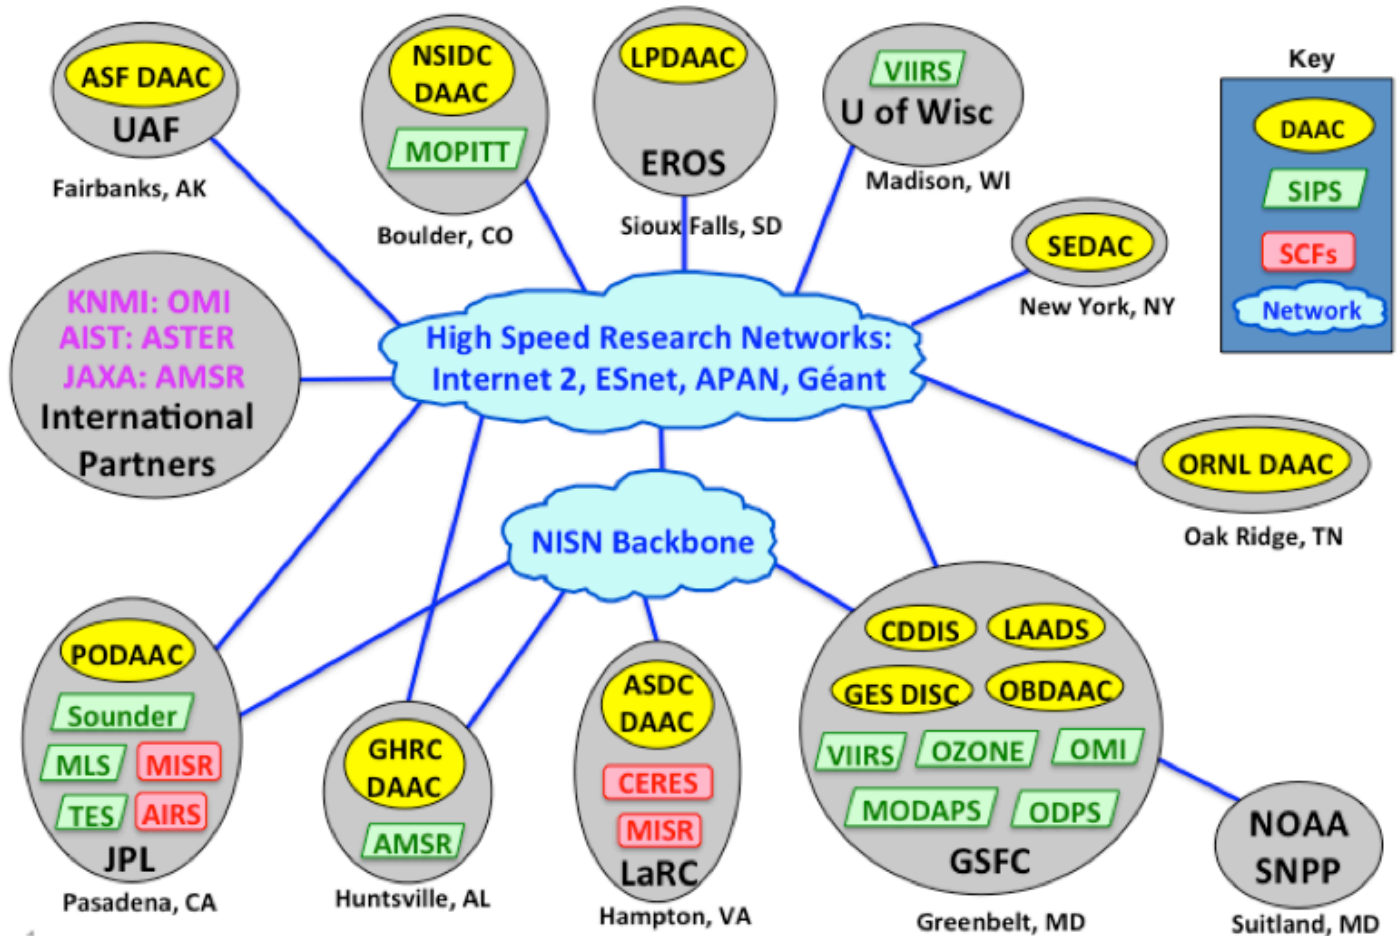


Figure 6: The ESSn logical network provides high speed data access to a diverse set of users. Acronym explanations can be found in the Appendix.

More information about EOSDIS network infrastructure is available in the Interface Control Document Between the Earth Observing System (EOS) Networks and the Earth Observing System Data and Information System (EOSDIS) Subsystems.

3.1.4. EOSDIS Common Services

A number of shared common services are offered by EOSDIS that in one way or another contribute to discovery, access, and use of EOSDIS data. Figure 7 provides a high level overview of the EOSDIS architecture and common services. Key services are described in more detail in the sections that follow.

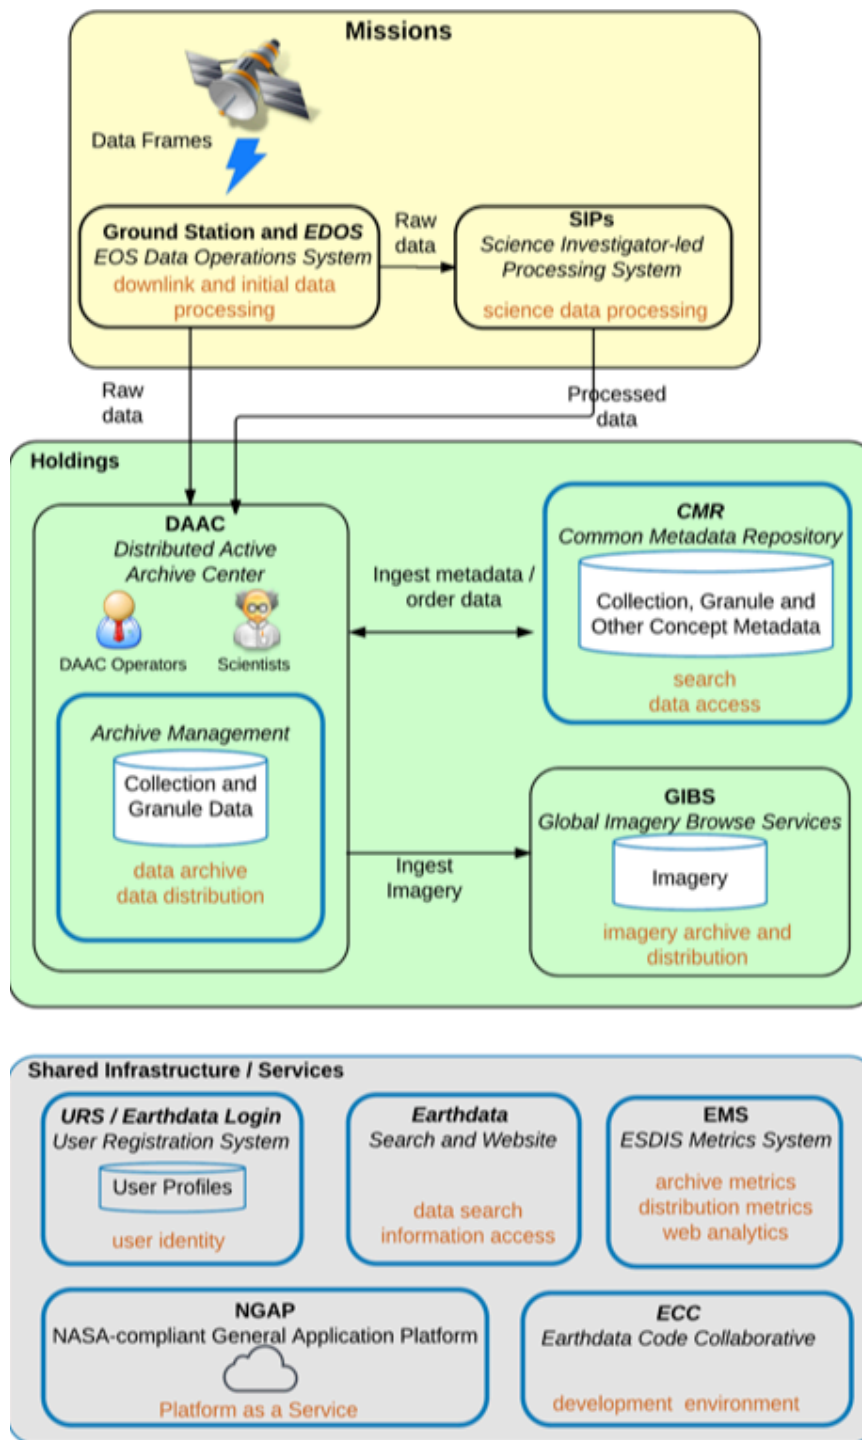


Figure 7: EOSDIS overview and common services

Common Metadata Repository (CMR)

The Common Metadata Repository (CMR) is a high performance, high-quality metadata system providing Ingest and Discovery capabilities for EOSDIS metadata. The CMR contains metadata for nearly 6400 EOSDIS collections and more than 270 million science granules as of March 2016. The CMR provides a single, authoritative metadata repository through the Unified Metadata Model (UMM). The UMM is an extensible metadata model that provides mappings between supported metadata standards and identifies EOSDIS compliant profiles.

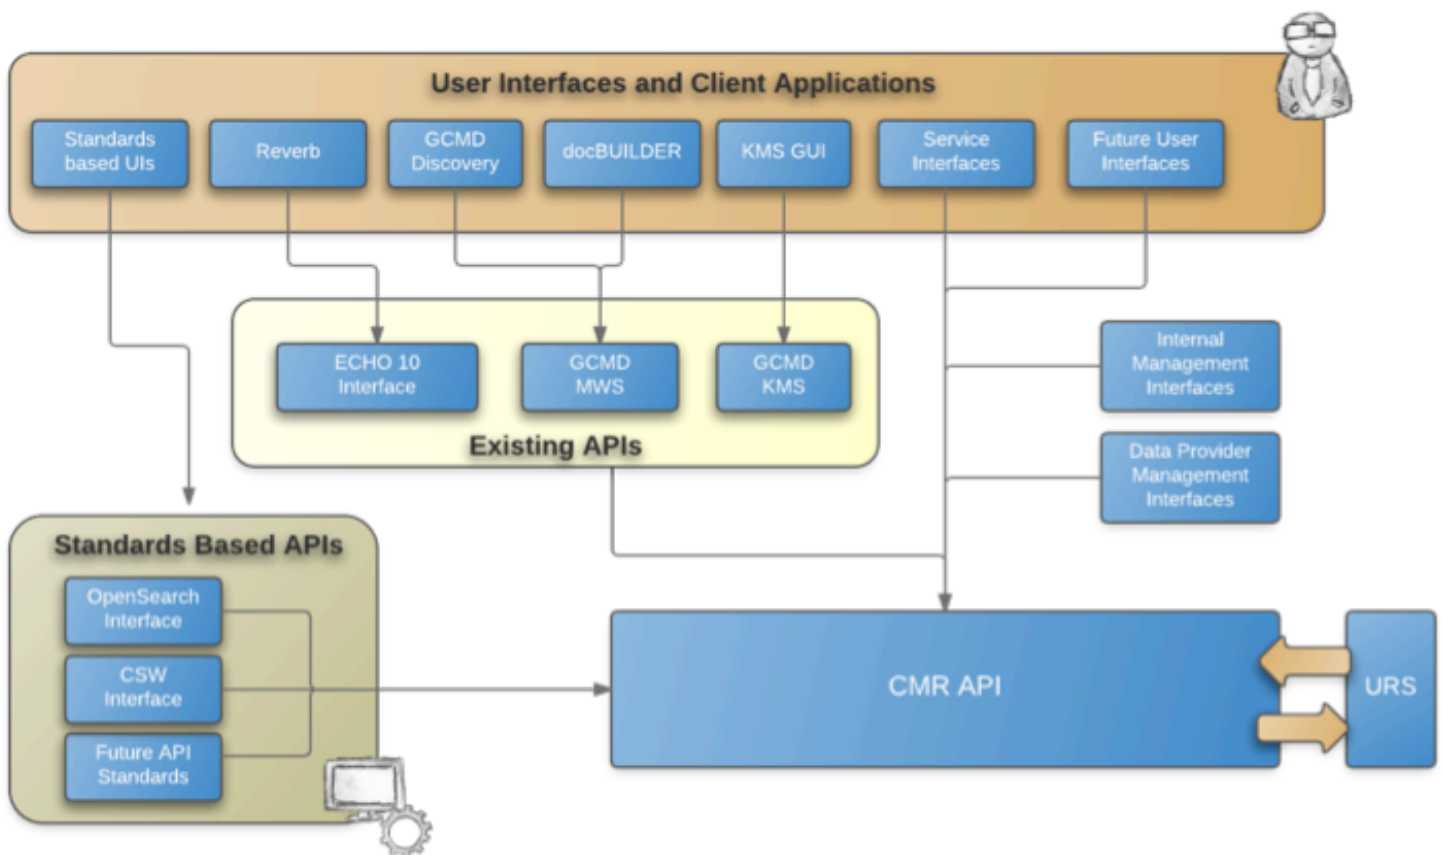


Figure 8: High level overview of key CMR interfaces

As data is made available at the EOSDIS DAACs, metadata is generated at the collection and, frequently, granule level in one of the supported EOSDIS metadata formats (see Section 4.1). This metadata is sent to the CMR to be ingested into the EOSDIS catalog for client discovery. Metadata is represented as a Metadata Concept in the CMR, with the CMR providing support for several core concepts including Collections, Granules, Services, Visual Metadata (e.g. browse) and Parameters.

As shown in Figure 8, metadata discovery is enabled through a series of APIs exposed by the CMR for cross-platform, cross application access. Through these APIs the CMR provides a number of high performance capabilities, including:

- Sub-second temporal, spatial, and keyword searching across hundreds of millions of science granule metadata records.
- Faceted responses to enable sophisticated client application data filtering and user navigation aids.
- Extensive metadata and response format support including ECHO10, GCMD DIF, ISO19115, Atom, and JSON. The CMR can automatically map disparate formats to a common response type based on client requests.
- Support for automated, rule based metadata generation, called Virtual Products, during ingest.

- Metadata Quality support capabilities including metadata adaptation during ingest, format specific validation, UMM rule validation, metadata revision tracking, meta-metadata tracking, and access control lists based permissions for metadata visibility and editing permissions.
- Support for relationships between metadata concepts, such as identifying which granules belong to which dataset.

The CMR is the hub for a number of core EOSDIS capabilities, providing a central point for navigating relationships between various EOSDIS data and application offerings.

Earthdata Web Site

The NASA Earthdata website is a collaborative environment integrating information from across EOSDIS. Earthdata is the entry point for EOSDIS data, articles, documentation, and collaboration and leverages NASA's CMR to provide comprehensive search capabilities. Earthdata offers new and experienced users an organized view of EOSDIS resources and latest events. EOSDIS DAACs are key contributors to Earthdata, providing the latest information on Atmosphere, Solar Radiance, Cryosphere, Human Dimensions, Land, and Ocean Science.

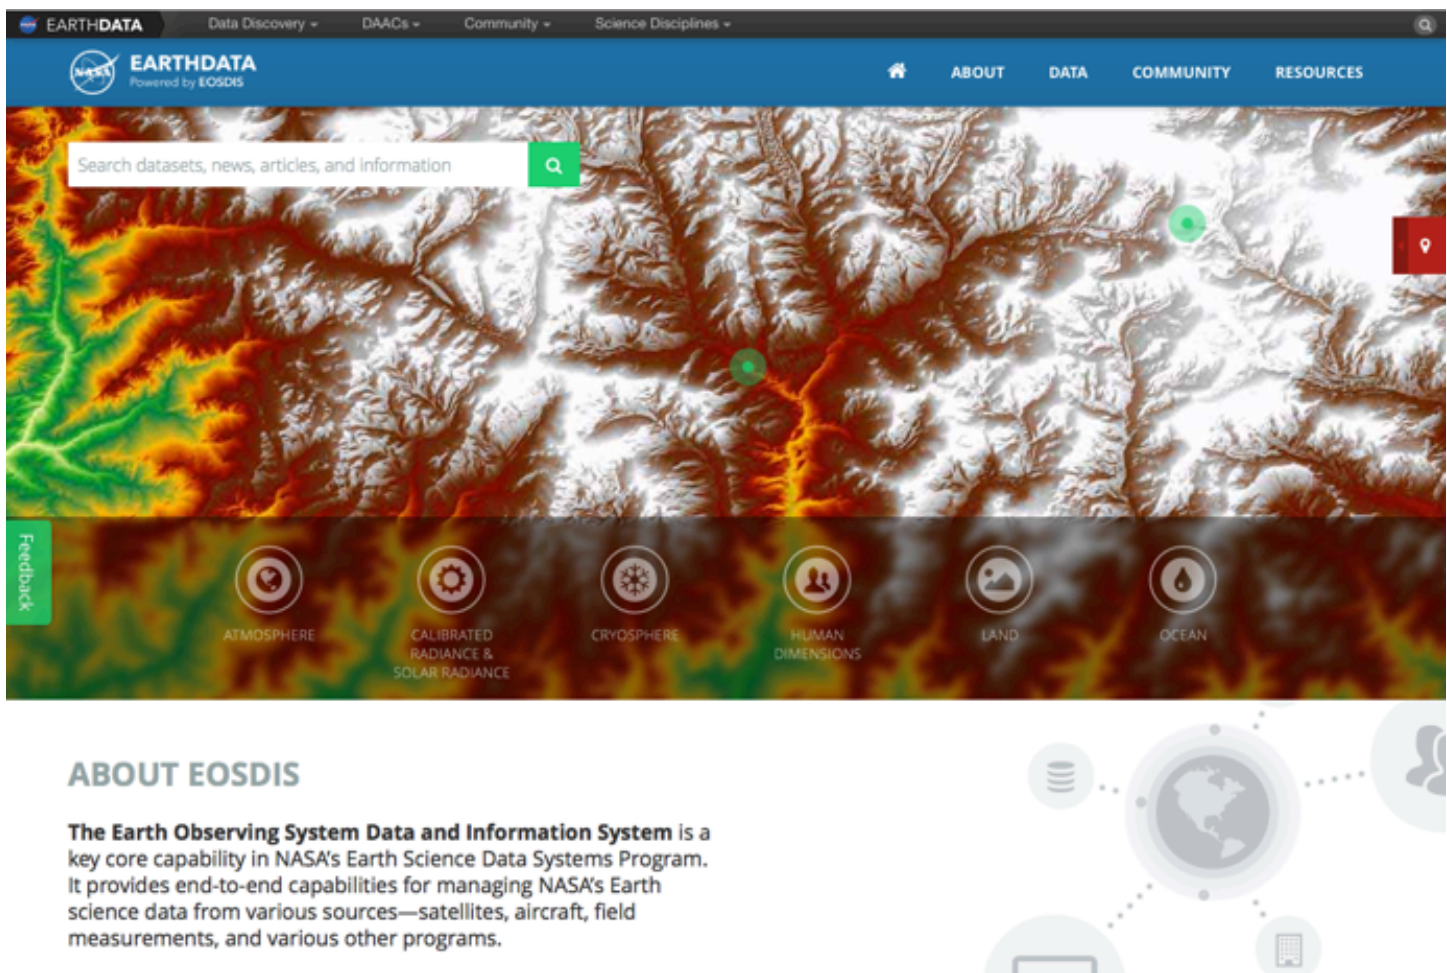


Figure 9: The Earthdata Web Site provides an entry point to EOSDIS information and holdings

Users can find a cross-DAAC, cross-domain view of EOSDIS resources along with the latest news and information, ultimately connecting them to the data, DAACs, and individuals best able to provide information. DAACs and researchers leverage Earthdata's collaboration capabilities to author and publish information to the EOSDIS community.

Earthdata leverages the CMR's information on relationships between data collections, individual granules, and unstructured content like web pages and documentation to provide users with the most relevant information available. Users frequently enter EOSDIS through Earthdata and then are handed off to more specific applications as they are presented with the information they're looking for.

Earthdata offers support to content authors in the form of a sophisticated Content Management System named Conduit that is capable of tracking drafts of content, associating it with data and other articles, and ultimately publishing the information to the global community.

The Earthdata website is available at <https://earthdata.nasa.gov>.

Earthdata Search

Earthdata Search is an open source EOSDIS web application that provides easy-to-use access to EOSDIS services for Earth Science data discovery, filtering, visualization, and access. It serves as a platform to feature planned EOSDIS services as they become available and integrates several core EOSDIS services into a unified user experience. Earthdata Search is able to take advantage of data services exposed by the EOSDIS community, in many cases simplifying data access and providing more targeted information to the user.

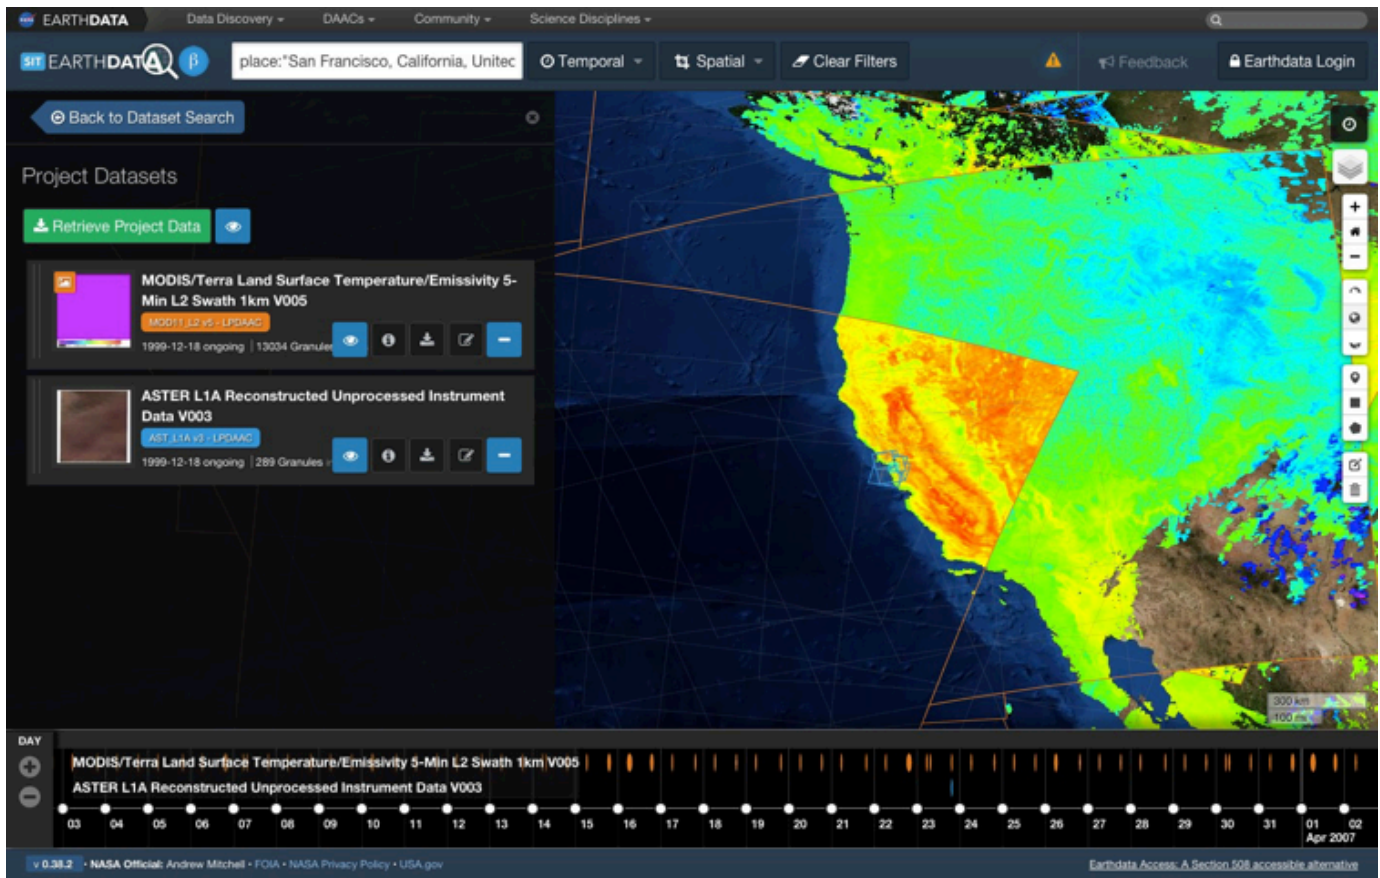


Figure 10: Earthdata Search provides an easy to use interface to EOSDIS data and imagery

Earthdata Search allows users to:

- **Search for Earth observation data:** Earthdata Search uses the CMR for sub-second search across the EOSDIS metadata catalog and simpler, faster access to faceted metadata. Earthdata Search will take advantage of new CMR capabilities as they're made available, such as parameter, visualization, and service metadata.
- **Preview Earth observation data:** Using Global Imagery Browse Services (GIBS), Earthdata Search enables high-performance, highly available data visualization. Earthdata Search provides data visualizations on a web-based map using the GIBS tile service. As GIBS adds support for granule-based visualizations, Earthdata Search will enable users to view even more precise map imagery and thumbnails for GIBS-enabled data collections.
- **Download and access Earth observation data:** In addition to direct download, Earthdata Search surfaces EOSDIS services for more customized access to EOSDIS data. Earthdata Search can interrogate and surface capabilities from OPeNDAP servers in addition to Order Form information from the CMR exposing advanced services like the HDF-EOS to GeoTiff Converter (HEG), the MODIS Reprojection Tool (MRT), and Geospatial Data Abstraction Library (GDAL) Translate.

Earthdata search is available at <https://search.earthdata.nasa.gov>.

Earthdata Code Collaborative

The Earthdata Code Collaborative (ECC) provides a full suite of development tools to EOSDIS projects that embody current best practices for software development.

The ECC offers tools and support for several key aspects of EOSDIS projects:

- **Requirements Gathering, Review, and Tracking:** The ECC offers a fully integrated requirements capture and tracking solution through Jama. Requirements, along with supporting documents, may be captured and tracked in the tool then opened for invitation only or public review. Once requirements are finalized they can be marked as approved and synchronized to the ECC's ticketing and work tracking software for implementation.
- **User Story / Work Tracking:** The ECC offers full Agile Development process support for implementation through Atlassian Jira. User Stories can be automatically synchronized with requirements to minimize process overhead while providing a consistent view of project progress and features. The ECC offers a standard workflow for projects with support for developers, internal testing, external testing, and deployments as well as custom workflows if needed.
- **Source Code Control:** The ECC offers full lifecycle source code control through Atlassian's Stash. This git-based source code control has support for branching and merging, code reviews, and linkage to Jira tickets.
- **Automated Builds and Continuous Integration:** The ECC offers sophisticated support for automated builds and continuous integration through Bamboo. The ECC is capable

of monitoring a project's source code repository for changes and then invoking custom builds and automated tests when detected. The ECC can monitor for branch creation as well, providing full support for feature-branch style development coupled with code reviews and subsequent merging. For approved, hosted EOSDIS projects, the ECC supports Continuous Deployment, automatically deploying the latest build of a project after a successful test execution and properly monitor versions by environment.

- **Team and User Collaboration:** Each registered ECC project gets a dedicated space in the Earthdata Wiki for team and user support. Team members can control access to their project wiki allowing for internally focused communication and collaboration as well as public facing support and communication. The ECC Wiki is aware of the other ECC components and can easily surface common agile graphs including burndowns, backlogs, build and test results, and latest deployment information.

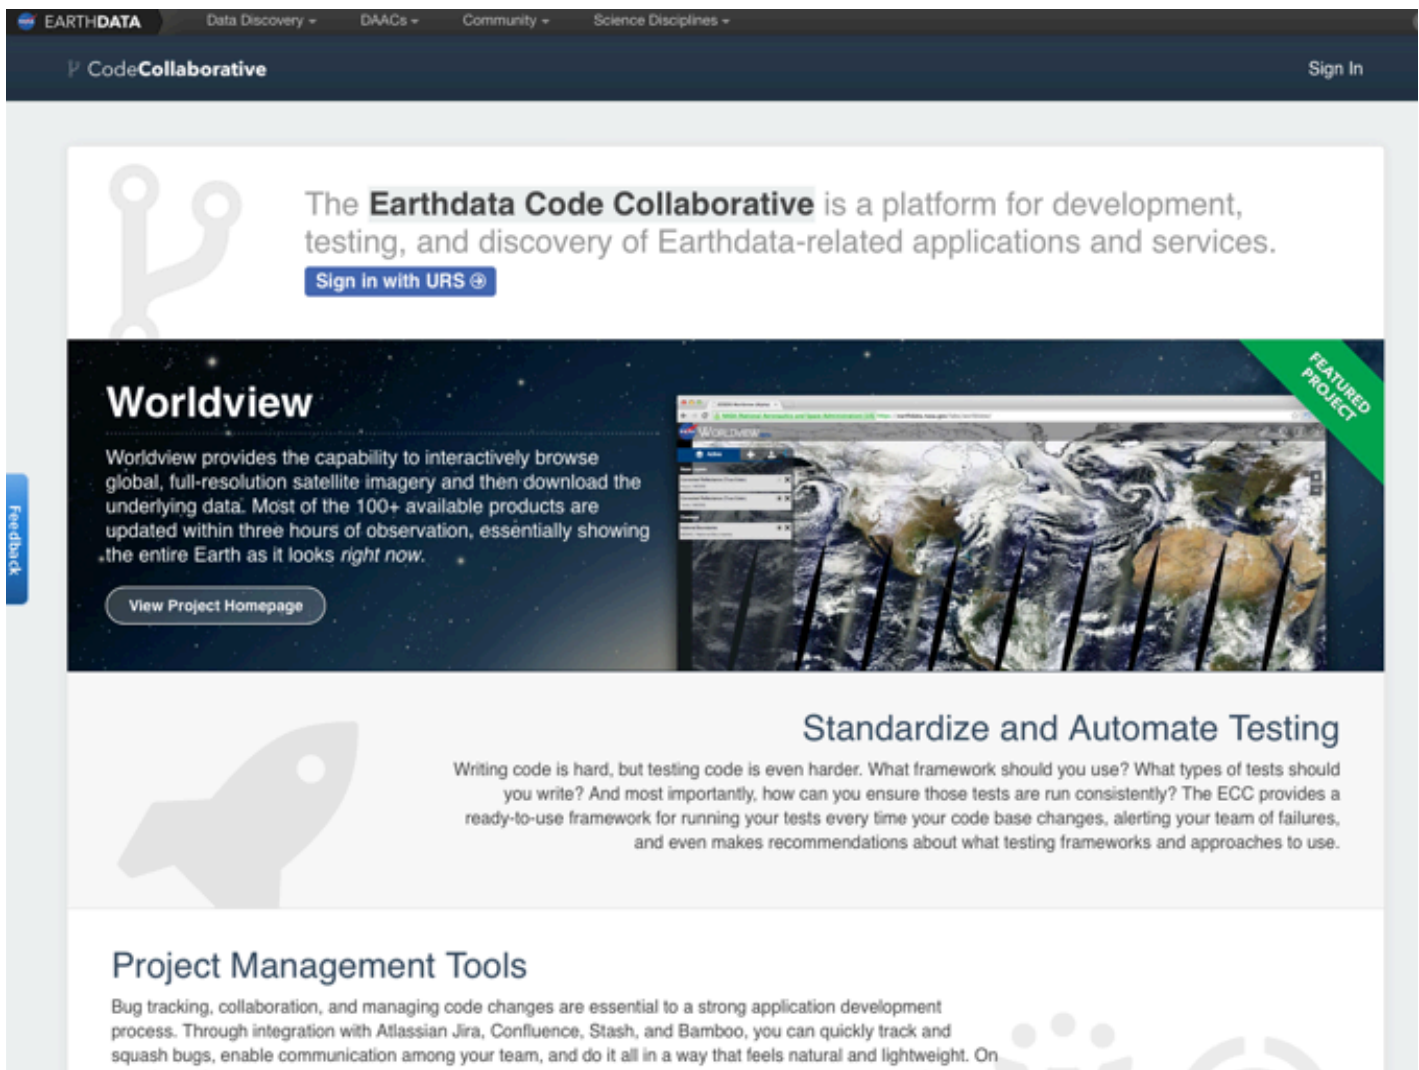


Figure 11: The ECC offers affiliated projects a highly capable development environment and process support

The ECC services are available to EOSDIS-affiliated developers and data centers, and require an Earthdata Login account. For more information or to request access please visit the ECC here: <https://ecc.earthdata.nasa.gov>.

Global Imagery Browse Services (GIBS)

The Global Imagery Browse Services (GIBS) system is a core EOSDIS component which provides a scalable, responsive, highly available, and community standards based set of imagery services. These services are designed with the goal of advancing user interactions with EOSDIS' inter-disciplinary data through enhanced visual representation and discovery.

These advancements are realized in the following ways:

- **Improved Approachability & Extended Reach:** Imagery greatly improves the usability of NASA Earth science data to new communities and improves cross-disciplinary data discovery through full-resolution, “no boundaries” (or “granule-free”) interaction patterns.
- **Cohesive Approach to Imagery:** As a core EOSDIS component, GIBS integrates with other core EOSDIS systems, components, and processes to provide a primary, authoritative source for EOSDIS imagery.
- **Improved Cross-Discipline Research:** GIBS leverages science expertise and interoperable standards to provide science-based products that enhance cross-discipline discovery and analysis.

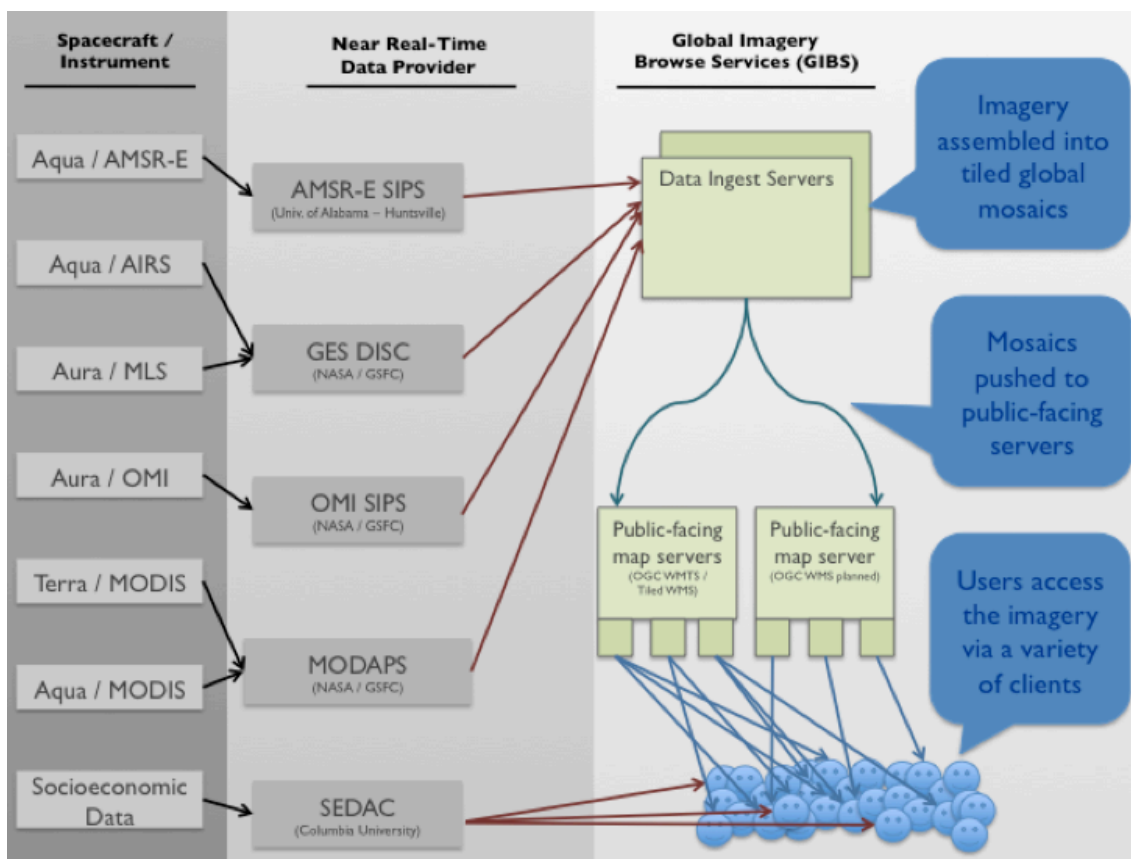


Figure 12: GIBS imagery comes from a variety of datasets and is made available through standard APIs

The GIBS imagery archive includes over 100 imagery products representing visualized science parameters from EOS. Each imagery product is generated at the native resolution of the source data to provide “full resolution” visualizations of a science parameter. GIBS works closely with the science teams to identify the appropriate data range and color mappings, where appropriate, to provide the best quality imagery to the earth science community.

As shown in Figure 12, many GIBS imagery products are generated by the Land, Atmosphere Near real-time Capability for EOS (LANCE) near real-time processing system resulting in imagery available in GIBS within 3.5 hours of observation (see Section 3.3.9). These products and others may extend from the present back to the beginning of the satellite mission. GIBS makes available supporting imagery layers such as data/no-data, water masks, orbit tracks, and graticules to improve imagery usage.

More information on GIBS is available at <https://earthdata.nasa.gov/about/science-system-description/eosdis-components/global-imagery-browse-services-gibs>.

Worldview

Worldview is an open source EOSDIS web application allowing users to interactively browse global satellite imagery within hours of it being acquired then download the underlying data. Most of the 100+ available products are updated within three hours of observation, essentially showing the entire Earth as it looks “right now”. This supports time-critical application areas such as wildfire management, air quality measurements, and flood monitoring. Arctic and Antarctic views of several products are also available for a “full globe” perspective. Browsing on tablet and smartphone devices is generally supported for mobile access to the imagery. Worldview uses the Global Imagery Browse Services (GIBS) to rapidly retrieve its imagery for an interactive browsing experience.

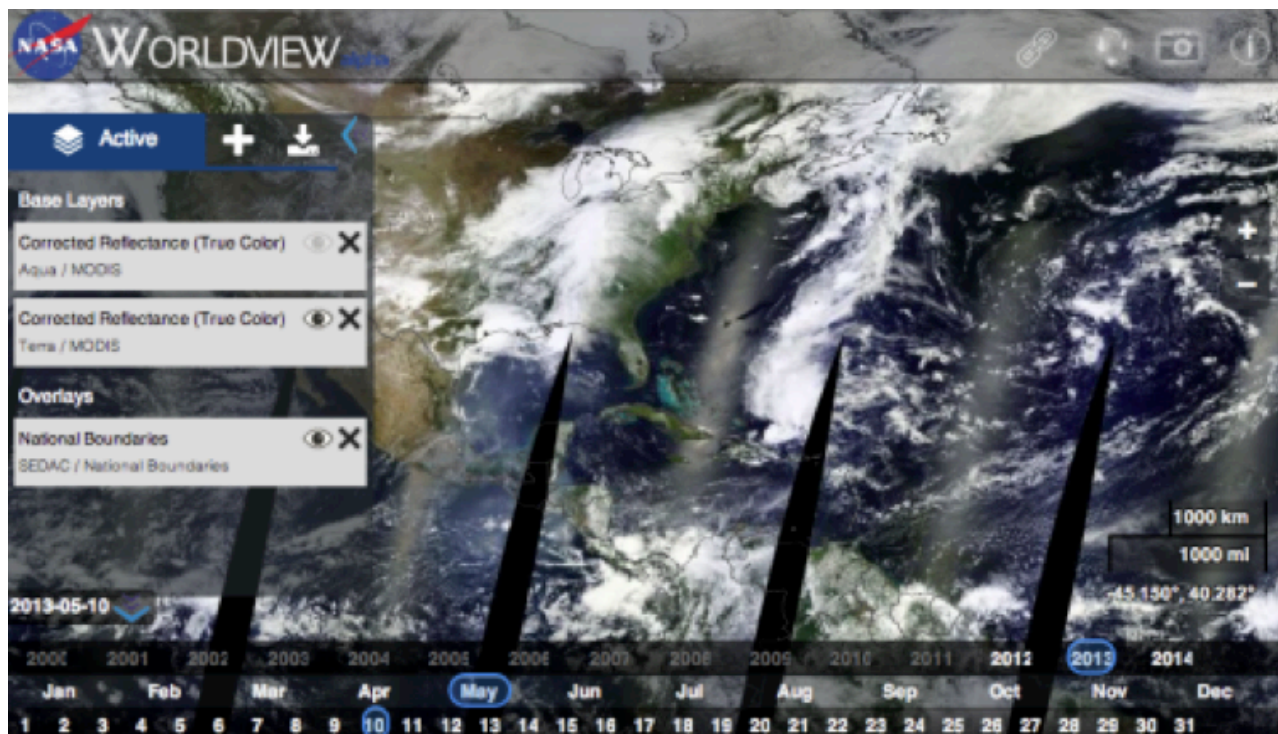


Figure 13: Worldview leverages near realtime imagery and GIBS to provide a “right now” view of the Earth

3.1.5. Metrics and Monitoring

The ESDIS Metrics System (EMS) supports ESDIS project management by collecting and organizing a variety of metrics from the DAACs and other Data Providers (such as the GCMD and ESDIS-sponsored websites). The EMS collects and supports the analysis of data on the usage of products and services delivered via the Internet or stored in databases. Specifically, the EMS allows the project and other users access to metrics from the Data Providers including such information as number of users, type and amount of data archived, type and amount of data distributed, and other related information. The EMS is a combination of a commercial web analytics package and custom code to track ingest, archival, distribution and user profile information collected from Data Providers.

The ESDIS Project monitors effectiveness in meeting user requests for data as measured by the American Customer Satisfaction Index (ACSI) for EOSDIS, which is evaluated annually via a survey of EOSDIS DAAC users by an organization independent of the ESDIS Project.

3.2. Standards and Conventions Organizations

EOSDIS is comprised of a number of systems from different organizations and a diverse set of users and contributors. Successful interoperability of these systems and supporting efficient use of EOSDIS offerings requires participation and support for broadly used standards and conventions. This section outlines how EOSDIS engages with the broader community and standards promoted within EOSDIS.

3.2.1. ESDIS Standards Office (ESO)

The ESDIS Standards Office (ESO) assists the ESDIS Project in formulating standards policy for NASA Earth Science Data Systems (ESDS), coordinates standards activities within ESDIS, and provides technical expertise and assistance to standards related tasks within the NASA ESDSWG. The link to the ESO page is found here: <https://earthdata.nasa.gov/about/esdis-project/esdis-standards-office-eso> . ESO maintains a list of standards approved for use in NASA Earth Science Data Systems. ESO also convenes ad hoc working groups to review candidate standards and technical notes that may be brought forward by ESDSWG, ESDIS, or others in the NASA Earth science community. The ESO recruits members for the review groups from ESDSWG membership and the broader Earth science informatics community.

3.2.2. ISO19115 Metadata Standards

EOSDIS has embraced ISO 19115 for high quality metadata describing EOSDIS holdings. ISO 19115 defines mandatory and conditional metadata structure and definition, minimum set of fields for given applications, and methods for extending metadata if needed, for specialized needs. Beginning in 2010, a cross-EOSDIS team was assembled for the Metadata Evolution for NASA Data Systems. The team recommended adoption of ISO 19115 as a model for metadata exchange and produced detailed translations of the ISO 19115 metadata model to EOSDIS's existing internal database models. The mappings were used to develop the Unified Metadata Model, which melded the ISO 19115 model with legacy models at both the data collection and data file level. The CMR was developed to support ISO 19115 as both an input and output model, with ISO 19139 providing the XML encoding. DAACs create ISO 19115/19139 compliant metadata for ingest into the CMR, and, in situations where legacy models and formats are still used, the CMR can automatically translate them into ISO19115/19139 for client application use.

3.2.3. Committee on Earth Observation Satellites (CEOS)

CEOS was established in 1984 to address the multidisciplinary nature of space-based Earth observations and the value of coordinating international Earth observation efforts to benefit society. CEOS coordinates and harmonizes Earth observations to make it easier for the user community to access and use data, spawning efforts and recommendations for interoperability, common data formats, the inter-calibration of instruments, and common validation and inter-comparison of products.

The ESDIS project leads NASA's participation in the CEOS Working Group on Information Systems and Services (WGISS), which creates and demonstrates prototypes supporting CEOS requirements. WGISS also addresses the creation of information systems, the management of Earth Observation data, and the delivery of interoperable services. WGISS's scope spans the information life cycle from requirements and metadata definition for satellite data through the incorporation of derived information into end-user applications.

3.2.4. Group on Earth Observations (GEO)

CEOS has played an influential role in the establishment and ongoing development of the Group on Earth Observations (GEO) and the Global Earth Observation System of Systems (GEOSS). GEO is engaged in the development of GEOSS to link Earth observation resources worldwide across multiple Societal Benefit Areas.

NASA, including the ESDIS project, contributes to GEO through two distinct but complementary channels: through the accomplishments of contributors to the CEOS as the space component of GEO, and through USGEO, which coordinates GEO participation across federal agencies and represents the U.S. as a member of GEO. Some NASA contributions are made in CEOS; some are made in USGEO; some complement one another in both forums; and still others exist independently to each other's benefit.

3.3. Key System Functions / Business Use Cases

In order to coordinate the development of disparate systems it is critical that the common goal is understood along with the role each system plays in reaching that goal. EOSDIS identifies several key system functions, or business use cases, that are essential to meeting the needs of EOSDIS users. This section discusses the key system functions, often realized across the enterprise by distinct participating applications and services, which define the core of EOSDIS.

3.3.1. Ingest

Data ingest is a fundamental use case between a data provider and a designed DAAC tasked with archiving and supporting the data for end users. It is important to note that while the 12 EOSDIS DAACs support the core system functions, DAAC specific configuration or implementations exist. This section describes the general flow for data ingest between a data provider and a DAAC, however there are variations in implementations.

The Ingest function for DAAC systems is a process that regularly checks for the arrival of a PDR in a file location that is configured for each data provider. These locations may either be located on a remote host or on a local host as agreed upon between the DAAC and the data provider. The interval between checks for new data is set by the DAAC based upon the delivery frequency for data by that provider.

When a new PDR is detected by the DAAC system, it transfers the record to an internal location for use by the Ingest process. PDRs are generally processed on a first come first served basis but selected providers may be

given a higher priority due to processing time requirements or other considerations.

Each PDR includes an optional data source identifier that is for the reference of the data provider, a total count and an aggregate length of all the files included in the PDR. This is followed by a series of file groups. Each file group consists of a data type and a version identifier data contained in the PDR, a metadata file name, location, length in bytes, and 1-to-n data files with their location and length in bytes that go with the metadata file. The process retrieves the files specified in the PDR verifying that the length of each file and the aggregate length agree with that in the PDR.

Next, the ingest process parses the metadata provided with each file to perform basic sanity checks for the data files described. The required fields are configured when the data type is initially configured in the system. These values are then used to populate a record in the Inventory database maintained by the system. If any of these tests fail, the ingest will be marked as failed.

Once this process is complete the process then generates a metadata file in XML format for storage in the system archive. The ingest process then sends the files as received to the archive process for permanent archive.

As the final step a success or failure notification is sent to the data provider and insert events are propagated to other system processes such as a subscription service and the metadata export task to send the metadata to the Common Metadata Repository. The ingest process is then marked as complete.

3.3.2. Science Data Processing

The science data processing function transforms Level 0 (raw) data into more usable products. Most data sets are processed to Level 1B, the term for calibrated and geolocated radiance. Figure 13 shows a Level 1B scene from MODIS on the Aqua platform for 28 March 2008 off the west coast of Africa.




Figure 14: Level 1B Calibrated Radiance from MODIS on Aqua, 28 March 2008. The image approximates true color, using MODIS Band 1 for Red, band 4 for Green, and band 3 for Blue.

Level 1B is then typically processed to Level 2, which are geophysical parameters in the satellite's orbital coordinates. Figure 15 shows retrievals for Aerosol Optical Depth from the above scene using the Deep Blue algorithm (for highly-reflective land areas).




Figure 15: Aerosol Optical Depth retrieval from a MODIS tile.

However, the orbital swath based coordinates in Level 2 data can be difficult to deal with, and individual files generally cover only relatively small areas. As a result, many of the science teams also create Level 3 products that aggregate the data files in space (i.e. on a regular grid) and time (e.g. averaged over a day or month) as shown in Figure 16. These files are generally easier for many users to work with, but at the cost of spatial resolution.

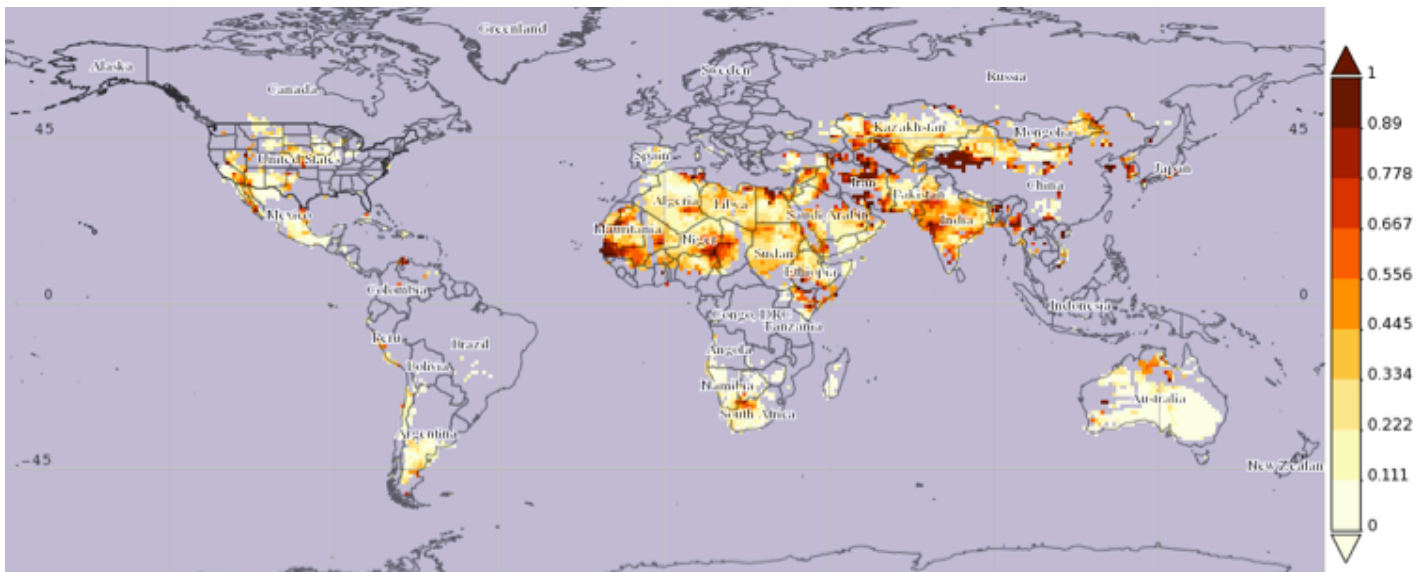


Figure 16: Level 3 gridded Aerosol Optical Depth (Deep Blue Algorithm) for MODIS on Aqua satellite, 28 March 2008.

Science data processing can be done on processing systems operated by the DAAC, at a SIPS managed by the instrument team, or a science processing system managed by the flight project or external scientist. In many cases, the data are transferred to the archive using the EOSDIS-standard PDR mechanism.

3.3.3. Archive

The archive process stores ingested data into a permanent archive of all data ingested by the system. At most of the DAACs, the data are stored on disk for immediate (synchronous) access by end-users. However, at a few of the DAACs, the volumes are (currently) too great to afford disk storage. In these cases, a robotic tape library is used for storage, supplemented by online cache space. For online storage, whether for basic archive or cache, are typically organized into directories by mission and dataset.

DAACs also typically store one or two copies of the data (depending on criticality), one of which is in another building, or ideally in another county or state for disaster recovery. (An EOSDIS-wide data center risk analysis documents the various backup policies at the DAACs.) The archive systems include utilities to verify data integrity of the data stored on the disk subsystem based on stored checksums. If a data set is found to have issues, it is flagged for operator action to correct the issue by restoring the data set from the backup copy.

3.3.4. Discovery

Discovery of EOSDIS data and related information is a critical step of leveraging EOSDIS data and is supported at several levels by EOSDIS systems. Over the last decade, EOSDIS has pursued a two-pronged approach, both developing a cross-EOSDIS search tool and enabling other organizations to build their own discovery tools with different focuses and targeted user bases by providing discovery at a service level.

The linchpin to EOSDIS discovery services is high quality metadata describing key EOSDIS assets including data holdings, services, visualizations and documentation. Metadata are generated or extracted during the data ingest process at the DAACs and submitted to the Common Metadata Repository (CMR) as collection and granule-level records for indexing. For non-data assets, metadata is either generated by hand with metadata authoring tools or automatically through tools Content Management Systems like Conduit (see Section 0). The

CMR indexes EOSDIS metadata and provides the foundation for discovery services, offering high performance search across the hundreds of millions of science granules held by EOSDIS.

The CMR exposes open access APIs to enable diverse search and discovery application development across EOSDIS information (see Section 5.1.1). Developers can leverage CMR APIs to produce broad general-purpose discovery clients like NASA's Earthdata Search or more specific, tailored applications such as the ASTER Volcano Archive.

In addition to basic metadata, data providers can choose to create high-resolution imagery in the form of Browse images or entire daily composite layers which are submitted to GIBS. Information about this high-resolution imagery is captured in metadata defined by the UMM-V and made available in the CMR for search.

By way of general-purpose search, discovery and access, EOSDIS maintains several core web applications. The Earthdata search client (<http://search.earthdata.nasa.gov/>) provides a familiar search interface exposing a map for visual spatial selection, calendar entries for time constraints, and a text field for free text searching. The Earthdata search client follows a traditional shopping cart metaphor for users to tag data of interest and is capable of leveraging ordering and web service offerings to further refine data selections by the user.

For a more visual presentation of data the Worldview application provides access to Near Real Time imagery and leverages the CMR to provide users access to the underlying data. Similarly, the Earthdata Search application leverages CMR to provide users access to the full EOSDIS holdings, rendering visual metadata from GIBS along with temporal information allowing for easy identification of temporally coincident information.

In addition to general-purpose client applications, EOSDIS common services provide backend support to a number of specialized and unique client applications, e.g.:

- **CWIC and CWICSmart:** The CEOS WGISS Integrated Catalog (CWIC) provides a brokerage service for granule level OpenSearch queries to a variety of disparate providers that may not have an OpenSearch API. CWIC also provides the NASA OpenSearch reference client CWICSmart. More information on CWIC can be found at: <http://ceos.org/ourwork/workinggroups/wgiss/current-activities/cwic/>.
- **LPDAAC2Disk:** The LPDAAC2Disk download manager (https://lpdaac.usgs.gov/data_access/data_pool) is both a web application and command line executable that searches for data products in the CMR. The LPDAAC2Disk tool provides a simplified search and HTTP download interface to the LP DAAC's online holdings, emulating an FTP directory hierarchy.
- **Giovanni:** Giovanni (Geospatial Interactive Online Visualization ANd aNalysis Infrastructure) is a Web-based application developed by the Goddard Earth Sciences Data and Information Services Center (GES DISC) that provides a simple and intuitive way to visualize, analyze, and access vast amounts of Earth science remote sensing data without having to download the data. Giovanni is available at: <http://giovanni.gsfc.nasa.gov/>.
- **ASTER Volcano Archive:** The ASTER Volcano Archive (AVA) comprises over 40,000 full-resolution ASTER Very Near Infrared (VNIR) daytime images of volcanoes, about a third of the ASTER volcano holdings, for 1542 active volcanoes listed in the Smithsonian Global Volcano Catalog. The AVA project uses volcano location information to search

CMR for daytime ASTER VNIR images. More information about AVA is available here: [https://science.jpl.nasa.gov/projects/ASTER%20Volcano%20Archive%20\(AVA\)/](https://science.jpl.nasa.gov/projects/ASTER%20Volcano%20Archive%20(AVA)/).

- **Simple Subset Wizard:** The Simple Subset Wizard (SSW) provides a simple, unified user interface for submitting subset requests for data at eight of the EOSDIS DAACs. More information about the Simple Subset Wizard is available here: <http://disc.sci.gsfc.nasa.gov/gesNews/new-simple-subset-wizard-makes-subsetting-simple>

3.3.5. Data Access and Ordering

EOSDIS data located at the DAACs is accessible for search, discovery and ordering through the CMR. The DAACs must export collection and, if desired, file level metadata for the data sets to the CMR in order to make them accessible to users. The CMR provides an API for search and order. The two primary EOSDIS supported clients, Reverb and Earthdata Search, provide means of selecting data to order through graphical user interfaces and allow the user to constrain results by many different parameters (temporal, spatial, instrument type, etc.). The clients also provide access to data services such as reformatting, subsetting etc. as described in section 3.4.6

In addition to the EOSDIS clients, DAACs may expose a custom interface for some data access, such as the EOSDIS Core System (ECS) Data Access utility. This utility provides direct access to data at ECS DAACs. A number of data services are also available through the Data Access interfaces.

Finally, CMR provides public support for the CMR search and order API. Users can use the API to access and order data directly using their own software written to interface with the API.

3.3.6. User Identification

EOSDIS includes a facility named Earthdata Login to help track detailed data usage at the individual user level, enabling both accurate metrics on data usage and the ability to send critical data update messages. Earthdata Login includes a User Registration site and protocols that allow both DAAC-developed and ESDIS-developed services to look up the user id for anyone accessing data. In general, the intent is to require registration for data access, but to leave it optional for search and browse functions. Users who register for an Earthdata Login account may potentially use the same account for all EOSDIS web services.

Applications within EOSDIS may use one of the software identification modules provided by the Earthdata Login team (i.e. Apache, Ruby, perl-cgi, etc.) for connecting to the Earthdata Login service. The host of the web service simply has to create an account with Earthdata Login and register the application they wish to use with the Earthdata Login service. Once the Earthdata team approves the application it is a matter of configuring the web service with the appropriate setting using the desired Earthdata module.

Configurations can be extended beyond simple user validation services, for example the Apache module can be configured to allow access to any Earthdata user to certain directories, but only certain specific users to access other data directories. Thus the application provider has a control over data accessibility using the Earthdata Login service.

3.3.7. User Services

A founding principle of EOSDIS is to provide user services for NASA Earth Science data for all user communities and science disciplines. Each DAAC provides support to the users of the data they distribute, including services to:

- Assist users in identifying and locating data sets
- Coordinate with other DAACs to assist cross-discipline users
- Track data order accuracy and completion of any data orders per system in use at the DAAC or any of the common distribution systems
- Compile user services statistics and metrics on data distribution usage
- Provide demonstrations of data center-specific tools
- Provide science expertise for the data and tools including any technical guide documents, dataset landing pages and web information pages
- Provide information on processing history and product revisions for data
- Provide up-to-date data validation information
- Collect and compile data quality comments from users
- Identify best practices, leverage strengths, and coordinate efforts
- Communicate user “pain points” to developers within EOSDIS
- Participate in system-wide social media activities including requests for conference support, participation in workshops, webinars, social media articles and website image support

3.3.8. Metrics

EOSDIS is measured continually to determine how well it is meeting users’ needs and how the system is performing. As mentioned in an earlier section, an annual survey is conducted of all current EOSDIS customers for whom an email address is available. The ESDIS Project contracts for a survey of the EOSDIS and DAAC Performance using the American Customer Satisfaction Index (ACSI). This evaluation has been performed annually by an organization independent of the ESDIS Project for over 10 years. The ACSI survey asks a series of questions to determine how satisfied customers are with EOSDIS and what problems they have encountered or suggestions they have for improvement. These are analyzed to produce a standardized score that is comparable both with previous years and with other organizations within the federal government. Detailed analysis of the survey results, including the freetext comments, sheds light on which areas of EOSDIS need the most improvement and suggests directions that improvement might take.

ESDIS also uses the EMS to maintain metrics for archive and distribution. The individual DAACs send archive logs and internet traffic logs to the EMS on a daily basis, which are then aggregated into a database. A user interface is provided for the DAACs to query the database for a variety of ingest, archive and distribution statistics,

both current and historical. Further explanation of EOSDIS metrics as well as annual reports and ACSI Survey results can be found on the Earthdata website at <https://earthdata.nasa.gov/about/system-performance>.

3.3.9. Near Real Time Support

In addition to science quality data products, NASA makes near real-time (NRT) data available for select datasets. The Land, Atmosphere Near real-time Capability for EOS (LANCE) supports application users interested in monitoring a wide variety of natural and man-made phenomena. NRT data and imagery from the AIRS, AMSR2, MLS, MODIS and OMI instruments are available much more quickly than routine processing allows. Most data products are available within 3 hours from observation. NRT imagery is generally available 3-5 hours after observation.

It is important to recognize the NRT data is not considered science quality. Science quality, or higher-level “standard” data products are an internally consistent, well-calibrated record of the Earth’s geophysical properties to support science. They are made available within 8-40 hours of satellite observation.

LANCE shares a number of services with the rest of EOSDIS, including user registration, web services, user services, data metrics, browse products, and tools for generating subsetted products and products using a variety of data formats. LANCE users span a variety of communities interested in a range of applications. Users come from both the civilian and military sectors, from government and non-government agencies and from universities and other research intuitions. LANCE provides NRT data supporting fire, floods, agriculture, air quality, and sea ice.

LANCE is reviewed and supported by a User Working Group (UWG) whose members reflect the various user communities served by LANCE as well as representatives of the Science Teams for the five LANCE instruments. The UWG meets at least once per year to review the status of LANCE operations and development activities and to provide guidance concerning future upgrades suggested by the user communities and the LANCE elements.

For more information on LANCE and EOSDIS NRT data please see: <https://earthdata.nasa.gov/earth-observation-data/near-real-time>.

3.4. Related Initiatives

This section discusses key relationships and initiatives undertaken within EOSDIS that encourage broad participation and contribution to the science community. These relationships and efforts inform EOSDIS use cases and services, ultimately helping to shape EOSDIS offerings.

3.4.1. Federation of Earth Science Information Partners (ESIP)

The Federation of Earth Science Information Partners (ESIP) is a broad-based, distributed community of data and information technology practitioners who come together to collaborate on coordinated interoperability efforts across Earth science communities. The ESIP Federation was founded by NASA in 1998 in response to a request to evolve new, distributed information structures for EOSDIS that would allow easier input by the Earth science community, which would include federal data centers, government research laboratories, research universities, and various nonprofit and commercial enterprises from more than 170 member organizations dedicated to developing methods to make Earth science data easy to preserve, locate, access, and use. The ESIP Federation is part of the nonprofit Foundation for Earth Science, which supports organizations involved in Earth science data by providing a collaborative infrastructure, development services, and event planning.

3.4.2. Earth Science Data Systems Working Group (ESDSWG)

Among EOSDIS, the DAACs, and investigators on NASA informatics-related projects, NASA has access to a wide variety of technical expertise. In order to leverage that expertise, NASA has developed an annual cycle for chartering working groups to facilitate technology infusion and develop community conventions for data systems. Each year, EOSDIS hosts a meeting at which the working groups present the previous year's output, which may include Use Cases, Technology Recommendations, Standards or Community Conventions. Standards and Conventions are further forwarded to the EOSDIS Standards Office (See Section 3.2.1) for review and disposition, while other recommendations are used by EOSDIS to plan for technology infusion and long-term strategy. At the meeting, new groups coalesce around current (and future) issues with potential impact on EOSDIS, and existing groups plan to take up a new set of topics within their theme.

3.4.3. Big Earth Data Initiative (BEDI)

The Big Earth Data Initiative (BEDI), initiated and sponsored by the White House Office of Science and Technology Policy, is an interagency effort to improve interoperability among the various federal Earth observation data systems and to increase the discoverability, accessibility, and usability of federally held Earth observation data. While the planned improvements will benefit the Earth science research community, one of the primary goals of BEDI is to facilitate the use of this data by application-oriented communities and to support decision-making. NASA's approach to implementing the goals and objectives of BEDI is guided by two specific principles. First, open community-driven standards are the key to interoperability. Realization of the BEDI objectives is dependent upon converging on a set of standards that govern and describe relevant data formats, interfaces, and protocols. Second, to the maximum practicable extent NASA will design and execute its BEDI-related work activities so that the output and products of those activities are beneficial and useful to other federal agencies as well as to NASA.

To this end, EOSDIS and the DAACs identified datasets relevant to 12 different OSTP-identified Societal Benefit Areas and targeted those datasets for metadata, browse, and web service access enhancements. These enhancements include assignment of Digital Object Identifiers (DOIs), ensuring that metadata was current and complete, making the datasets available through OPeNDAP and, where appropriate, generating browse imagery to be included in the Global Image Browse Service (GIBS).

4. Information Viewpoint

4.1. Earth Science Information Models

EOSDIS has a plethora of information moving from system to system – data, metadata, service requests, documents, imagery, etc. Information models play a critical role in interoperability between systems. This section discusses key information models used within EOSDIS and the relationships between them.

4.1.1. UMM and Relationship to ISO metadata model

EOSDIS generates, archives, and distributes massive amounts of Earth Science data, which in turn is made available to the science community and the public at large. To aid in the search and discovery process, this data must be organized and cataloged, which makes accurate, complete, and consistent metadata a requirement for efficient accessibility. To improve the quality and consistency among its metadata holdings, EOSDIS has developed a model describing metadata that it archives and maintains. This model documents elements that may be represented across various metadata standards and unifies them through core elements useful for discovery. This Unified Metadata Model (UMM) has been developed as part of the EOSDIS Metadata Architecture Studies conducted between 2012 and 2013.

EOSDIS has a significant investment in existing metadata and systems built on existing metadata schemas unique to EOSDIS. Specifically, the GCMD DIF and the ECHO10 metadata models have been in widespread use for Collection and Granule metadata for many years and describe petabytes of existing data holdings. At the same time, interoperability with international systems and compliance with international standards like ISO19115 of are high value to EOSDIS.

The UMM provides a schema agnostic representation of the fields necessary to provide high quality metadata for EOSDIS holdings. The UMM defines key EOSDIS assets as UMM Profiles. For example, Collections are described by the UMM Collection Profile (UMM-C), Granules are described by the UMM Granule Profile (UMM-G), services by UMM-S, etc.

The UMM documentation identifies required fields, constraints on those fields, optional fields, overall descriptions of the semantics of those fields and guidance for how to develop and maintain high quality metadata. For each element defined in the UMM there is an element specification, description, profile utilization, cardinality, analysis, specification of any conflicts that were identified by the unification process, a mapping to existing standards, examples, and future recommendations.

At the time of this writing, the UMM also provides mappings to and from various formats.

While the UMM provides guidance on manually creating records in a given format, the more common approach is to leverage the CMR for automated transformations. The CMR fully implements the EOSDIS UMM and can provide automated UMM-compliant transformations between the supported formats. For example, a DAAC can submit metadata from a legacy system in ECHO 10 metadata format and a new, ISO19115 compliant application can consume that metadata through the CMR APIs leveraging the CMR to apply the appropriate UMM mappings. Similarly, a legacy client application can ask the CMR to provide new ISO19115 metadata in a legacy format such as the GCMD DIF.

Details and documentation on all of the UMM profiles described below can be found here: <https://wiki.earthdata.nasa.gov/display/CMR/UMM+Information>.

4.1.2. Collection metadata

Collection metadata is the anchor point for EOSDIS data holdings. The UMM-C serves as a reference model for geospatial science metadata for collections. This reference model is referred to as the UMM-C, where ‘C’ indicates that this is the collection model. The UMM-C attempts to unify several metadata formats (DIF9, DIF10, ECHO10, EMS, ISO 19115-2:2009, ISO 19115-1:2013). The model breaks down collections into elements or classes closely aligned to the ISO 19115-* Geographic Information Metadata schemas. As mentioned above, the CMR is used to provide a consistent format view across disparately formatted holdings.

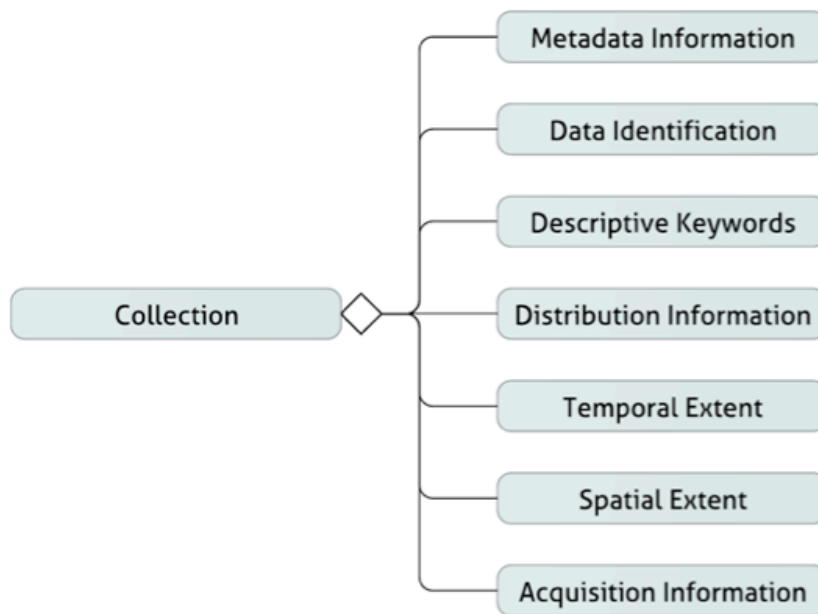


Figure 17: Top level elements of a Collection in the UMM-C

Collection metadata elements describe an entire set of data products or files; values of collection metadata apply to all of the granules in a specific collection. Collections may represent a release of any given data product, sets of data generated during an experiment, a campaign or an algorithmic test.

Any collection instance governed by the UMM-C may have relationships to other metadata formats. For example, each collection may have child granules, associated collections and may even include a parent collection that can serve as a discovery mechanism for closely related data products. In addition, each collection has a single instance of a “meta-metadata” profile, which has its own model, the UMM-M. The meta-metadata model provides information about the metadata to which it is associated. Meta-metadata includes fields such as metadata quality assurance information, associated tags, and metadata revision history.

The ISO 19115-2 and ISO 19115-1 mapping paths and extracts of XML encoding of this standard used by the UMM-C are derived from the NASA Best Practices ISO translation from ECHO to 19115-2 and a similar ECHO to ISO 19115-1 translation.

4.1.3. Granule Metadata

The UMM-G, where ‘G’ indicates granule model, serves as a reference model for geospatial science metadata for individual granules that are part of a larger data collection. As with the UMM-C, the UMM-G attempts to unify several metadata formats: ECHO 10 Granules, 19115-2:2009 , and ISO 19115-1:2013. The model breaks down granules into elements or classes closely aligned to the ISO 19115-* Geographic Information Metadata schemas.

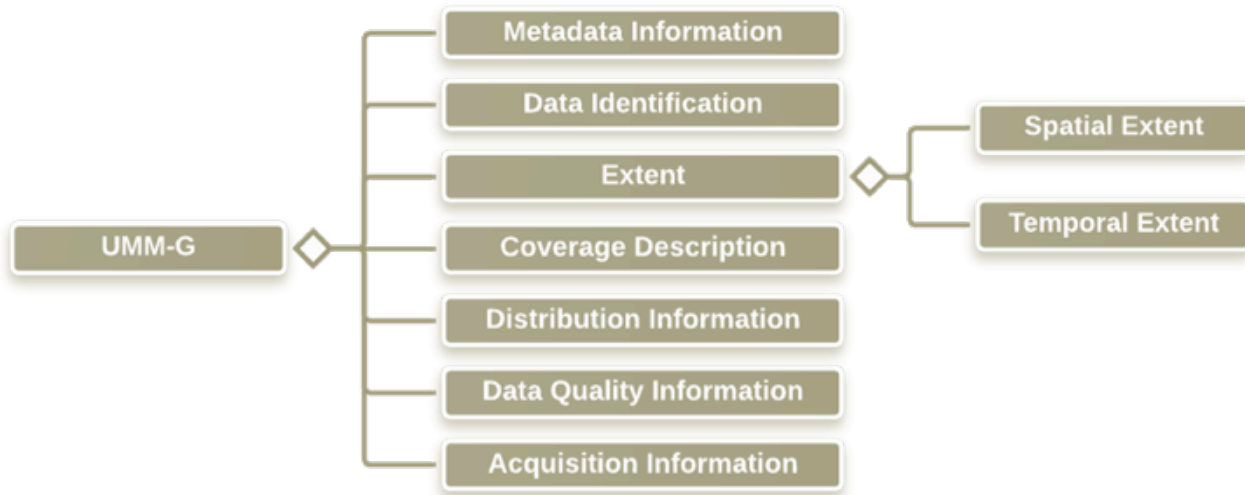


Figure 18: Top level elements of a Granule in UMM-G

4.1.4. Browse model

Being able to preview or visualize data is a major component of a user’s search and discovery workflow within EOSDIS. This preview rendering is called Browse Imagery in EOSDIS and exists for millions of granules. Browse imagery, or browse for short, is produced in several variants: thumbnail browse, full resolution browse, and high resolution “interactive” browse.

Thumbnail browse are simple low-resolution images typically mapping one-to-one with a science granule. Thumbnail browse are of varying scales, are not georeferenced, and may have graphical or textual overlays such as watermarks, coordinate system grids, or legends. Thumbnail browse are linked to science granules by Universal Resource Locators (URLs) in a UMM-G record. Thumbnail browse are intended to provide users with static snapshots of a granule helping them determine if the underlying data is appropriate for their needs.

Full resolution browse is similar to thumbnail browse except that it is generated at the same resolution as the underlying data. Full resolution browse may or may not be georeferenced and could contain overlays. Full resolution browse is typically used when a user wants a full resolution rendering of a predefined set of variables in a format more easily used than the underlying data. For example, full resolution browse are frequently made available as Georeferenced Tagged Image File Format (GeoTIFF) or Portable Network Graphics (PNG) files for common variables. Full resolution browse are linked to science granules by URL in a UMM-G record.

Full resolution interactive browse, sometimes called visual metadata, are full resolution georeferenced images suitable for map overlays or integration with other imagery and mapping toolkits. Within EOSDIS, GIBS is the largest source of visual metadata served over a standards compliant Web Map Tile Service (WMTS) API.

At present, imagery within GIBS is constructed into daily composites providing daily, seamless imagery for supported collections.

4.1.5. Usage model

The EMS is designed to support the ESDIS project management by collecting and organizing various metrics from the EOSDIS DAACs and other Data Providers. The EMS metrics are reported to NASA Headquarters to evaluate performance and effectiveness of the EOSDIS. Product Attribute Metadata flat files are the foundation for the EMS's reporting capabilities. Product attributes are metadata such as instrument, mission, product level, and discipline describing the characteristics of a data product. Accuracy of reports is highly dependent on providing comprehensive, consistent, and timely updates to the product attributes.

4.1.6. Processing Levels

EOSDIS data products are processed at various levels ranging from Level 0 to Level 4 . Level 0 products are raw data at full instrument resolution. At higher levels, the data are converted into more useful parameters and formats. All EOS instruments must have Level 1 products, most have products at Levels 2 and 3, and many have products at Level 4.

Data Level	Description
Level 0	Reconstructed, unprocessed instrument and payload data at full resolution, with any and all communications artifacts (e.g., synchronization frames, communications headers, duplicate data) removed. In most cases, the EOS Data and Operations System (EDOS) provides these data to the data centers as production data sets for processing by a DAAC or SIPS to produce higher-level products.
Level 1A	Reconstructed, unprocessed instrument data at full resolution, time-referenced, and annotated with ancillary information, including radiometric and geometric calibration coefficients and georeferencing parameters (e.g., platform ephemeris) computed and appended but not applied to Level 0 data.
Level 1B	Level 1A data that have been processed to sensor units (not all instruments have Level 1B source data).
Level 2	Derived geophysical variables (e.g. ice concentration, soil moisture) at the same resolution and location as Level 1 source data (e.g. swath).
Level 3	Variables mapped on uniform space-time grid scales, usually with some completeness and consistency.
Level 4	Model output or results from analyses of lower-level data (e.g., variables derived from multiple measurements) which are not directly measured by the instruments and spatially gridded.

Table 2: Data processing levels and their descriptions

4.1.7. Data Model

Generally speaking most EOSDIS data fall into three types of structures: Grid, Swath and Point, which are somewhat aligned with the processing levels. These three structures are also the core concepts around which the HDF-EOS data format (see Section 4.2) is constructed.

Grid

The grid data structure represents data at locations on a regular grid, based upon one of several Earth projections. This data type is most commonly used for products at EOS processing levels 3 or 4. Map projections presently supported include: Mercator(space oblique), Universal Transverse Mercator (UTM), interrupted Goode, polar stereographic, Lambert azimuthal (equal area), Lambert conformal conic, polyconic, transverse Mercator, Hotin oblique Mercator, and integerized sinusoidal grid.

Swath

The swath data structure comprises time-sequences of scan lines, profiles, or other array data. A swath has an “along-track” direction that is monotonic with time and usually a “cross-track” direction associated with each time instance. The cross-track direction need not be a straight line; a circular scanning instrument can utilize this data type by specifying its coordinates as a function of time. Swath data products are most often found at EOS processing levels 1 or 2.

Point

The point data structure is used for networks of locations that are irregularly spaced. Data product designers may define hierarchical “levels” that are similar to linked tables in a relational database. This feature allows the point data structure to be used with ocean buoy data that change locations with time.

4.2. Data Format Standards

From its inception, the EOSDIS system has leveraged standard data formats to contain system complexity and to enhance usability and interoperability. The initial data format standard used by EOSDIS was the Hierarchical Data Format, version 4 (HDF-4). HDF-4 enables a rich set of data structures, which are often important for handling the breadth of data structures needed to represent data from a wide variety of instruments and platforms. HDF is also a self-describing format, enabling the inclusion of rich metadata at every level of the data file. HDF-4 has since been largely supplanted by its successor, HDF version 5, or HDF5. HDF5 is also the storage format for netCDF-4, a widely used community standard, netCDF-4; thus most HDF5 files are also valid netCDF-4 files, meaning that they are usable in many of the tools available for netCDF, thus increasing the HDF format’s interoperability commensurately.

In addition to leveraging HDF as a storage format, EOSDIS conventions on data structures have been developed and used widely by NASA Earth Observing systems. HDF-EOS2, developed for HDF-4, provides for three main types of structures that are designed to account for the three main types of the data model: Point, Grid and Swath (see Data Model in the previous section). The more recent version is HDF-EOS 5, developed for HDF5, but presenting the same API to calling programs as HDF-EOS 2. These conventions standardize the representation of geolocation in the files, thus serving a role similar to that of the Climate-Forecast (CF) coordinate conventions. With the convergence of HDF5 and netCDF-4, use of the CF conventions is also beginning to grow within EOSDIS. CF and the HDF-EOS datasets have the advantage that the internal geolocation standardization makes the data amenable to generalized data subsetting, visualization and usage tools, both on the client side (e.g., Panoply) and the server side (e.g., OPeNDAP).

In addition to HDF and netCDF, a few datasets utilize other standard formats, often from a specific community or organization, shown in Table 3.

Type of Data	Format Standards
Level 0 packets	Consultative Committee for Space Data Systems (CCSDS)
Synthetic Aperture Radar (SAR)	Committee on Earth Observing Satellites SAR format
Global Navigation Satellite System and Doppler Orbitography and Radiopositioning Integrated by Satellite	Receiver Independent Exchange Format (RINEX)
Laser Ranging	Consolidated Laser Ranging Data Format (CRD)

Table 3: Data types and formats

In addition, a number of unique binary and ASCII formats are used, primarily for heritage datasets dating from the previous century, which often need case-by-case handling for data ingesting software and data usage tools. Work is ongoing within the ESO to develop more standardized forms of ASCII formats for use within EOSDIS. Further information on EOSDIS data format standards can be found at <https://earthdata.nasa.gov/about/esdis-project/esdis-standards-office-eso>

5. Service Viewpoint

5.1. Discovery Services

EOSDIS services are thoroughly documented to make it as easy as possible to consume them in diverse applications. Discovery services are generally the entry point to EOSDIS data and are exposed through several different APIs to fit a variety of use cases.

All discovery services are publicly accessible though not all data is open to anonymous access. Some data require user registration and an Earthdata Login profile (see Identity Management Services in Section 5.3 for more information). Clients wishing to track usage by user type can do so by programmatically requesting access tokens that are used with search requests to identify the user. In addition to user identification, some discovery services support client identifiers, notably the CMR. Client identifiers are self-assigned and are used by EOSDIS services to understand usage patterns and provide support when necessary.

5.1.1. CMR

The most capable, but largest, discovery API is the CMR Search API. The CMR search API is a RESTful API designed to be simple to use from a broad range of technologies. Like many RESTful web services, the CMR supports returning results in multiple formats. The CMR is capable of providing search results as:

- Simple JSON
- References only XML
- Collection native XML
- ECHO 10 XML
- ISO19115 XML
- DIF 9 XML
- DIF10 XML
- Comma-separated Value
- Atom formatted XML
- Opendata JSON (collections only)
- Keyhole Markup Language
- UMM based JSON (collections only)

The CMR search API centers around UMM Profiles, e.g. there are search paths for collections, granules, etc. All query strings support a basic set of parameters such as temporal conditions, spatial conditions, sort order, page result length, starting page, and query options. Individual profiles support additional options such as facets and highlighting support on collection searches and equator crossing date and online accessibility on granules searches.

For much more information on the CMR discovery API please see the CMR client developer documentation

available here: <https://wiki.earthdata.nasa.gov/display/CMR/Client+Developer+User+Guides+-+API+-+Technical+Documentation#>.

5.1.2. OpenSearch

EOSDIS provides an implementation of the Earth Science Information Partners (ESIP) Federation set of conventions around the Atom syndication specification and OpenSearch query/response specification to enable lightweight mechanisms of data advertising, discovery and access. This framework supports both a query/response pattern as well as a publish/subscribe pattern called collection casting. The OpenSearch specifications from www.opensearch.org have been adopted along with the draft specifications for Time and Geo(spatial) extensions, the Atom specification, and Dublin Core's specification for date/time range in the response and syndication cases. The result is a data advertising and discovery framework with low buy-in for search providers and low learning curves for client developers.

OpenSearch queries are based on an OpenSearch Description Document (OSDD), which describes how queries are made to a given search service. Response and casting conventions follow the Atom specification, with customizations to adapt the representation of data items in Atom. Some of these customizations are designed to help machine-level clients parse the search results.

OSDDs are XML files that are the key to specifying to clients how queries are to be done with a given search provider. The OpenSearch specification of OSDDs is available in detail on the OpenSearch site at http://www.opensearch.org/Specifications/OpenSearch/1.1#OpenSearch_description_document. The query specifications are contained within XML elements of type <Url>, which includes a template for the URL query and a mime-type for the expected response. The template gives the URL to execute the query, with placeholders indicated by curly braces '{' and '}'. The searchTerms parameter is expected to be a simple set of keywords to be used in a freetext search.

OpenSearch APIs are available from a number of providers and services, including the CMR and CWIC.

5.1.3. CWIC

The Committee on Earth Observation Satellite (CEOS) addresses coordination of the satellite Earth Observation (EO) programs of the world's government agencies, along with agencies that receive and process data acquired remotely from space. The Working Group on Information Systems and Services (WGISS) is a subgroup of CEOS, which aims to promote collaboration in the development of systems and services that manage and supply EO data to users world-wide. To realize a federated catalog for data discovery from multiple EO data centers, the CEOS WGISS Integrated Catalog (CWIC) system has been implemented. CWIC provides inventory search to WGISS agency catalog systems for EO data.

CWIC uses a mediator-wrapper architecture that has been widely adopted to realize the integrated access to heterogeneous, autonomous data sources. A data source archives data and disseminates it through the Internet. The wrapper on top of the data source provides a universal query interface by encapsulating heterogeneous data models, query protocols, and access methods. The mediator interacts with the wrapper and provides the user with an integrated access through the global information schema.

Based on the mediator-wrapper architecture, current version of CWIC has been developed and operational

with following data partner catalog systems: NASA Common Metadata Repository (CMR), NOAA Group for High Resolution Sea Surface Temperature (GHRSST), United States Geological Survey Landsat Surface Imaging (LSI) Explorer, Brazil Institute for Space Research Catalog System and Canada Centre for Mapping and Earth Observation. CWIC exposes these catalogs to client applications through a single OpenSearch endpoint.

For more information on CWIC please see: <http://ceos.org/ourwork/workinggroups/wgiss/current-activities/cwic/>

5.2. Online Data Access and Transformation Services

Efficient data access and transformation is key to utilization of EOSDIS data. To that end, EOSDIS provides a number of data access mechanism and support services. In the simplest case, a significant portion of the EOSDIS holdings are available online. Online access URLs are included in granule metadata and simply point to the underlying archived file, typically HDF-EOS or netCDF. Client applications can surface online access URLs directly to users or automate the retrieval of data on their behalf.

For some data online access isn't immediately available and the data must be requested from the responsible DAAC via a system of data ordering. The CMR is capable of brokering data orders to avoid ordering applications or users needing to know DAAC specific ordering systems. In some cases, data reformatting, subsetting or transformations are possible and can be requested along with a data order through ordering options. Ordering options allow clients to take advantage of offline tools like the HDF-EOS to GeoTIFF converter (HEG), MODIS Reprojection Tool (MRT), and Geospatial Data Abstraction Library (GDAL) Translate. Ordering is currently supported through a Simple Object Access Protocol API and is publicly accessible. For more information on data ordering APIs and options, please see <http://api.echo.nasa.gov/echo/ws/v10/OrderManagementService.html>.

In addition to basic online access to the underlying data, many DAACs offer sophisticated data access services such as OPeNDAP, w10n, WCS, WMS, and WMTS. These services are collectively capable of:

- Data reformatting
- Data reprojection
- Spatial, temporal, and parameter subsetting
- Image generation
- Direct data access

Not all data is available through all services; the responsible DAAC is the best source of information for what services are available for given data. For more programmatic discovery of EOSDIS access services the CMR will eventually associate UMM Service Profile (UMM-S) entries with data metadata allowing clients to discover supported services dynamically.

5.3. Identity Management Services

The EOSDIS Earthdata Login provides a centralized and simplified mechanism for user registration and profile management for all EOSDIS system components. End users may register and edit their profile information in one location allowing them access to the wide array of EOSDIS data and services. The EOSDIS Earthdata Login also helps the EOSDIS program better understand the user demographics and access patterns in support of

planning for new value-added features and customized services that can be directed to specific users or user groups resulting in better user experience.

Earthdata Login provides a user registration service allowing users to self-register with the EOSDIS. The registration is provided free of charge to the user. The user needs to set up a user ID, password and provide additional information, such as, user name, affiliation, country and a valid email address, in order to become a registered user.

A user that has not yet been registered with Earthdata Login may be directed to the Earthdata Login registration web page when they attempt to access the EOSDIS or a DAAC data service that requires user registration. The user is referred back to the original application after the registration has been completed successfully.

5.3.1. User Identification

Earthdata Login provides a user identification interface allowing applications to identify a user accessing their data and services when providing access. This allows EOSDIS to provide metrics on the various user categories obtaining data from EOSDIS. Earthdata Login provides users with Single-Sign-On (SSO) access to Earthdata applications. The SSO capability enables an Earthdata Login registered user to log in to Earthdata Login once, and access multiple Earthdata Login-integrated EOSDIS applications without being prompted to sign in for each application separately. The SSO capability in Earthdata Login is implemented using the OAuth2 protocol.

For more information on Earthdata Login please see: <https://urs.earthdata.nasa.gov/documentation>.

6. Outlook and Conclusion

Over the last 2+ decades, the EOSDIS system has continually evolved to keep up with both requirements and technology. As of 2016, EOSDIS's evolutionary goals are twofold:

1. Develop the capability to leverage elastic computing and storage in the cloud
2. Improve the User Experience at each step from data discovery to data access and usage

The first evolutionary direction is targeted toward cost containment and fine-grained cost control. In addition, the on demand resources and unprecedented scalability available to researchers provides an opportunity to significantly better exploit the vast amount of Earth remote sensing data in EOSDIS. Tools, systems, and infrastructure offer new and exciting techniques to deal with the petabytes of data available. Prototype implementations of systems architected to exploit elastic computing demonstrate the potential of applying scalable databases and machine learning techniques to decades of recorded science data.

The User Experience effort begins with the new discovery and access clients like Worldview and Earthdata Search, both designed to be interactive and responsive, giving users the ability to visually navigate through massive data holdings with their web browser or mobile device. EOSDIS emphasis on user experience goes much deeper than just visual web applications, however. A combination of exploiting elastic computing capabilities for data holdings and user experience work opens up exploration of other ways of interacting with EOSDIS holdings. For example enabling researchers to ask “When did this area last burn and how long did it take to recover?” across 40 years of Landsat data with rapid results and links to the science granules supporting the response.

In order to ensure efforts are aligned with user needs, usage analytics and metrics driven results are being embraced throughout EOSDIS to effectively measure what is actually happening in the enterprise and how things can be improved. For example, behind the scenes, efforts are also underway to improve the relevancy ranking of search engine results. Questions are being asked around how to best model user intent, how user interface selection should influence result ranking, how to best incorporate the deep knowledge in DAAC user services when providing information to users. For documentation centric applications, such as the Earthdata site, work is beginning to allow the CMR to index unstructured content coupled with structured metadata captured in the UMM-D. Having access to diverse metadata describing EOSDIS assets, the CMR can model relationships between items such as documentation about an instrument, collections captured using that instrument, and papers written leveraging that data. The CMR could surface this information in a manner navigable by EOSDIS Discovery clients, enabling users to find exactly the information they need across the significant EOSDIS holdings.

With a solid, scalable base of core services and the rich depth of knowledge shared throughout its DAACs, NASA's EOSDIS is leading exciting advances in data discovery, information access, and climate research.

7. Appendices

The following appendices have been identified as desired future expansion to the Technical Manual and are captured here for reference.

7.1. Adding new data to the EOSDIS Collection

Many of the EOS data sets were assigned to EOSDIS back in the 1990s through a review process, called EOSDIS Version 0. The software system was principally built to support the EOS on-orbit mission platforms. Most of the new on-orbit NASA Earth Science missions are directly assigned to EOSDIS by NASA Headquarters. The ESDIS Project engages early in the NASA mission lifecycle to manage the ingest, archive and distribution of data as soon as it becomes available from the new spacecraft platform (example ICESat-2 to launch in December 2017). Headquarters can also assign missions directly to ESDIS based on negotiations with other national and international agencies based on agreements (example the European Space Agency Sentinel series). The distributed nature of NASA Earth science requires that EOSDIS also be able to support new data acquired through all of the varied NASA Earth Science Division programs. NASA investigators can also request to add data to the DAACs following a review process. The process below describes how new datasets are selected to be added to the EOSDIS collection.

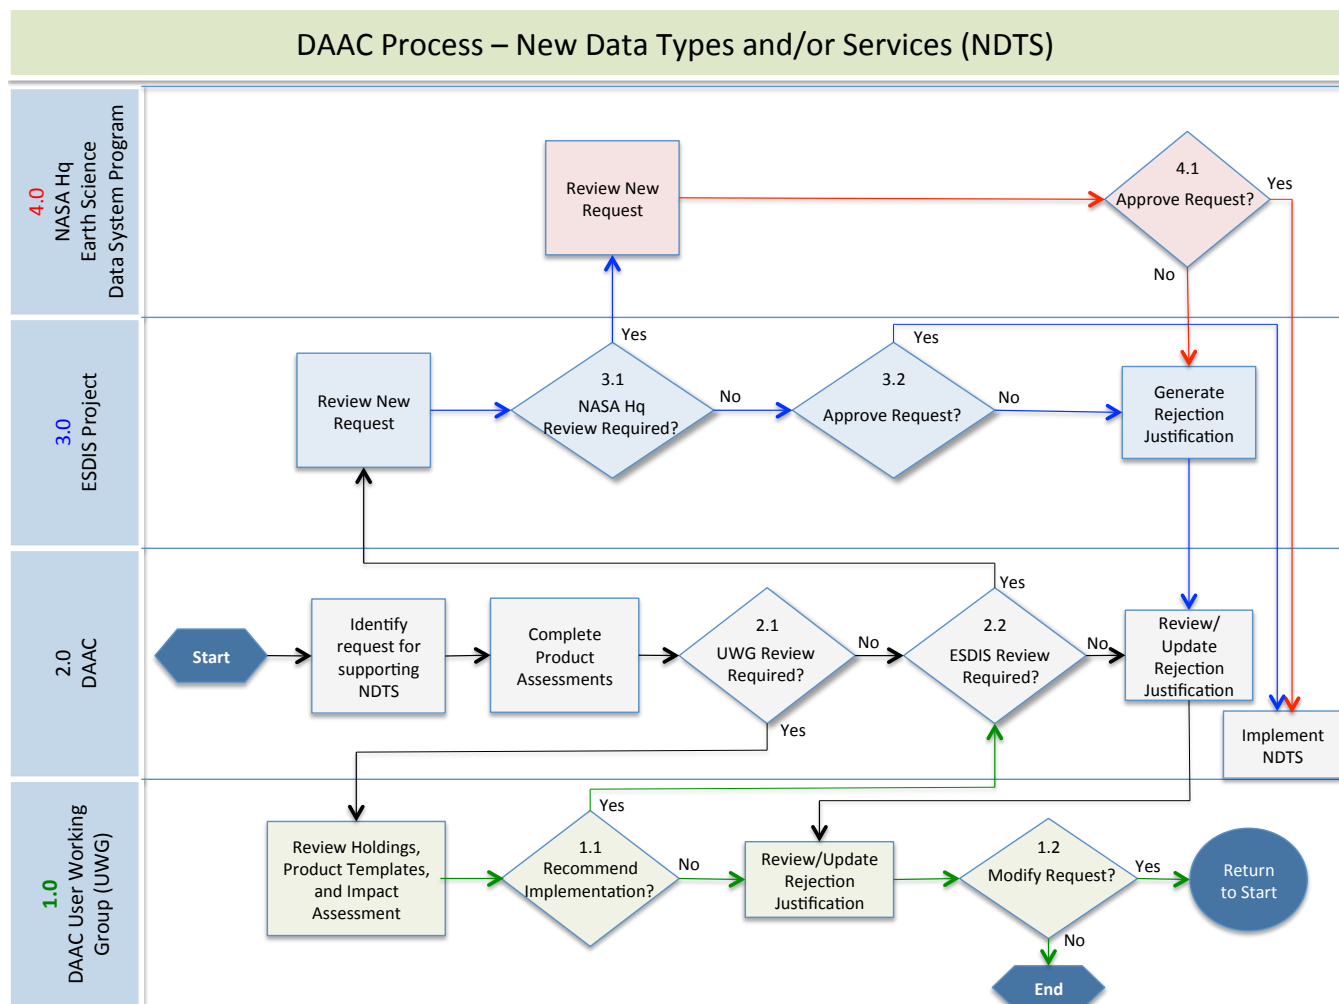


Figure 19. DAAC process for adding new data types or services to the collection.

7.2. Acronyms

ACRIM	Active Cavity Radiometer Irradiance Monitor
ACSI	American Customer Satisfaction Index
AGS	Alaska Ground Stations
AIRS	Atmospheric Infrared Sounder
AIST	(National Institute of) Advanced Industrial Science and Technology
AMSR	Advanced Microwave Scanning Radiometer
AMSR-E	Advanced Microwave Scanning Radiometer - Earth Observing System
API	Application Programming Interface
ASCII	American Standard Code for Information Interchange
ASDC	Atmospheric Science Data Center
ASF	Alaska Satellite Facility
ASTER	Advanced Spaceborne Thermal Emission and Reflection Radiometer
BEDI	Big Earth Data Initiative
CCSDS	Consultative Committee for Space Data Systems
CEOS	Committee on Earth Observation Satellites
CERES	Clouds and the Earth's Radiant Energy System
CF	Climate-forecast
CMR	Common Metadata Repository
CRD	Consolidated Laser Ranging Data Format
CSO	Communication Service Office
CSW	Catalogue Services for the Web
CWIC	CEOS WGISS Integrated Catalog
DAAC	Distributed Active Archive Center
DB	Direct Broadcast
DIF	Data Interchange Format
DOI	Digital Object Identifier

ECC Earthdata Code Collaborative

ECHO EOS Clearinghouse

ECS EOSDIS Core System

EDOS EOS Data and Operations System

EDSC Earthdata Search (Client)

EMSESDIS Metrics System

EMSn EOS Mission Support network

EO Earth Observations

EOCEOS Operations Center

EOS Earth Observing System

EOSDIS Earth Observing System Data and Information System

ERPS EOS Real time Processing System

ESDSWG Earth Science Data Systems Working Group

ESIP Earth Science Information Partners

ESMO Earth Science Mission Operations

ESO ESDIS Standards Office

ESSn EOS Science Support Network

FDS Flight Dynamics System

GCMD Global Change Master Directory

GDAL Geospatial Data Abstraction Library

GEO Group on Earth Observations

GEOSS Global Earth Observation System of Systems

GeoTIFF Georeferenced Tagged Image File Format

GES DISC Goddard Earth Sciences Data and Information Services Center

GHRC Global Hydrology Resource Center

GIBS Global Imagery Browse Service

Giovanni Geospatial Interactive Online Visualization AND aNalysis Infrastructure

GLAS Geoscience Laser Altimeter System

HDF Hierarchical Data Format

HDF-EOS Hierarchical Data Format - Earth Observing System

HEG HDF-EOS to GeoTiff Converter

HIRDLS High Resolution Dynamics Limb Sounder

ICD Interface Control Document

IST Instrument Support Terminal

JAXA Japan Aerospace Exploration Agency

JPL Jet Propulsion Laboratory

JSON JavaScript Object Notation

KMS Keyword Management System

KNMI Royal Netherlands Meteorological Institute (Koninklijk Nederlands Meteorologisch Instituut)

LANCE Land, Atmosphere Near real-time Capability for EOS

LIS Lightning Imaging Sensor

LP DAAC Land Processes Distributed Active Archive Center

LSI Landsat Surface Imaging

MISR Multi-angle Imaging SpectroRadiometer

MLS Microwave Limb Sounder

MODAPS MODIS Adaptive Processing System

MODIS Moderate Resolution Imaging Spectrometer

MOPITT Measurements of Pollution in the Troposphere

MRT MODIS Reprojection Tool

netCDF network Common Data Form

NGAP NASA-compliant General Application Platform

NISN NASA Integrated Services Network

NOAA National Oceanic and Atmospheric Administration

NRT Near real-time

NSIDC National Snow and Ice Data Center

OB.DAAC Ocean Biology Distributed Active Archive Center

OBPG Ocean Biology Processing Group

ODPS Ozone Data Processing System

OMIOzone Monitoring Instrument

OMPS Ozone Mapping & Profiler Suite

OPeNDAP Open Source Network for a Data Access Protocol

ORNL Oak Ridge National Laboratory

OSDD OpenSearch Description Document

PDR Product Delivery Record

PDS Production Data Set

PNGPortable Network Graphics

PO.DAAC Physical Oceanographic Data Center

RINEXReceiver Independent Exchange Format

SAGE Stratospheric Aerosol and Gas Experiment

SAR Synthetic Aperture Radar

SEDAC Socioeconomic Data and Applications Center

SGS Svalbard Ground Station

SIPS Science Investigator-led Processing Systems

SNPP Suomi National Polar-orbiting Partnership

SORCESolar Radiation and Climate Experiment

SSW Simple Subset Wizard

TDRS Tracking and Data Relay Satellite

TES Tropospheric Emission Spectrometer

UAF University of Alaska - Fairbanks

UMM Unified Metadata Model

UMM-C Unified Metadata Model - Collection

UMM-G Unified Metadata Model - Granule
UMM-M Unified Metadata Model - meta-metadata
UMM-S Unified Metadata Model - Service
URL Universal Resource Locator
URS User Registration System
USGEO United States Group on Earth Observations
UTM Universal Transverse Mercator
UWG User Working Group
VIIRS Visible Infrared Imaging Radiometer Suite
WCS Web Coverage Service
WFF Wallops Flight Facility
WGISS Working Group on Information Systems and Services
WMS Web Map Service
WMTS Web Map Tile Service
WMTS Web Map Tile Service
WSC White Sands Complex
XML eXtensible Markup Language

