EOS Science Working Group on Data

Report of the Data Distribution Workshop held at Fort Lauderdale, Florida, on February 1, 2001

Date of report: 10 September 2001

CONTENTS

1. INTRODUCTION

The EOS Science Working Group on Data (SWG) held its inaugural workshop on June 01-02, 2000, at the Goddard Space Flight Center (GSFC). That workshop concentrated primarily on data processing requirements for NASA's Terra mission. It also concluded that, while at that time there had been only limited data ordering beyond the needs of the science teams conducting validation, innovative approaches would eventually be needed to meet data distribution needs. It was recommended "that there be a meeting about six months from now to address data distribution status and archive access needs." This document reports on the resulting Data Distribution Workshop, which was held at Fort Lauderdale, Florida, on February 1, 2001, in conjunction with the EOS Investigators Working Group (IWG) meeting.

The objectives of the workshop were:

- To assess the current status of Terra data distribution.
- To identify immediate and foreseeable obstacles to meeting user data needs.
- To identify critical needs and areas for improvement and approaches for new development.
- To include DAAC Managers, DAAC User Working Group Chairs and EOS Instrument Science Data Representatives and ESDIS Project representatives in the deliberations.
- To develop a report summarizing current status and future needs, to be used for ESDIS and NewDISS program planning.

The workshop was designed to clarify a range of issues including meeting user needs at minimal cost; review current and planned distribution capacity; identify improvements to the ordering system; identify data subsetting needs; examine the user model applicable to tracking data orders and the dynamic response of users; identify ways to improve user services; assess user feedback on distribution; and assess alternative distribution options.

The workshop was led by the SWGD Chair, Dr. Chris Justice. A full list of participants is included in Appendix A, and the agenda is in Appendix B. The proceedings of the workshop can be divided into two major section:

A. Gathering status and feedback from all the parties participating in EOS data distribution, including NASA Headquarters; the EOS Program Office and ESDIS Project at GSFC; the DAAC User Working Groups (UWGs); and DAAC management.

B. Using discussion groups in a breakout session to address specific topics in data distirbution, producing specific suggestions and recommendations.

Most of the individual presentations given at the workshop are accessible on the SWGD's web site (http://swgd.gsfc.nasa.gov). Some of these presentations include considerable detail with respect to system status, requirements, and issues. This report does not attempt to replicate every item in those presentations, but serves as a means of summarizing the overall workshop proceedings and its recommendations.

2. PART A: STAKEHOLDER STATUS AND FEEDBACK

2.1 Current EOSDIS status

This topic was addressed by Skip Reber (ESDIS Project Scientist), Mike Moore (ESDIS Project), and Vanessa Griffin (ESDIS). Metrics for EOS data production and distribution are available on line at http://spso.gsfc.nasa.gov/edgrs. For all data centers combined, the number of orders delivered increased from 4 million in FY98 to 6 million in FY99, and 9 million in FY00. Currently, over 1 Terabyte/day is being distributed from the DAACs at NASA Goddard Space Flight Center, NASA Langley Research Center, and the USGS's Eros Data Center (EDC). This includes output from both the EOSDIS Core System (ECS) at all three sites, and the Science Information Partner Systems (SIPSs) at Goddard and Langley. Nevertheless, the requests for data are far below the system capabilities, presumably because the initial validation of Terra data products is still in process and because users are still becoming familiar with the data products. It was also reported that: Only a small number of media types are available; there have been only limited outreach activities by the DAACs and instrument teams; the subsetting capability is not yet available; the EOS Data Gateway (EDG) data ordering system both has a steep learning curve and is awkward to use; many products do not have associated browse image data; tools to handle HDF and HDF-EOS data are limited and not widely known; and the high percentage of failed orders discourages all but the most determined users

Augmentation of the ECS is under way that will provide improvements in a number of areas:

- *Product Distribution System (PDS):* To be operational in the April time frame, this will extend the range of physical media to include CD-ROM, DVD, and DLT, as well as the initial 8 mm tapes, at GSFC and LaRC, with the number of drives can be extended. The PDS is an EDC development that is being installed at all DAACs.
- *Subsetting:* Spatial subsetting for images is being implemented by Sara Graves's group at The University of Alabama at Huntsville, and will be integrated into the ECS. It will be hosted initially at GSFC. While the requirements are being based on inputs from individual teams, concern was expressed by some present that if these inputs are not adequately comprehensive, then the present initiative should be viewed as only a partial solution.
- *Data pools:* These will be caches of frequently requested data available on line to users via the Internet. ESDIS is working with the DAACs to define these data pools. At this stage only concepts are being developed, and there are not yet any prototypes. ESDIS plans to have some of the data pools operational later this year.

2.2 Feedback from DAAC User Working Groups

The DAAC User Working Groups (UWGs), Science Advisory Panels (SAPs), or similar user organizations provide liaison between DAAC/EOSDIS staff and the scientific community. They have a vital role, not only in facilitating the requirements and issues of the users, but also in assisting the DAACs to determine user needs, interface requirements, and priorities. They can also provide a feedback mechanism between the science community and the EOS instrument teams regarding the products.

UWGs are active at all of the DAACs. Reports from the UWG representatives indicated that the ordering system works as well as might be expected within its limitations. After gaining experience with the interface, users appear to find data easy to select. The UWG representatives also believe that user services support is generally good. There are also numerous issues still being worked or remaining to be worked, such as:

- the limited throughput of the LaRC DAAC;
- difficulties in ordering particular orbits;
- difficulty in ordering collocated data from multiple instruments;
- the steep learning curve for new users to order data;
- the lack of similarity in file naming conventions between Terra instruments;
- difficulty in tracking calibration changes;
- concern about limiting the volume of distributed products to 1X;
- concern about eventual migration of data from EOSDIS to the long term archives;
- concern about the way larger orders (greater than 15 GB) need to be broken up;
- the slowness of FTP data deliveries;
- the limited time for which some ordered data is available for FTP pull;
- the lack of subscription services, particularly geographically qualified subscriptions.

These are just some of the issues raised. Many of them are being worked by the DAACs/SIPSs and/or by ESDIS/ECS, while others represent system limitations that have yet to be addressed. Further work is needed to categorize and prioritize these issues.

It was not the intention of the workshop to repeat the experience of the EOS tools workshop held three days earlier. Nevertheless, limitations of the available tools for data access were mentioned repeatedly, and this is seen as one of the major obstacles to increased data use. Clearly there are tools available but they do not yet have the primary functionality needed. Tools are addressed below in Part B of this report.

Much of the UWG material for the workshop was based on surveys of UWG members. The number of respondents was relatively low; for example, the LaRC UGW received only five responses. This was presumably indicative of the relatively low number of persons actively ordering data at this stage of the mission, or the weakness of surveys to obtain the user feedback.

2.3 Immediate needs for the EOS distribution system

The needs of the respective DAAC-based distribution systems have highly individual characteristics. At the JPL DAAC, current distribution capacity is adequate, and there is an extensible architecture that could be readily augmented if additional throughput were needed. Both the LaRC and GSFC DAACs are experiencing archive and distribution throughput well below the installed capacity, and therefore do not have issues in that area. For example, data distribution at the GSFC DAAC during the year 2000 averaged 23 GB/day compared with an installed capacity of 446 GB/day. A survey carried out by this DAAC during the recent AGU meeting revealed that reasons for the lower-than-projected distribution include:

- waiting for the products to mature;
- difficulty in using the system for accessing data;
- lack of a subsetting mechanism that will deliver files of manageable size;

- unfamiliarity with the HDF-EOS file format.

All of these factors may be associated with the newness of the Terra mission, and can be addressed through improving available tools and facilities, and keeping the community informed of current system capabilities.

Immediate distribution issues raised include:

- Orders larger than 15 GB typically fail or are rejected. As a result, such large orders need to be handled manually, e.g. by dividing into a series of smaller orders.
- More popular media types are required. This issue is being addressed by the forthcoming Product Distribution System (PDS), which will add CD-ROM, DVD, and DLT. However, the PDS will not address ASTER on-demand products, and that need remains.
- The FTP distribution capabilities at the EDC DAAC need to allow a 72-hour on-line residency for the distributed products, to accommodate delays in the user pull, such as occurs over long weekends.
- The distribution system at EDC requires additional capacity to meet requirements, and this expansion is in progress.

Other distribution needs include user-accessible subsetting; improved EDG searches; improved order tracking and error recovery; improved data availability mechanisms such as data pools, data mining, regional redistribution, allowing users to create bookmarks that other users can use to search and order archived granules, and on-demand processing.

Many of the needs described in this section either are being worked or are addressed explicitly or implicitly in Part B of this report. For that reason, recommendations from the workshop are considered to be covered in Part B.

2.4 The NewDISS perspective on data distribution

The NewDISS team was chartered in August 1998 to produce a plan for how NASA's Earth Science Enterprise can best make data and information available in a timely manner during the coming decade. While the NewDISS plan is still undergoing review, its concept revolves around creating a flexible and responsive system that allows for a spectrum of heterogeneous approaches, utilizing standard interfaces to facilitate a workable and highly distributed infrastructure.

Within this context, individual investigators will be empowered to utilize the open interfaces to facilitate system operability, including distribution. A published, open "Standards and Practices" will be included in Research Announcements for Mission Data Systems and Science Data Centers.

At this time, we have an existing system with its combination of centralized and distributed systems that make up a "federation" of facilities. While this provides the many "lessons learned," from which NewDISS is specified, it is also the starting point for the evolution to NewDISS. This is evident in various innovative, low cost solutions already operating, such as the SeaWIFS reprocessing

NewDISS is a work-in-progress that will bring greater innovation and extensions to the existing operational architecture, providing a means to implement requirements trades regarding science priorities, a faster process for product refinement, and the augmenting of existing distribution capacity.

3. PART B: SELECTED ISSUES AND RECOMMENDATIONS

The second of the two major parts of the workshop involved discussing the process for resolving the various issues in data distribution, and making specific recommendations. To facilitate this, the workshop divided into three discussion groups to address selected priority issues: software tools to facilitate distribution and early use of data; user modeling of community access patterns; and creative solutions to current and projected obstacles to distribution

3.1 Software tools to facilitate distribution and early use of data

The EOS data products, which are in the HDF-EOS format, are not always regarded as easy to read and handle. Some tools to assist with this are available, but their capabilities are limited. The workshop therefore considered the process for developing software tools, and for providing access to those tools via the DAAC User Services. Three categories of tools were identified: data product search and order tools; format conversion tools; and data manipulation tools. It is not considered appropriate that the DAACs should provide EOS data users with comprehensive image processing and data analysis packages.

A three-step process is recommended for providing software tools that offer basic capabilities:

1) *Survey of available tools:* The first step is for each DAAC to make a comprehensive survey of tools to support their respective data products. It should include tools already available at the DAAC, plus tools from other DAACs and external sources. The Global Change Master Directory (GCMD) has an extensive source of information about tools that would be useful in this survey. The results should be discussed at a forthcoming DAAC managers meeting, after which each DAAC manager should present a list of available and potentially applicable tools to their User Working Group.
2) *Evaluation:* The available tools should be evaluated, primarily by the UWGs, to identify those tools of greatest value. Each UWG should recommend which tools should be supported by their respective DAAC. The required support would vary, and could include options such as posting of share-ware tools for user downloads; maintenance of the tools; user assistance through the DAAC User Services; and referral to commercial vendors. Funding to provide this support would need to be included within the DAAC budget, either as reprogrammed funds or as an augmentation. In completing its evaluation, each UWG should identify critical capabilities not afforded by available tools.
3) *Development:* Tools that are not available will need to be developed. Because the provision and support of software tools will fall ultimately to the DAACs, much of the responsibility for tool development is likely to reside with the DAACs, in conjunction with their UWGs and the Instrument Teams, with funding sought through budget augmentation. At the same time, cross-DAAC requirements need to be identified and coordinated by the DAAC managers at their regular meetings. There are also other ways that tools development could be sponsored, such as direct solicitation by NASA Headquarters; the provision of NASA seed money for commercial vendors to incorporate HDF and HDF-EOS functionality into their existing products; or through volunteers. The open competition of funding for tool development should also be considered. In general, providing the DAACs with responsibility and funding to develop and/or support tools is likely to be an expeditious approach to putting

tools into the hands of those that use EOS data products, but other options should also should be addressed.

3.2 User modeling and community access statistics

The primary application of user modeling is in making resource allocations at the DAACs more effective, such as to aid in identifying current and future stress points; to develop mitigation strategies; and to help users and DAACs become more efficient. Without a successful modeling technique, systems can be wrongly sized and funding mistakes may occur.

Prior to the discussion group, an approach to predicting quantitative user access and data distribution rates was presented to the workshop by Dr. Bruce Barkstrom. This method uses a generally accepted marketing research model, known as “innovation diffusion” to estimate the time evolution of the fraction of a population that will adopt a new way of doing things, such as ordering data from a DAAC. While Barkstrom has conducted numerical experiments with several difference potential user populations, the current EOSDIS ordering statistics are consistent with the notion that most current EOSDIS users come from the sc ientific community. The estimate of the size of this population is based on the numbers of members of professional scientific societies who publish data. It includes a model for various order patterns, allowing the model user to apportion different ordering patterns to different product types. Thus it can explore differences in the ordering patterns for long time series versus limited space-time data sets. The typical EOS file of around 80 Megabytes is taken into account. The resulting model gives an encouragingly close comparison of actual versus predicted data volume distribution. It also indicates that most of the people who will use the data are already doing so. Current statistics are consistent with the view that science data users are currently the majority user community, and will continue to be so, compared with students and commercial users. An implication from the model is that data distribution should be segregated based on user ordering patterns, such as in judicious use of on-line data caches, selectable subsetting, and special services for very high volume users.

Although there was discussion about the correctness of the model, it was pointed out that the collection of real-mission statistics is only beginning, so that there will be an adaptive system that has flexibility to use current resources in the most appropriate manner. Meanwhile, models such as Barkstrom’s are important, and demonstrate the need to tailor the system to meet the need. The question then arises whether the best statistics are being gathered to allow a sophisticated tailoring of the system, in say a year from now. The ESDIS Project does not have a model that is used at this time to forecast system evolution.

The discussion group identified a process for developing user models that would require work by the DAAC User Services, particularly through the DAAC UWGs, DAAC User Services Working Groups (USWGs), and the DAAC managers. Typical steps in this are.

1) *Use of current statistics:* Improve the use and collection of metrics specific to DAAC user “Tribes”. The effectiveness of this approach will depend upon what statistics exist now, and the ways these can best be used to improve our understanding of current users, to predict how current users will behave in the future, and to identify better other user communities. There is a need to determine what additional statistics should be collected.

2) *Alternative models:* It would be useful to develop alternative models to the one presented at this workshop by Dr. Barkstrom, and to extend his work, so that there is a range of options to draw upon.

3) *Presentation:* A presentation or a workshop on the findings should be presented to one or more of the DAAC UWGs or USWGs, so that the models can be refined. This may be possible in the May 2001 time frame, particularly in conjunction with the USWG biannual meeting in Huntsville, Alabama.

4) *Tribal profiles:* The development of "tribal" profiles of users would permit improved modeling techniques, depicting the various patterns of data access with respect to both the types of data and the types of data files used. A formalizing of the individual cases could result in services individual to each particular community, such as coding examples for file use.

3.3 Creative solutions to current and projected distribution obstacles

This topic was designed to suggest processes and techniques for resolving issues in data distribution that are not addressed by current plans or developments. The guiding principles of such solutions are to maximize science results and to maximize the resulting science-based applications. The range of challenges requiring creative solutions is significant, and includes issues such as:

- Maximizing distribution in line with the expected rapid growth in capacity
- Maximizing the usefulness of EOS data
- Overcoming consumer resistance
- Reaching out to more users without proliferating products and proliferating volume
- Ensuring a balance between making sure there is enough use but not overwhelming use
- Remaining flexible in the face of rapidly evolving and changing user communities
- Convincing the user community to use the available data sets, educating the user

3.3.1 Classes of solution

The workshop discussion group provided a forum to brainstorm regarding specific types of creative solutions. In that context, the following are possibilities that could be developed into specific opportunities for solutions.

1) *Other groups:* Contributions from groups other than the existing DAACs and UWGs can be encouraged, e.g. from ESIPSs, data brokers, relevant interdisciplinary science (IDS) groups, RESACs, and others. Contributions might include value added products; helping to distribute standard products, e.g. partial or full mirror sites.

2) *Direct broadcasting:* Direct broadcasting from EOS satellites can be encouraged, along with software tools to use the data received. Basic processing tools and algorithms are already available. Those who receive the direct broadcasts or use the data can be encouraged to redistribute their data; this practice can be developed from the existing direct broadcasting. The greater availability of direct broadcasting may lead to greater general availability of near-real-time data sets. As higher-level science algorithms/products developed by the EOS

science teams mature, they should be added to the standard suite of programs that run within the direct broadcasting systems.

3) *Greater distribution efficiency:* Existing data and service providers can be encouraged to be more efficient, and to provide better services and products. Typical technological mechanisms for this include:

   - Data pools
   - Knowledge-based data extraction, i.e. data mining
   - Implementation of existing solutions using novel approaches, e.g. large numbers of PCs and computers working at night
   - Data compression
   - Coincident data search from multiple sources

4) *Creative financing:* One possibility may be commercial not-for-profit organizations

5) *Breaking mechanisms:* This technique may not be needed, but is included for completeness. It involves sliding scales of data availability depending on the size of the data sets required. Typically this relates to pricing policy, and could include consideration of the genuineness of need. It may be pertinent to place barriers in the way of some types of user, such as requiring written proposals from non-government users, and/or Memoranda of Understanding (MOUs) with other agencies.

6) *Innovative products:* Certain products might be developed in an innovative manner. They include development of refined data sets, such as time series, maps, continental scale mosaics, and other value-added products.

7) *Improved data user tools:* This is addressed by an earlier section of this report.

8) *Anticipating the user model:* This is addressed elsewhere in this report. It can ensure that the science users' needs are met and that data applications users are not promoted unduly at the expense of science users

3.3.2 Processes and mechanisms

Here are some candidate ideas for ways in which the process for developing and maintaining creative solutions can be encouraged.

1) *Leadership by the DAACs:* The DAACs should take the lead in better defining existing and future needs of the users. Multiple mechanisms should be used, including the UWGs, using user responses to define the needs. The resulting options should be prioritized and categorized according to feasibility. This activity can be linked with or otherwise use the Federation and the RESACs.

2) *Maintain the momentum:* The momentum of the workshop in defining innovative ideas and reviewing progress should be continued. The mechanism for this needs to be defined, and is not necessarily the participants of the discussion group at this workshop. The full spectrum of participants needs to be involved, including NASA Headquarters, ESDIS, DAACs, DAAC users, science users, instrument teams, and the broader community. The IWG meetings may provide the needed opportunity to sponsor ongoing meetings on this topic. Specific activities

include conceptualizing the innovations; implementing the innovations; and reviewing and culling the evolution of the innovations.

3) *Continued data sets:* There can be a mechanism to ensure a continued supply of data sets from non-traditional supplies, with a mechanism for assessing the risks of stopping the supply of new products.

4) *Criticality assessment:* A program management mechanism is needed to assess the criticality of data sets and products, ensuring that long term archiving is possible, and deciding what resources are needed to maintain the data holdings.

5) *A suitable environment:* It is important to have an environment that allows for innovative solutions, even with mechanisms that allow for unconventional ideas, so that innovative technology is engaged where appropriate. Special, low-level funding for innovative ideas is a way to do this.

6) *Engage the non-traditional:* The data system climate between NASA and non-NASA groups can be changed by engaging the non-NASA community directly in our activities, and utilizing their in-built organizational capabilities.

In conclusion, some form of "ecosystem diversity" needs to be encouraged to maintain sustainability. While the same size does not fit all, there is a limited number of sizes that suits nearly everyone. The challenge is to provide the range of products and services that facilitates the distribution of EOS data to the widest range of users in the most efficient and productive manner.

4. CONCLUSIONS AND NEXT STEPS

This report, along with material presented at the workshop, will be made publicly available via the SWGD's web site (http://swgd.gsfc.nasa.gov) so that discussion of the recommendations can proceed. It is hoped that a constructive dialog can ensue that will result in a formal management response to the workshop recommendations. The success of the SWGD depends upon a genuine interaction between the EOS teams represented by the SWGD and the program and project management, so that issues relating to success of the respective EOS missions are resolved effectively. The prospects for this are promising, based on the senior level participation in the workshop by representatives from NASA Headquarters and the GSFC EOS Program and Project Offices.

This was the second SWGD workshop, and the success of this type of forum indicates that future workshops should be considered. Some potential topics include:

- *EOS user models*: This may be in conjunction with UWGs. Presentations and discussions at USWG meetings in May are already planned.
- *Software tools:* This workshop would follow completion of the tools survey recommended above, and may be possible as early as July 2001. It would include both NASA Headquarters and the GSFC ESDIS Project as potential funding sources, as well as the responsible DAAC managers.
- *Terra long-term archiving:* The cost of long-term archiving is not an issue now, but it is projected to become so in the next few years. This topic may belong more appropriately to NewDISS as a "creative solutions" issue.

The core of future SWGD considerations is embodied in the above discussion on creative solutions. The discussion at the current workshop represents only a first step, and is primarily at the level of potential possibilities that will require ongoing elaboration, review, and maintenance to ensure a successful evolution of progress. A meeting focusing on creative solutions is therefore proposed. The timing for such a meeting needs to be worked out. It could, for example, be held in conjunction with the next IWG meeting, which is nine months from the current workshop.

The next meeting should also address the recommendations of the previous workshop on production, and of this workshop. The present workshop did not include adequate coverage of the previous recommendations on production. There is a need to continue focusing on implementation of the SWGD recommendations, and the priority they take in relationship to other ESDIS funding.

In conclusion, a reminder about the working of the SWGD is timely. This is a voluntary group dedicated to ensuring the success of the EOS series of NASA missions. The group has no resources to direct the development of capabilities and processes, but it does seek to work with the existing organizational, developmental, and operational structures to get the job done. In doing this, it is pertinent to work for a community consensus and to assist in communicating that consensus and its related proposals to NASA Headquarters. It is important that no opportunities be left unutilized, or underutilized.

The final business of the workshop was an announcement that the chairmanship of the SWGD will pass from Chris Justice to Graham Bothwell at conclusion of the workshop. The SWGD charter specifies that this position is an annual appointment. The Working Group is especially grateful for the dedicated efforts of Chris Justice in facilitating this and the previous workshops.

APPENDIX A: ATTENDANCE AT THE WORKSHOP

Name	*Organization*	*E-mail*
Bruce Barkstrom	NASA Langley	b.r.barkstrom@larc.nasa.gov
Andy Bingham	JPL	andrew.bingham@jpl.nasa.gov
Graham Bothwell	JPL	
David Bromwich	Oregon State University	bromwich.1@osu.edu
Donald Collins	JPL	djc@seaanchor.jpl.nasa.gov
John Dwyer	EROS Data Center	dwyer@usgs.gov
Bob Evans	University of Miami	r.evans@miami.edu
Jennifer Francis	Rutgers University	francis@imcs.rutgers.edu
Vanessa Griffin	GSFC	vanessa.l.griffin.1@gsfc.nasa.gov
Jim Irons	GSFC	jim@ltpmail.gsfc.nasa.gov
Chris Justice	University of Virginia	justice@maddux.evsc.virginia.edu
Tom Kalvelage	EROS Data Center	kalvelage@usgs.gov
Steve Kempler	GSFC	steve.kempler@gsfc.nasa.gov
Katalin Kovacs	GSFC	kkovacs@forest.gsfc.nasa.gov
Martha Maiden	NASA HQ	mmaiden@hq.nasa.gov
Richard McGinnis	LaRC	r.s.mcginnis@larc.nasa.gov
Dolly Perkins	GSFC	dolly.perkins@gsfc.nasa.gov
Moshe Pniel	JPL	moshe.pniel@jpl.nasa.gov
H. K. Ramapriyan	GSFC	ramapriyan@gsfc.nasa.gov
Jon Ranson	GSFC	jon.ranson@gsfc.nasa.gov
Skip Reber	GSFC	reber@skip.gsfc.nasa.gov
Greg Scharfen	NSIDC	scharfen@kyros.colorado.edu
Chris Scolese	GSFC	christopher.scolese@gsfc.nasa.gov
Diana Starr	NSIDC	starrd@kyros.colorado.edu
John Townshend	University of Maryland	jtownshe@glue.umd.edu
Diane E. Wickland	NASA HQ	dwicklan@hq.nasa.gov
Robert Wolfe	GSFC	robert.wolfe@gsfc.nasa.gov
Daniel Ziskin	NCAR	ziskin@ucar.edu

APPENDIX B: AGENDA OF THE WORKSHOP

- Introduction and objectives — Chris Justice
- Summary of OES data distribution status — Skip Reber
- User modeling: Empirical and Theoretical Study of Data Ordering — Bruce Barkstrom
- Feedback on Terra distribution by the DAAC User Working Groups
  - NASA Langley DAAC — Jennifer Francis, Dan Ziskin
  - NASA GSFC DAAC — Wayne Esaias
  - EROS Data Center DAAC — Jim Irons, John Dwyer
  - NSIDC DAAC — David Bromwich
  - JPL DAAC — Bob Evans
- Improvements and Immediate Needs for the EOS Distribution System
  - NASA LaRC DAAC — Richard McGinnis
  - NASA GSFC DAAC — Steve Kempler
  - EDC DAAC — Tom Kavelage, John Dwyer
  - NSIDC DAAC — Greg Scharfen
  - JPL DAAC — Don Collins
- Planned EOS distribution capacity and capabilities — Mike Moore, Vanessa Griffin
- Data Distribution: The NewDISS Perspective — Martha Maiden
- Prioritization of Issues — Jon Ranson
- Discussion Groups Addressing Selected Priority Issues
  - Software tools to facilitate data distirbution — Jim Irons
  - User modeling — Bruce Barkstrom
  - Creative solutions to distribution obstacles — John Townshend
- Reports by the Discussion Groups — Discussion Group Chairs
- The Way Forward — Chris Justice