170-WP-005-002

Data Assimilation System Integration with the ECS

White Paper

White paper - Not intended for formal review or Government approval.

May 1996

Prepared Under Contract NAS5-60000

RESPONSIBLE ENGINEER

George Elengical /s/ 5/30/96

George Elengical, Systems Engineer Date
EOSDIS Core System Project

SUBMITTED BY

R. E. Clinard /s/ 5/31/96

Robert E. Clinard, ECS CCB Chairman Date
EOSDIS Core System Project

Hughes Information Technology Systems
Upper Marlboro, Maryland

This page intentionally left blank.

Abstract

This is a Data Assimilation System (DAS) specific White Paper covering all aspects of the procurement of DAS equipment, the definition of the interface between the DAS and the Release B version of the Earth Observing System Data and Information System (EOSDIS) Core System (ECS), deployment of the DAS equipment, and the check out of the interface between DAS and ECS.

This document starts with an overview of the ECS and the DAS. It then describes the interface between ECS and DAS and goes on to explain how the DAS takes advantage of the services provided by ECS. This section is followed by a description of the architecture of ECS and how this architecture provides an optimum environment to support the generation and distribution of DAS products. Appendices A and B elaborate on the key drivers that were used for arriving at the interface between DAS and ECS. Next, the deployment schedule with various dependencies is discussed and an optimum plan is presented. An Operations Concept for DAS is presented with details of how ECS streamlines the DAS production process. The White Paper concludes with a plan for the Science Software Integration and Test (SSI&T).

Keywords: Operations, Scenarios, DAACs, Mission Support, Release B, Interfaces, DAO, DAS

This page intentionally left blank.

Contents

1. Introduction

2. Related Documentation

3. The ECS and DAS

4. Architecture

5. DAS Schedule

6. DAS Operations Concept

7. Science S/W Integration and Test

8. Open Issues

Figures

Tables

1. Introduction

1.1 Purpose

The purpose of this White Paper is the following:

- Define the interface between the DAS and the ECS beyond the Release B implementation of the ECS
- Identify the ECS utilities and Toolkit that will considerably simplify the architectural design of the DAS
- Identify areas where ECS will need enhancements in order to support DAS
- Provide an end-to-end description of how ECS supports different modes of operations related to the DAS
- Document the design rationale for the areas where ECS interfaces with DAS
- Describe the deployment plan and schedule for getting DAS integrated with ECS in time for the AM-1 launch
- Present an operations concept for the DAS after deployment at the Goddard Space Flight Center (GSFC) Distributed Active Archive Center (DAAC)
- Describe integration and testing methodology for getting science algorithms incorporated into DAS

The Data Assimilation Office (DAO) is responsible for developing advanced assimilation algorithms used to produce research-quality assimilated data products like the multiyear global atmospheric data sets. In the Operations mode of DAS, data from various sources (e.g. National Oceanic and Atmospheric Administration (NOAA)) are provided to the DAS. Some of the DAO-acquired datasets are required by ECS as ancillary data for ECS production. This document will explain how the ECS, located at the GSFC DAAC, stages data needed by the DAS, archives data needed by DAS, and distributes DAS generated products to the end-users. In addition, this document will describe how the ECS at EROS Data Center (EDC) performs post-processing of Moderate-Resolution Imaging Spectrometer (MODIS) data (using software provided by the DAO) needed by DAS before being forwarded to the ECS at GSFC.

1.2 Scope

This document defines various aspects of the deployment of DAS related equipment that will be installed at the GSFC DAAC. The document establishes the schedule of activities that will lead to the integration of DAS equipment at the GSFC DAAC. The operations concept defined in this paper will supersede the current operational concept for the DAS. This White Paper is not a deliverable item as identified in the Contract Data Requirements List (CDRL) for the ECS project and therefore, doesn't have a Data Item Description (DID) associated with it. The ECS architectural description included in this paper is for reference purpose only and in case of discrepancy, documents generated under DID 305 will supersede this white Paper.

1.3 Organization

This paper is organized as follows:

Section 1 -- Statement of purpose, scope, and overall organization.

Section 2 -- Listing of referenced and applicable documentation

Section 3 -- The DAS and ECS

Section 4 -- Architecture

Section 5 -- DAS Schedule

Section 6 -- DAS Production

Section 7 -- Science S/W Integration and Test

Section 8 -- Open Issues and Roadmap to Closure

Appendix A -- DAS Data Flows

Appendix B -- Trades and Studies

1.4 Review and Approval

This White Paper is an informal document approved at the Office Manager level. It does not require formal Government review or approval; however, it is submitted with the intent that review and comments will be forthcoming.

2. Related Documentation

2.1 Parent Documents

The following documents are the parents from which this document's scope and content derive:

162-WP-001-002	Operations Concept for Integration and Test of Science Data Production Software, White Paper for the ECS Project
423-10-01-00	Goddard Space Flight Center, Earth Science Data and Information System (ESDIS) Project -- Level 2 Requirements, Volume 0, February 18, 1993
423-10-01-01	Goddard Space Flight Center, Earth Science Data and Information System (ESDIS) Project -- Level 2 Requirements, Volume 1, January 27, 1993
423-10-01-05	Goddard Space Flight Center, EOSDIS Core System Statement of Work, May 21, 1993 (as updated by CCR 95-0053)
423-41-02	Goddard Space Flight Center, Functional and Performance Requirements Specification for the Earth Observing System Data and Information System (EOSDIS) Core System

2.2 Applicable Documents

The following documents are referenced herein and are directly applicable to this document. In the event of conflict between any of these documents and this White Paper, these documents listed here-in shall take precedence.

209-CD-008-004	ICD Between ECS and the GSFC Distributed Active Archive Center
305-CD-020-002	Release B SDPS/CSMS Design Overview Specification for the ECS Project
210-TP-002-006	Technical Baseline for the ECS Project

2.3 Information Documents

The following documents, although not directly applicable, amplify or clarify the information presented in this document.

102-CD-002-001	Maintenance and Operations Configuration Management Plan, Vol. 2 for the ECS Project

160-TP-004-001	User Pull Analysis Notebook for the ECS Project
161-TP-001-001	EOSDIS Product Use Survey, Technical Paper for the ECS Project
194-00311TPW	ECS Scenario Notebook for the ECS Project
194-00312TPW	User Characterization and Requirements Analysis for the ECS Project
194-00313TPW	ECS User Characterization Methodology and Results for the ECS Project
194-00548TPW	User Scenario Functional Analysis for the ECS Project

3. The ECS and DAS

3.1 EOSDIS Core System

Figure 3.1-1 shows a block diagram of a generic DAAC architecture. The objectives of the DAAC include archiving of the science data as well as providing support services to its various users. With the Release B implementation, the ECS becomes the largest component of the DAAC. ECS establishes a processing environment that facilitates archiving and distribution of data and products. In addition, ECS provides planning and scheduling services and ensures that the data is staged before the Product Generation Executables (PGE) are initiated. ECS de-stages the products of the PGEs, deletes unwanted files from the staging area, and reports the PGEs' processing status for human interpretation and control. This processing, archiving, and communications environment established by ECS, relieves the Science Processors from having to manage their own environments. The PGEs related to DAS run on Science Processors located at two of the DAACs, namely GSFC DAAC and EDC DAAC. Even though it is too early to speculate DAS equipment layout, majority of the DAS related processing equipment is expected to be located at the GSFC DAAC.

3.2 Data Assimilation System

DAS, a set of assimilation algorithms developed and managed by the DAO, is a substantial user of the ECS resources installed at the GSFC DAAC. In addition to utilizing resources of the GSFC DAAC, the DAS has a large collection of dedicated processors and storage equipment to accommodate its work load. PGEs related to the DAS run on heterogeneous platforms. Logically the DAS processing work load is grouped into Operations, Re-analysis, and Research & Development (R&D). The DAO has estimated that 20% of the processing power will be used by Operations, 40% by Re-analysis, and the remaining 40% by the R&D. The Operations and the Re-analysis work of DAO is similar to the first time processing and reprocessing done by other science algorithms. The sharing of the Science Processors with the R&D effort is unique for DAO. Also, the concept of a single PGE running on multiple computers with each computer having multiple CPUs, is unique for the DAS.

3.2.1 DAS Product Generation

The DAS combines observations from many different instruments using a General Circulation Model (GCM) to produce gridded data sets identified as Level 4 data products. Currently the PGEs associated with the assimilation run on Parallel Vector Processors (PVP) located at GSFC. The proposed 1998 system will add land surface and simple chemistry models to the existing GCM. Also, the future analysis will consider the influence of the data globally. To accommodate these growing processing needs, the DAO estimates that its processing power has to increase from about 0.5 GFLOPS to over 50 GFLOPS by 1998. Results from the ongoing prototyping at DAO and the simulation at HAIS will be used to perform a life cycle cost analysis to determine the optimum processing architecture as well as the platforms that will best fit the DAO's anticipated equipment needs. The initial conclusion from the analysis shows that the PVPs will not meet DAO's projected 1998 processing needs. Also, to bring down the cost per sustained GFLOPS, DAO is evaluating several platforms that are commercially viable.

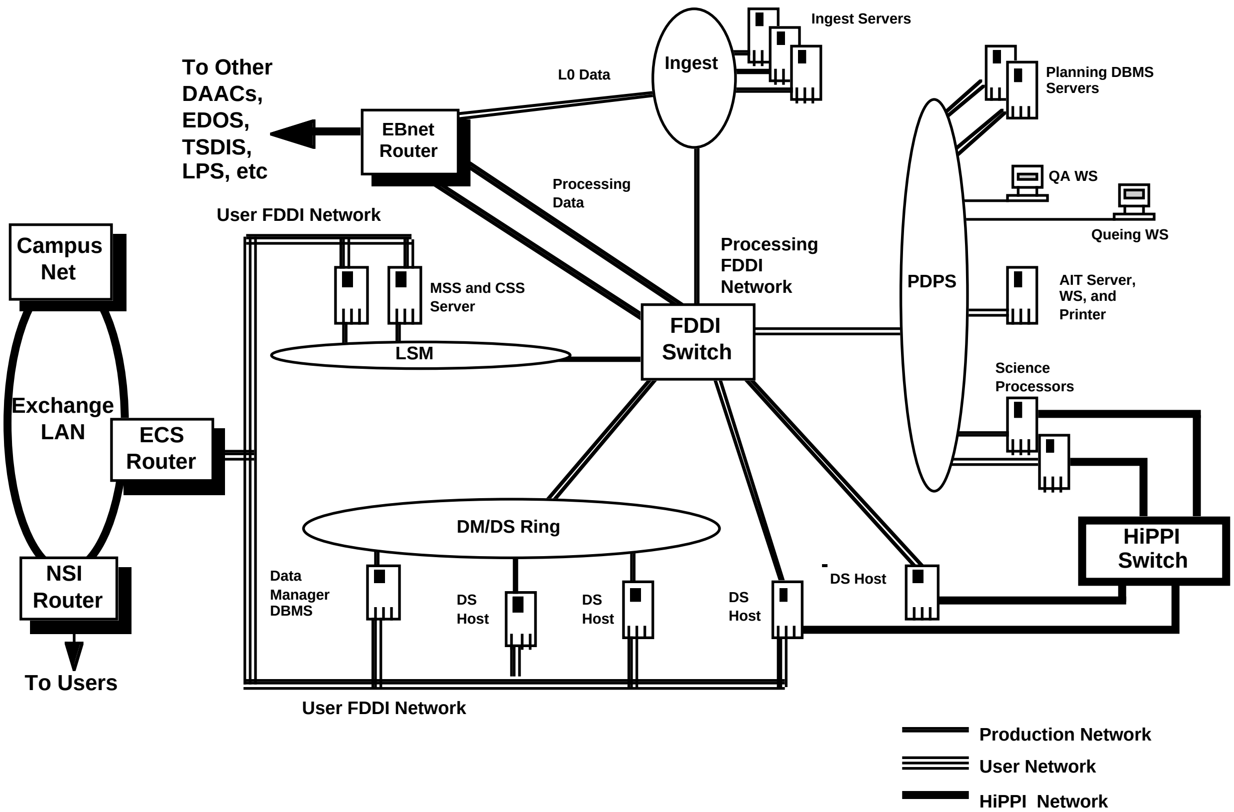

Figure 3.1-1. Release B Generic DAAC Architecture

3.2.2 DAS Performance Goals

The DAO requires that the final configuration selected for the DAS should make optimum use of the resources, especially the processing power. Table 3.2-1 shows the current estimate of DAO's resource needs. Given the unpredictability of future funding levels, DAO plans to procure the most amount of processing power for its allocated budget. The architectural design of the DAS environment will be geared to take maximum advantage of this processing power by minimizing data latency and by parallel processing of the PGEs.

Table 3.1-1. DAO Computing Requirements

Year	GFLOPS (Sustained)	Storage in TB
95	0.5	13
96	0.5	20
97	35	25
98	50	43
99	75	70
00	150	110
01	225	150
02	500	200

3.3 DAS Interface with ECS

There are two underlying themes in the definition of the interface between DAS and ECS. First, the scope of the interface is bounded to the extend that the DAO retains enough flexibility in selecting the architecture that best suits the DAS. Second, the interface between ECS and DAS will be treated as the standard ECS interface with other Science Processors. This latter consideration will result in the nailing down of the ECS - DAS interface so that when the DAS environment is completely defined, no redesign of this interface will be needed. Only the scalability aspect of the ECS - DAS interface will need to be scrutinized for integrating the DAS with ECS.

3.3.1 Release B and DAS

The Release B Version of ECS provides all capabilities needed to integrate DAS with ECS. The Planning Subsystem of ECS provides a planning capability and manages the resource utilization. The ECS Scheduler is responsible for ensuring the staging of the input data and the scheduling of the PGEs when all the relevant run conditions are met for each PGE. Upon completion of the execution, the Scheduler forwards the outputs from the PGE to appropriate destinations, and performs clean up activity on the working storage area used by the PGE. The processing status is reported by the Scheduler to the Management Subsystem (MSS) for appropriate action. In case of PGEs that use Same Program Multiple Nodes (SPMN) paradigm, the Scheduler submits the PGE to a front-end processor that is part of the resources allocated to that “Super” PGE. This allows the flexibility of using a detailed scheduler for the SPMN cluster.

The importing of the data from different data locations is coordinated by the Data Processing Subsystem of the ECS. This data get loaded into the staging area of the DAO processors before the PGEs are scheduled for execution. Abnormal conditions, when reported by the DAS, are passed to the Management function of the ECS for proper action. The distribution and archiving of the DAS products is carried out by ECS. The MSS will also use a Simple Network Management Protocol (SNMP) interface to monitor the health of the processors that make up the DAS environment. Thus the ECS environment allows the DAS to carry out the assimilation function without being burdened with the logistics of the operation.

3.3.2 DAS Inputs & Outputs

Table 3.3-1 shows the input and output data rate that need to be supported in order to fulfill the DAS's processing needs. Table 3.3-1 also shows the projected time frame during which each of the DAS processes is to begin running under the DAS - ECS integrated environment. The data rates that need to be supported for the Off-line Reanalysis and the R&D functions are not available at this point. In coming up with the estimates, "Super-Obbing", which reduces the volume of data to be sent to the DAS, has not been considered. Refer to Appendix A for more details about the processes that constitute the DAS.

Table 3.3-1. Data Rates for DAS

Process Name	Scheduled Period for Initial Run	Input Volume in GB/Day	Output Volume in GB/Day
First Look Analysis	1st Quarter '98	0.5	4.0
Ten Day Forecast	1st Quarter '98	0.001	2.0
Final Platform Analysis	3rd Quarter '99	416.0	3.6
20 Year Reanalysis	1st Quarter '00	740.0	108
Pocket Analysis	1st Quarter '00	414.0	7.2
Off-line Reanalysis	TBS	TBS	TBS
Research & Development	TBS	TBS	TBS

3.3.3 DAS Use of ECS Services

Table 3.3-2 lists the ECS services that will be used by DAS. The DAS environment may consist of heterogeneous platforms. Not all ECS services will be available on all of the DAS platforms. For example, the only ECS service the super computing platforms will use is the SNMP Agent to report the equipment status. Most equipment vendors support the COTS products identified in Table 3.3-2. When comparing the cost/performance of equipment from different vendors, the development/porting cost of ECS services and COTS products will be used as one of the evaluation criteria. Appendix B provides a more detailed account of the ECS services.

Table 3.3-2 ECS Services Used by DAS

Service	ECS Subsystem	Description	COTS Products / Vendors
Schedule	Data Processing	Scheduling and status collection of the PGEs	Auto Sys/Platinum
Production Plan	Planning	Generates plans for DAS production and equipment use	None (Part of ECS)
Equipment Status	Management	Provides equipment status of the super computing platforms	SNMP Agent/Peer Auto Xpert/Platinum
Communications Tool	Management	Integrates DAS component status with that of ECS	Sentry/Tivoli HP Openview/HP
File Transfer	Communication	Provides staging of data and distribution of products	FTP over FDDI FTP over HiPPI
Distributed Comp. Environment	N/A	Provides compatibility with the ECS environment	DCE Equipment vendor
DCE Application Development Tool	N/A	Facilitates application development for DCE	Equipment vendor

This page intentionally left blank.

4. Architecture

4.1 DAAC Architecture Overview

The objective of the DAAC is to archive science data and provide support services to its users. With the Release B implementation of ECS, the GSFC and EDC DAACs will provide interfaces between ECS and DAS, enabling DAS to take advantage of the ECS infrastructure. Though technically not an EOS mission/instrument, the ECS baseline for Release B considers DAS as a supported EOS mission.

4.1.1 GSFC DAAC

The Release B GSFC DAAC Design Specification for the ECS Project (Document Number 305-CD-030-002) describes the architecture of the GSFC DAAC. Figure 4.1-1 shows the LAN topology of the GSFC DAAC. Currently identified as the GSFC V0 DAAC, the GSFC DAAC operates multiple systems, with ECS representing the largest component of the DAAC. The DAS processing environment at the GSFC DAAC will have the capability to perform Operational Assimilation, Re-Analysis, and Research and Development (R&D) activities. Research and Development, in addition to facilitating the integration and test of algorithms, will provide a Science Computing Facility (SCF) for development and validation of the DAS Algorithms.

To accommodate these DAS activities, the GSFC DAAC will maintain a dedicated equipment suite for DAS. Categorized as production processes, operational assimilation and re-analysis will be performed using this equipment suite. R&D will share this equipment suite with the production processes and ECS will be responsible for ensuring that these different modes of operation do not interfere with each other. The input data will be made available to the processes in the DAS suite in a manner which optimizes the equipment usage. DAS will interface with ECS for the archiving and distribution of the products generated by DAS.

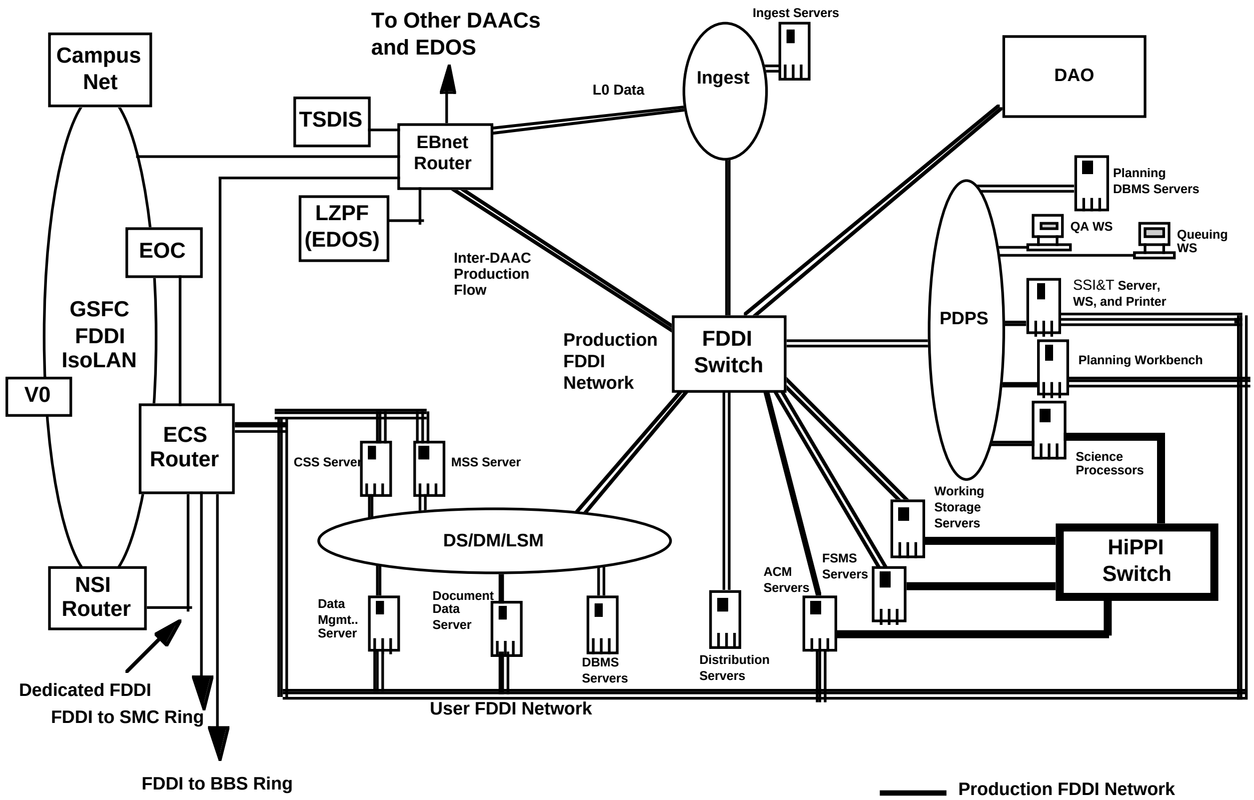

Figure 4.1-1. GSFC DAAC LAN Topology

4.1.2 EDC DAAC

The Release B EDC DAAC Design Specification for the ECS Project (Document Number 305-CD-033-002) describes the architecture of the EDC DAAC. Figure 4.1-2 shows the LAN topology of the EDC DAAC. DAS will require MODIS data processed and archived at EDC. Because of the substantial volume of this data, a form of data compression referred to as "Super-Obbing" will be used to reduce the data volume. Super-Obbing is based on the fact that the resolution of the data collected by the MODIS instruments is much finer than what DAS will use. Therefore, the data can be reduced substantially before transmission to the GSFC DAAC and still not degrade the quality of the assimilation process. MODIS instruments collect data at a resolution of one kilometer by one kilometer. The DAS assimilation algorithms use grids spaced at 1.25 degrees * 1.25 degrees. At the equator this translates to about 135 kilometers by 135 kilometers. At the time of MODIS product generation, these Super-Obbing algorithms fromDAO will be run on the EDC science processing hardware. This Super-Obbed data will then be forwarded to GSFC DAAC for archiving and eventual use by DAS.

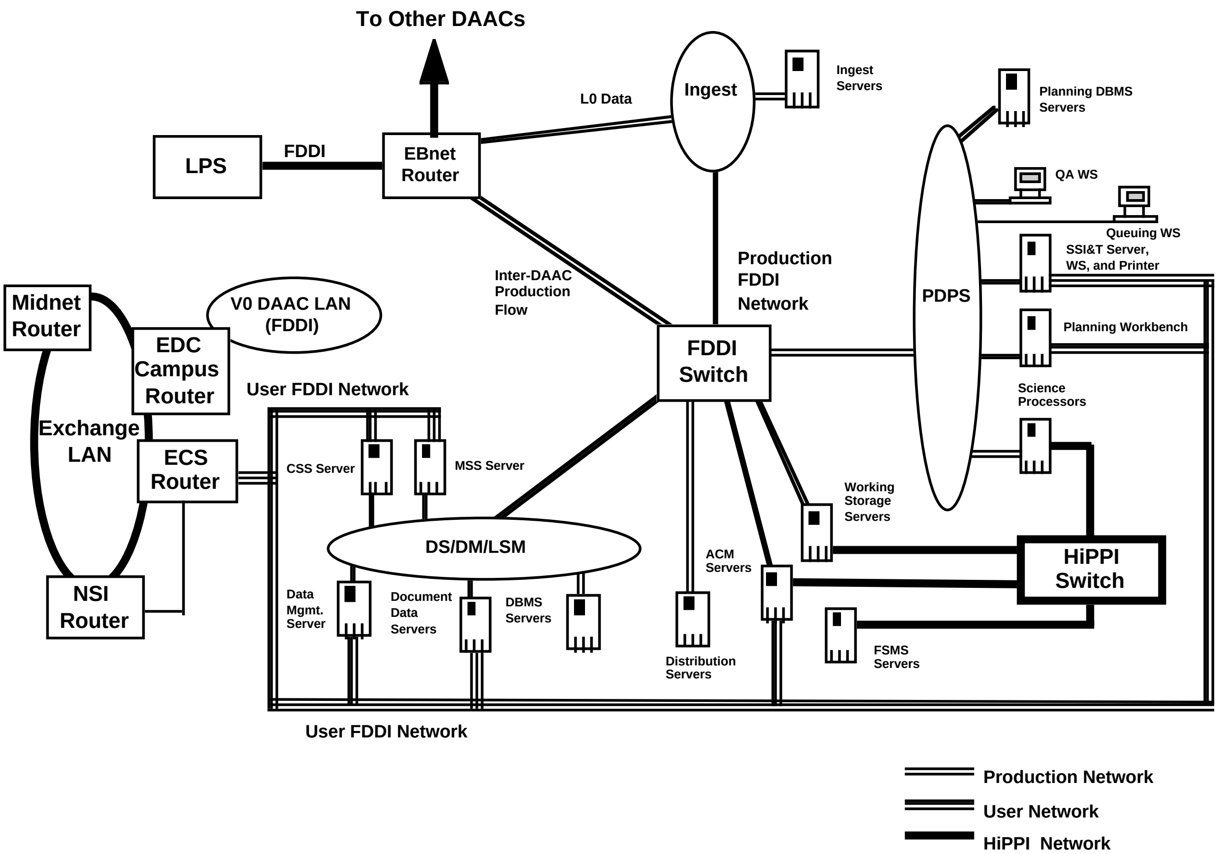

Figure 4.1-2. EDC DAAC LAN Topology

4.1.3 LAN Configuration

The LAN topology under Release B ECS consists of a User FDDI Network, a Production FDDI Network, and at the EDC, LaRC, GSFC DAACs a HiPPI Network. The creation of separate User and Processing networks allows processing flows to be unaffected by user pull demands, and the introduction of the high-speed HiPPI Network provides adequate bandwidth to the Processing and Data Server subsystems to transfer high volumes of data. The block diagram in figure 4.1-1 shows DAS being connected to ECS through the Production FDDI Network. Refer to Appendix B.11 for a detailed discussion of the rationale for this approach. Any future need for the higher bandwidth HiPPI connection for DAS will not require any redesign of the ECS infrastructure.

4.2 DAS Interface with SDPS/CSMS Subsystems

The following SDPS and CSMS subsystems will interface with DAS and will have a bearing on the performance of the PGEs that comprise DAS: the Data Processing Subsystem (DPS), the Planning Subsystem (PLS), the Data Server Subsystem (DSS), the Management Subsystem (MSS), and the Communications Subsystem (CSS). The paragraphs below describe the contributions made by each of these subsystems to DAS processing.

4.2.1 Data Processing Subsystem

The basic functions of the Data Processing Subsystem are to:

- Stage data for PGEs in the job queue and initiate their execution;
- Perform predictive staging to reduce processing latency; and
- Monitor all the ongoing processing for abnormal behavior.

To support the parallel DAS computational needs, no change to the basic design of PDPS will be needed. However, additional functions will be implemented in the Toolkit to support parallel computation in general. Before a PGE that performs parallel computing is initiated, the required CPU resources (perhaps on multiple SMPs) and the staging area needed by the PGE are all allocated. The necessary communication facilities between the archiving facility of the DAAC and the staging area are already provided as a part of the infrastructure of the ECS.

To meet the R&D requirements of the DAO, multiple sessions of interactive operations will be supported. This includes the capability to perform editing, compiling, and debugging.

4.2.2 Planning Subsystem

The basic functions of the Planning Subsystem are to :

- Make tentative schedules for routine processing based on the projected availability of processing resources, inputs for PGEs for a time period, and production rules;
- Release a PGE to the job queue when all required inputs become available locally;
- Schedule On-Demand production requests based on their priorities and the availability of resources reserved for this purpose; and
- Re-plan the tentative schedules as needed.

As such its basic design will not change as a result of bringing DAS into the operational envelope of ECS.

4.2.3 Data Server Subsystem

As a result of having DAS use the ECS archive and retrieval facility, the Working Storage Hardware (WKSHW) and the Deep Repository Hardware (DPRHW) Configuration Items of the GSFC DSS will be expanded. The ECS design for the GSFC DAAC uses storage and retrieval capacities estimated for the third quarter of 1999, including the added capacity needed for supporting DAS data and products. Appendix B.5 documents results from trade studies done to optimize DSS service to DAS. It concludes that the current physical configuration, in which DAS shares the DSS equipment, is the optimum use of the storage hardware budget.

The Super-Obbing algorithms that are planned to be run on the science processing equipment at the EDC DAAC forward their outputs to the GSFC DAAC. The Super-Obbing by the Science processors is not expected to add significantly to the DSS load at EDC.

4.2.4 Management Subsystem

Appendix B.14 documents the results from a trade study conducted to evaluate management schemes for the DAS equipment. It concludes that since the DAS environment consists of heterogeneous platforms, two different approaches will be used. The DAS supercomputing platforms will report their status using the SNMP agent provided by the supercomputer vendor. The rest of the DAS equipment will use the standard ECS management scheme for reporting its status.

4.2.5 Communications Subsystem

Appendix B.11 describes the analysis for the communications interface between ECS and DAS. Based on this analysis, an FDDI interface between ECS and DAS at the GSFC DAAC provides adequate bandwidth for the estimated DAS data rates. The data rates projected for fourth quarter '99 were used for this analysis. If data rates substantially increase after that period, the ECS design does facilitates the use of HiPPI connections to selected DAS hosts. The communications interface at EDC DAAC is not significantly impacted by the Super-Obbing processes.

4.3 DAS Architecture

Figure 4.3-1 shows a block diagram of the DAS environment. Since the exact configuration of the equipment in the DAS environment has not been finalized, the interface descriptions in thissection between DAS and ECS are to be considered as possible options. These options provide sufficient flexibility in the selection of COTS products for DAS while ensuring adequate infrastructure support to maximize DAS equipment throughput. At the same time, the specification of these interfaces prevents any ripple effect on the ECS design due to equipment selections for DAS.

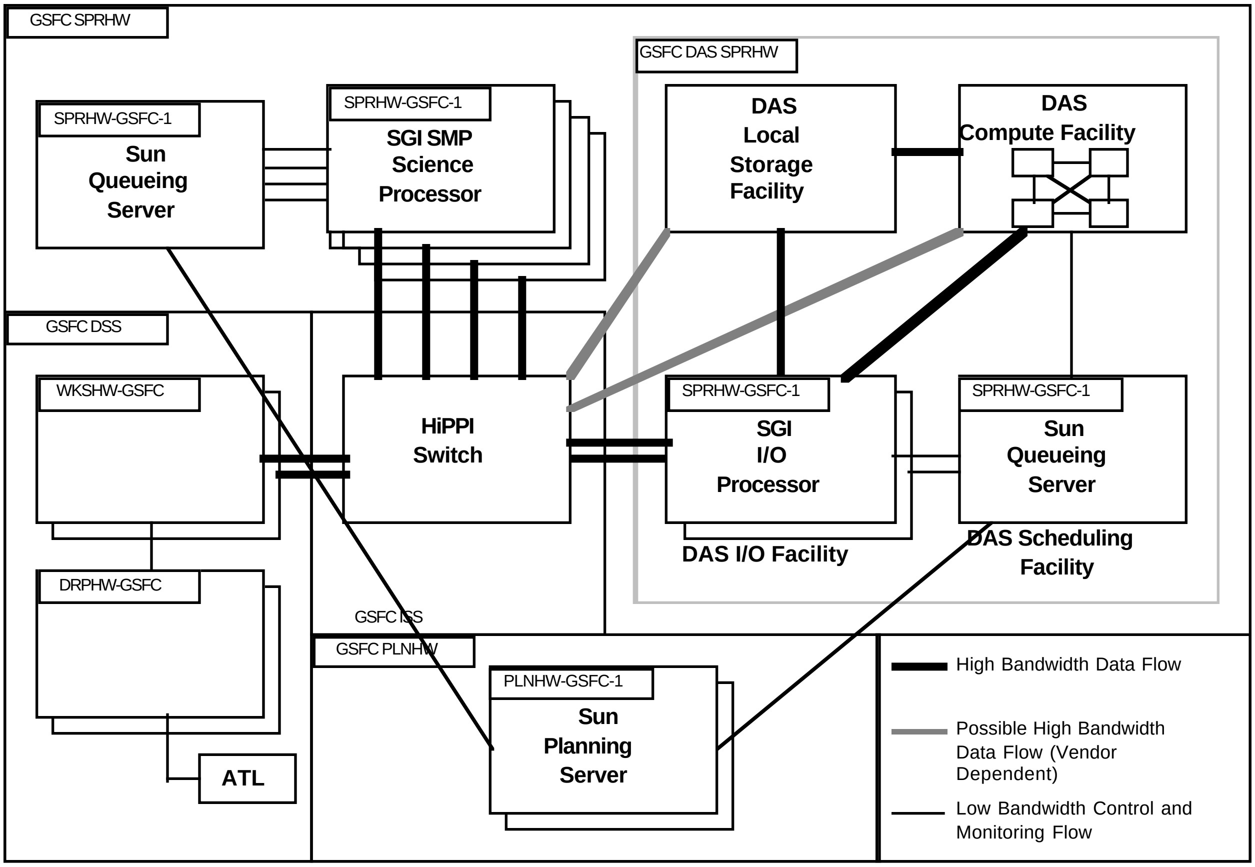

Figure 4.3-1. Data Assimilation System Environment

4.3.1 DAS Science Processing Hardware

The heart of the Science Processing hardware is the Compute facility. The estimate for the 1998 time frame is for processing power of 50 GFLOPS. The current architecture used by DAO, Parallel Vector Processor (PVP), is found to be inadequate and economically not feasible. DAO is leaning toward a Message Passing paradigm implemented on clusters of Symmetric Multi-processing (SMP) systems. Processing power estimates beyond '98 are somewhat speculative; the procurement task will essentially concentrate on obtaining the most processing power for the allocated budget. The design of the DAS environment will then concentrate on getting the maximum throughput by providing sufficient I/O bandwidth and storage capacity.

Unlike other ECS Instrument Teams, DAS has PGEs identified as SPMNs (Single Process, Multiple Nodes) running on multiple SMPs. These SMPs will use the vendor provided Message Passing Interface (MPI) for communicating with each other; ECS does not provide MPI support. The data distribution and the process monitoring for the SPMNs will be coordinated by the PGE internally with no ECS-level support.

4.3.2 Planning/Scheduling of PGEs

The GSFC ECS Planning/Scheduling functions will also support DAS product generation. Appendix B.7 has a detailed discussion of how the Planning and Data Processing subsystems of ECS manages and synchronizes the PGEs that constitute the DAS. These subsystems together, creates plans for the PGEs, stages and destages data involved in the running of the PGEs, and collects the status of the execution for further processing and display. A Sun Scheduling Server is proposed for the DAS Scheduling Facility. ECS uses Autosys, a COTS product, for the scheduling function.

4.3.3 DAS I/O Facility

One or more SGI I/O processors are proposed for the staging and de-staging of the data needed by the DAS operations. The SGI Challenge platform was selected by ECS for other Science Processors because of its high bandwidth backplane and its scalability. The DAS scheduler will initiate the PGEs for the I/O processor and the I/O processor in turn will manage the transfer of data between DSS and the DAS Local Storage. It is anticipated that any data segmentation, gridding, and post-processing prior to delivery to ECS will take place in this facility.

4.3.4 DAS Local Storage

Data transfer between the supercomputing platforms of DAS and DSS will use the facility identified in figure 4.3-1 as the DAS Local Storage Facility. This Local Storage will be used as the central repository for all of the input data and products from DAS operations. Any distributions to individual nodes in an SPMN environment will be handled by front-end processors that are part of the SPMN cluster. Both network attached storage and host attached storage are candidate technologies for the local storage. Cost, flexibility, and scalability will be used as criteria for selecting the optimum storage.

This page intentionally left blank.

5. DAS Schedule

5.1 Deployment Plan Overview

This section contains the deployment plan for all DAS equipment that will be procured through the end of 1999. In developing this deployment plan, the launch of MTPE's AM-1 satellite was used as a critical event. The AM-1 launch is scheduled for June of '98 and it is intended that the first installment of the DAS equipment be operational by this time frame. Given that this operational deadline can be met, this deployment plan then, focuses on ways to delay the selection of DAS related COTS products in order to take advantage of advances in technology and reduction in unit cost thereby, providing great flexibility to DAO for obtaining the best suited and most cost effective equipment for the data assimilation task. After the initial installation of equipment, this deployment plan envisions two incremental upgrades of equipment for DAS. In order to highlight the critical events during the selection of equipment for the initial installation, the deployment schedule is being presented in two sections. Section 5.2, Pre Procurement Schedule, describes the events leading up to the procurement of the initial installment of equipment. The Section 5.3, DAS Deployment, shows the overall schedule for the procurement of DAS equipment including the two upgrades planned for ’98 and ‘99.

5.2 Pre Procurement Schedule

The Pre Procurement Schedule describes activities from the beginning of April '96 to the end of June '97. Figure 5.2-1 shows the tasks that will be worked as well as the milestones that will be met during this period. One of the milestones scheduled for this period has already been completed with the briefing conducted on April 11, ’96 at the ECS development facility. This White Paper itself was issued as a work-in-progress document prior to the briefing and will have its first version published by the end of May '96. As shown in Figure 5.2-1, this White Paper will remain as a "living document" during the Pre Procurement phase with the final version to be published by end of June '97. A Pre Procurement review is scheduled for the end of May '97 with the intent of having DAO issue the Bill of Material (BOM) for the first installment of equipment by the end of June '97.

5.2.1 DAS Prototyping

The on-going prototyping of the DAS is the process which will identify the most suitable computing architecture as well as platform(s) for the DAS. Performance and operational data from prototyping and benchmarking, are being used in creating models for the DAS equipment. Simulation runs are used for extrapolating the prototyping/benchmarking results to estimated future loads. These simulation results are providing valuable information about infrastructure impacts resulting from different DAS architectures. Also, the simulation results and other pertinent data will be used to perform a cost/performance analysis to generate estimates for lifecycle cost and software development/porting cost.

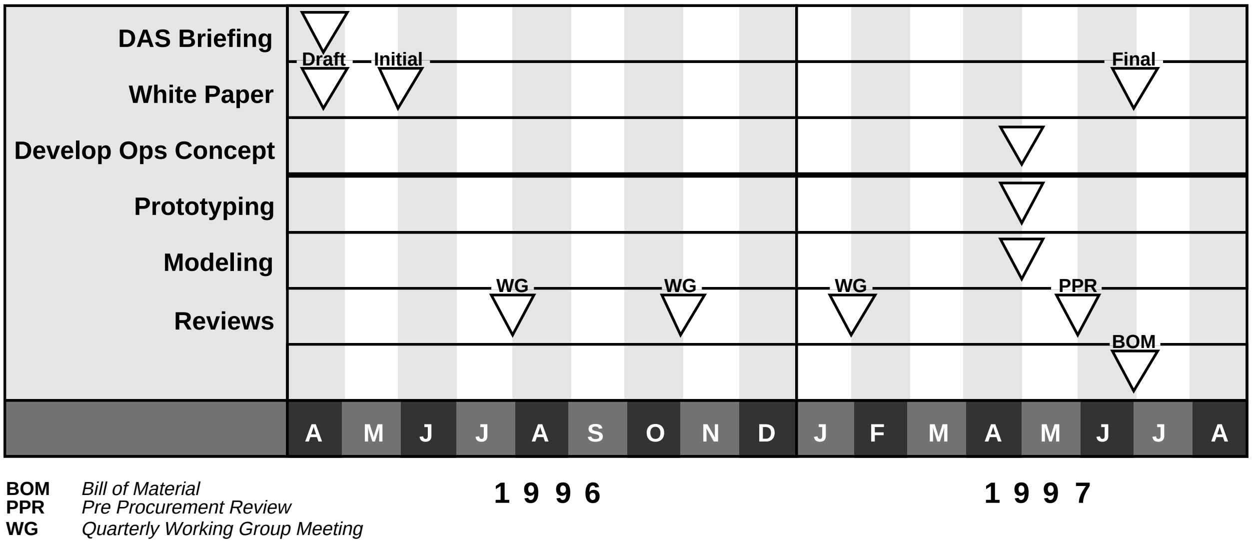

Figure 5.2-1. Pre Procurement Schedule

5.2.2 DAS Equipment Selection

The DAO and ECS working jointly will provide a detailed BOM for the equipment to be procured. As mentioned earlier, this BOM will be issued by the end of June '97. The DAO and ECS will also have obtained any necessary configuration change approval so that ECS can immediately proceed with purchase of the equipment. It is very clear that this schedule is very ambitious and therefore, any departure from anticipated equipment configuration will adversely affect the ability to field the system in time for the AM-1 launch. In order to reduce the risk of equipment selection surprises, quarterly working group meetings between DAO, ESDIS, and ECS will be conducted. The following three paragraphs describe additional steps that are being taken to mitigate the risk of selecting incompatible equipment for DAS.

5.2.2.1 Vendor Qualification

Until now the ECS program has put five (5) vendors, namely Silicon Graphics Inc. (SGI), Hewlett Packard (HP), International Business Machines (IBM), Digital Equipment Corp. (DEC), and Sun Microsystems, under contract to supply COTS equipment including software for the ECS program. Current prototype effort at DAO has identified CRAY as another potential vendor. Any change in the vendor list lengthens the procurement lead time due to the need for competitive bidding. It takes approximately nine (9) months for the procurement process to get a vendor under contract. This means that any new Request For Proposal (RFP) effort has to be started, at the latest, third quarter of ‘96 so that the potential vendor(s) will be under contract by the time the DAS equipment selection is finalized. Again, this procurement time-line assumes no developmental effort on the procured system.

5.2.2.2 OS/DCE Version Compatibility

The five vendors listed above have available the Open System Forum's (OSF) Distributed Computing Environment (DCE). ECS program, which adopted the Object Oriented Distributed Computing Environment (OODCE), has put HP under contract to implement the enhancements needed for getting the products from the above mentioned vendors, compatible with the OODCE. New vendors will be made to undergo a qualification process to ensure that their equipment and software are compatible with that of ECS.

5.2.2.3 ECS Utilities Porting

Table 5.2-1 identifies the COTS software products that are to be available on the DAS equipment. Not all utilities need to run on all equipment (e.g. the SMPs don’t need to support most of the products). Since the DAS consists of equipment from more than one vendor, a case by case evaluation will be made as to the COTS product supported by each of the equipment. Every effort will be made to get the equipment vendor support these COTS software products. Incompatibility of COTS products and ECS environment will impact the procurement schedule as well as the life-cycle cost of the DAS equipment.

Table 5.2-1. COTS Products for DAS

Product	Description
Auto Sys	Process Scheduler
Auto Xpert	Process Monitoring
Fortran 90	Compiler
Fortran 77	Compiler
C	Compiler
SNMP Agent	Resource Monitor

5.3 DAS Deployment

The overall deployment schedule for the DAS is depicted in Figure 5.3-1. The figure itself is logically divided into two halves. The bottom half of the figure shows key milestones and events during the deployment time period. The top half shows deployment and operations activities, and indicates approximate hardware proportions to support them. The horizontal axis of the figure shows time in months and the vertical axis on the top half shows GFLOPS of processing power available for the DAS. For example, the initial deployment consists of 50 GFLOPS of computing to be followed by 25 GFLOPS of incremental upgrades in each of the two following years.

5.3.1 Deployment Assumptions

In coming up with this schedule, the following assumptions were made:

1. The DAS will be up and running as early as possible relative to the launch of the EOS AM-1 satellite. This means that the DAS will be integrated into EOSDIS production at the GSFC DAAC.
2. The initial DAS deployment is a separate mini-release, subsequently referred to, in this paper, as Release B.05
3. Release B.05 requires minimal ECS functionality over Release B. Specifically, nearly all DAS required ECS functionality will be provided in Release B.
4. Release B will provide anticipated infrastructure to support DAS. As an example, robotics and read/write stations required for DAS staging and product distribution will be provided in Release B.

5.3.2 Overall Deployment Schedule

Release B.05 will basically be restricted to providing the computing environment, and, possibly, dedicated mass storage for the DAS in 1998. There will be subsequent incremental upgrades to the DAS in each of the succeeding years. These will be associated with as yet unnamed releases.

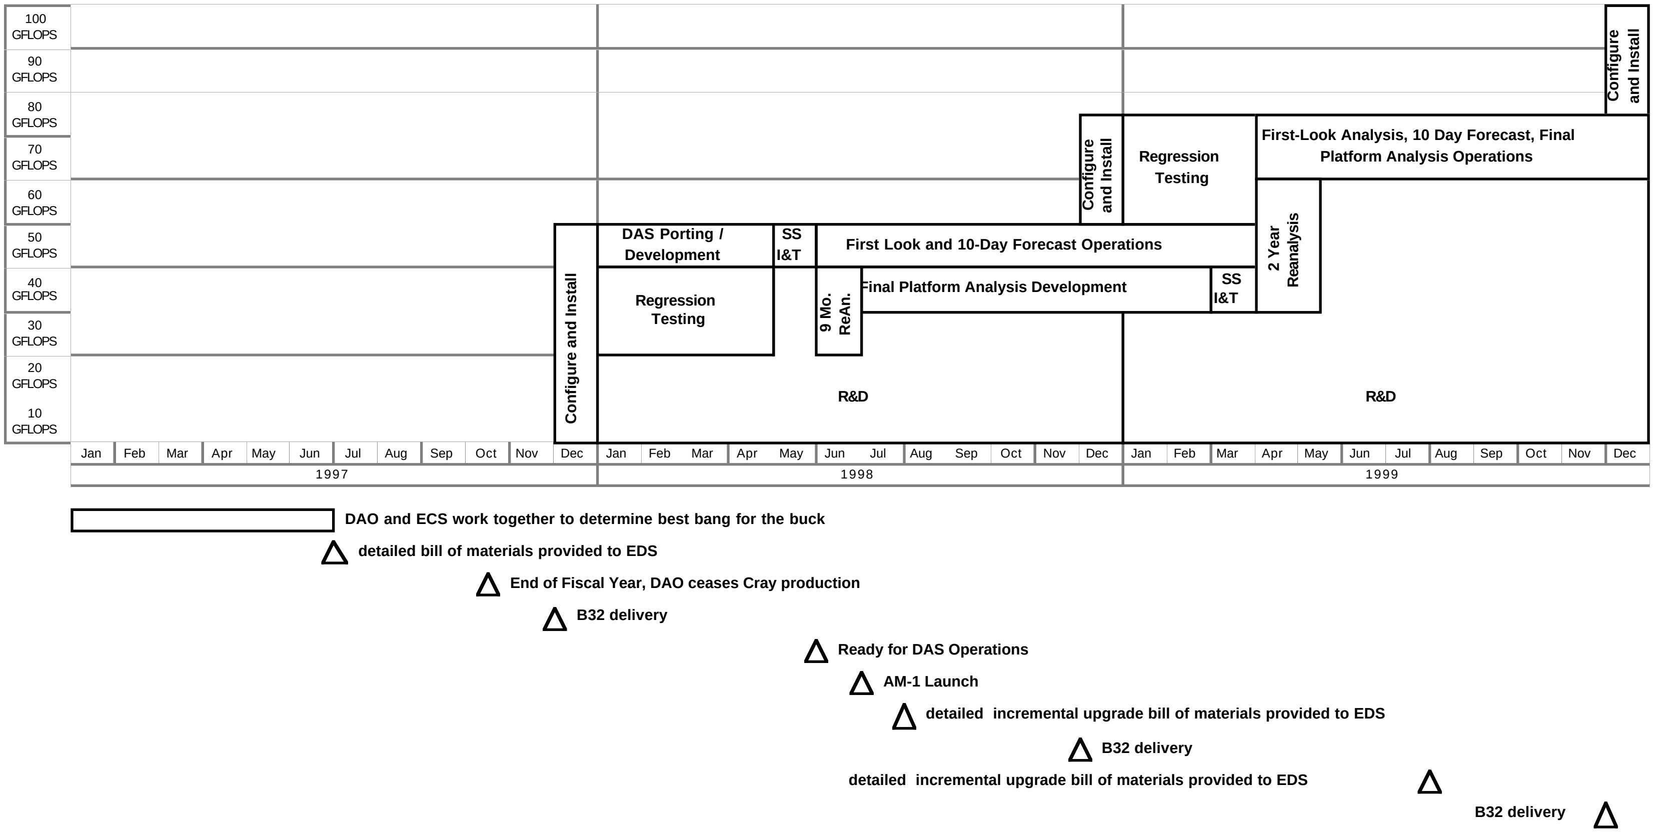

Figure 5.3-1. Overall Deployment Schedule for DAS

5.3.3 DAS Regression Test

The scope of testing will have a great deal of dependency on the vendor equipment. Release B testing schedule shows four (4) separate test related activities namely Integration & Test (I&T), Acceptance Testing, Independent Verification & Validation (IV&V), and ECS Ground Segment (EGS) operations. If DAS equipment is selected from one of the approved vendors, only a limited amount of these tests need to be conducted. A recommendation is being made to ship the equipment directly to the site saving considerable time. Certain degree of unpredictability exists when new vendor products get integrated into the ECS environment. If DAS platform selection results in new type of hardware, an assessment will be made about the feasibility of getting the hardware delivered directly to the DAAC.

5.4 Scheduled Activities

Described below are the DAS related procurement and installation activities based on an yearly basis going up to the year 2000.

5.4.1 1997 Activities

The kickoff event for Release B.05 deployment is the approved, detailed BOM being provided to EDS at the end of June 1997. This will involve minimal additional ECS infrastructure components, since ECS components needed for DAS will already have been procured and delivered with Release B. Five months are allotted until the delivery of the material directly to the GSFC DAAC (Building 32) facility. The last month of 1997 is relegated to installing the DAS equipment and configuring it to run in the extant DAAC. This includes installing ECS software. The goal is to perform the installation and configuration without interfering with ongoing DAAC operations at GSFC, or at any other DAAC for that matter.

5.4.2 1998 Activities

During 1998, three parallel activities, namely regression testing, DAS porting, and R&D, will be conducted on the DAS equipment. The Regression testing is expected to last about four months. This regression testing is intended to include all manner of contractor and government tests needed to insure that the DAS equipment is both, properly installed and capable of functioning with the GDAAC and the rest of EOSDIS.

The First-Look Analysis is the first candidate for porting to the new DAS environment. The manner in which these three activities share the DAS environment is negotiable. The following are some of the possible options: partitioning, in which some of the resources are used for testing, and other resources are used for porting; time division, in which testing gets all the equipment for part of the day, and porting gets it at other times. In any case, resources not used for testing and DAS porting, are available for DAO’s R&D effort.

At the end of the development period, which coincides with the end of regression testing, one month is allocated for SSI&T. Its purpose is to validate that the First-Look Algorithm is properly integrated with ECS.

Following SSI&T, the First-Look and Ten-day Forecast algorithms begin production operations: ingesting NMC observations; and producing DAS outputs for distribution through GDAAC. A brief period is also allocated for processing observations that may have been accumulating since the DAO ceased production on the NCCS system currently being used. However, the DAO feels that this may not be necessary, since other facilities are being considered for meeting this need.

In addition to production, the balance of the year is spent on porting of the Final Platform Analysis algorithm. This algorithm uses AM-1 instrument retrieval products, but, at this point in time, they are not expected to have sufficient quality to support assimilation production. However, they can be used for development support.

December is the time for installing the first DAS equipment increment, estimated to be in the order of 25 GFLOPS. This increment also includes infrastructure anticipated for DAO's Final Platform Analysis algorithm. Such infrastructure may include robotics, read / write stations, and communications. Since some of the inputs to Final Platform Analysis are archived at sites other than GSFC, notably EDC, any infrastructure impact is likely to affect multiple DAACs and inter-DAAC communications. The goal is to perform the installation and configuration with minimal interference to on-going operations. The equipment will be ordered four to five months prior to installment, but achieving the minimal interference goal requires careful planning that must begin long before that.

5.4.3 1999 Activities

1999 begins with parallel activities of testing, operations, development, and research activities. Regression testing is performed to assure that the new 25 GFLOPS and infrastructure are operational and integrated with the balance of EOSDIS. This testing may involve multiple DAACs, since some of the infrastructure may affect sites beyond GSFC.

Final Platform Analysis algorithm development will culminate when the quality of the instrument retrieval products is deemed adequate. This could be a single episode. Alternatively, individual instruments could be phased in incrementally over a period. That is a science issue to be resolved. When ready, SSI&T is performed to assure proper operation of the Final Platform Analysis algorithm with ECS. Final Platform Analysis is then added to the DAS production suite which already includes First-Look Analysis, and 10-Day Forecasting.

Reanalysis is allocated a couple of months to assimilate the instrument data since launch. This must be coordinated with reprocessing of the instrument data over that same period.

Another equipment increment is slated for the end of the year. As in the previous year, the increment must be ordered several months earlier, and careful planning is needed for what will likely affect multiple DAACs.

5.4.4 Year 2000 and Beyond

The 20'th Century's last year, not shown on the schedule, is expected to look much like 1999. The new increment must be tested. Production continues. By the end of 2000, however, the computer line reaches its obsolescence point. DAO anticipates reevaluating the situation and re≠competing the processing environment. This could mean a wholesale replacement of the solution adopted in 1997. DAO expects this to occur on a regular 3-year cycle.

6. DAS Operations Concept

6.0 Introduction

As mentioned earlier in this document, the Data Assimilation environment has three distinct processing modes: Operations, Re-analysis, and R&D. Once integrated with the ECS infrastructure, various PGEs of DAS running in different modes will use the following operations concepts.

6.1 Operations

Operations component of DAS provide gridded global analyses and forecasts on a daily basis for various users within the ECS. Also fly-through data sets will be produced for the EOS platforms. The following paragraphs detail end to end scenarios for DAS Operations.

The Operations mode of DAS will have three processing components. The first component is the First-look analysis that runs 12 to 24 hours after data acquisition, and uses primarily non-EOS inputs. This run provides background information to the AM-1 instrument retrieval algorithms. The second component is the Ten-day forecast which runs once a day following the first-look analysis. It is used to support NASA field experiments and system validation by providing estimates. The third component is the Final Platform analysis which runs several months after data acquisition and uses data from the AM-1 instruments in addition to the data used by the First-look analysis. Final platform analyses may include standard DAS assimilation and off-line analyses.

6.1.1 First-Look Analysis

First-look analysis is the 'real time operational' processing of the DAS and uses the NMC observational data as input. First-look analysis runs as a single PGE and performs a complete 24≠hour assimilation. The 24-hour NMC observational data that is used by the First-look analysis, is generated in six (6) hour increments at 0Z, 6Z, 12Z and 18Z.

The output of the First-look analysis will include global gridded data as well as specially tailored products for the AM-1 instruments. The tailored products will include special "fly-through" data sets matching the satellite sub-track. Appendix A lists the inputs and outputs related to the First≠look analysis.

Table 6.1-1 lists, in chronological order, the activities that take place during the First-look analysis. The initiation of the PGE is based on the availability of all the inputs. In the event some of the data needed by the First-look analysis is not available, the PGE is initiated based on a timer event (12 to 24 hours after data time). When either all the data have been staged to the DAS staging area or the initiation time limit has expired, the ECS Scheduler will submit the

First-look analysis PGE for execution on the processors allocated for this purpose. The first-look analysis, however, will only utilize about 7% of the total DAS processing resources each day. (Analysis based on Operations = 20% of day, assuming First-look analysis and Ten-day forecast will use 1/3rd of the processing resources and Final-platform analysis will use the remaining 2/3rd of the resources) . The First Look Analysis processing load is thus very predictable.

The Planning Subsystem (PLS) of ECS will determine when First-look analysis is to run and block out resources for it to use in that time window. PDPS will ensure that adequate resources (discs and processors) are set aside for the First-look analysis mode.

6.1.2 Ten Day Forecast

The Ten-day forecast job is spawned automatically by the first-look analysis. It is merely a ten≠ day extension of the assimilation and does not use observational input. The 1 MB/day of input to the Ten-day forecast is from boundary-conditions such as the sea surface temperature climatology. The output from this PGE, on the other hand, is much larger than the assimilation, approximately 20 GB for a Ten-day forecast. Most of the output from this PGE is saved with provision to retain up to ten Ten-day forecasts (roughly 200 GB worth of storage). Only about 4 GB data from each Ten-day forecast will be permanently archived at the GSFC DAAC.

6.1.3 Final Platform Analysis

The Final Platform Analysis is driven exclusively by the availability of all the input data and is initiated only after all needed AM-1 data are available. The initiation will likely be several months after the data observation time. Nominally, there is one Final Analysis PGE per day that assimilates a 24-hour period.

The input data for this PGE will be considerably more than the first look because of the added AM-1 data. There is still uncertainty as to whether the Final-platform analysis will use the raw data from the AM-1 radiance measurement or the retrieved level 2 product for that instrument. If direct radiance are used, the amount of input data will be significantly more than if the level 2 products are used. In addition, off-line analysis techniques may be employed to assimilate some new data types. Since the time constraints on this product are less than the First-look, the Final≠ platform analysis will run at a lower priority but still, will run at priorities higher than that for Re-analysis and R&D. The final platform analysis will utilize about 13% of the total DAS processing resources.

Table 6.1-1. First-Look Analysis Scenario (1 of 2)

From	To	Action
NMC	GSFC DAAC (Larry)	0630Z - NOAA Radiosondes for -0600Z to 0000Z arrive
GSFC DAAC (Larry)	Ingest Subsystem	0635Z - NOAA Radiosondes ingested
Ingest Subsystem	Data Server Subsystem	0640Z - NOAA Radiosondes archived
Data Server Subsystem	PDPS - DAS Scheduling Facility	0643Z - Scheduler notified of NOAA Radiosondes existence. Scheduler determines that it is input to DAS data preparation PGE. Scheduler initiates Radiosondes staging.
Data Server Subsystem	PDPS - DAS Local Storage Facility	0644Z - NOAA Radiosondes data copied to Operations Partition
PDPS - DAS Scheduling Facility	PDPS - DAS I/O Facility	0650Z - On conclusion of data transfer, Scheduler determines that it can run DAS data preparation PGE. Scheduler allocates an Operations Partition processor and starts job. It is assumed that the appropriate binary executable image resides as a permanent file on either the local storage facility or processor attached disc in the I/O facility.
PDPS - DAS Local Storage Facility	PDPS - DAS I/O Facility	0651Z - DAS Data Preparation PGE reads input Radiosondes, sorts / re-samples to the DAS grid, and writes the output to the Local Storage Facility operations partition. Data preparation PGEs will be included in the AHWGP baseline for the DAS. These functions are currently performed on Larry/ Hera for the NCCS supported DAS.
PDPS - DAS I/O Facility	PDPS - DAS Scheduling Facility	0653Z - Job ends with normal completion signal.
PDPS - DAS Scheduling Facility	PDPS - DAS Local Storage Facility	0653Z - Scheduler determines that input Radiosondes are not needed for any pending or near term expected jobs and deletes file. Scheduler determines that re-gridded radiosondes are needed for First Look analysis and leaves them. Scheduler determines that, however, additional inputs are also needed that do not exist, so no further action is taken at this time.
. . .	. . .	Over the next 18 or so hours, more NMC data arrives, is ingested, is archived, is data prepared, but there is still not sufficient input to run First Look Analysis
PDPS - DAS I/O Facility	PDPS - DAS Scheduling Facility	2455Z - Data preparation job that vends last input needed to run First Look Analysis ends with normal completion signal.
PDPS - DAS Scheduling Facility	PDPS - DAS Local Storage Facility	2455Z - Scheduler determines that last input NMC file is not needed for any pending or near term expected job and deletes it. Scheduler determines that re-gridded output file is needed for First Look Analysis and leaves it.

Table 6.1-1. First-Look Analysis Scenario (2 of 2)

From	To	Action
PDPS - DAS Scheduling Facility	PDPS - DAS Compute Facility	2456Z - Scheduler determines that all necessary input for First Look Analysis exists on Local Storage Facility. Scheduler allocates n processors on Compute Facility Operations Partition and starts SPMD (Single Program Multiple Data) job. It is assumed that the appropriate binary executable image resides as a permanent file on either the local storage facility or processor attached disc in the compute facility..
PDPS - DAS Local Storage Facility	PDPS - DAS Compute Facility	2470Z - DAS First Lookup Analysis SPMD job runs, reading its inputs from and writing its outputs to local storage. This is the First Look Assimilation for observation period 0000Z-2400Z relative. In this rendition it is assumed that the SPMD job has a master process, or other mechanism, for parallelizing its inputs and serializing its outputs. Alternatively, these functions could be carried out on the I/O facility before and after the First Look Analysis job. Also, the assumption that the job uses the local storage facility, rather than directly attached storage in the compute facility, is made for purposes of simplification in this presentation. A more complicated situation might apply, depending on the final design for the DAS compute environment.
PDPS - DAS Compute Facility	PDPS - DAS Scheduling Facility	2590Z - First Look Analysis Job ends with normal completion signal.
PDPS - DAS Scheduling Facility	PDPS - DAS Local Storage Facility	2590Z - Scheduler deletes input files. Scheduler does not delete any interim files, such as restart file, needed for next First Lookup Analysis execution, or output files. Scheduler initiates transfers of output files for archival.
PDPS - DAS Local Storage Facility	Data Server Subsystem	2591Z - Appropriate output files copied for permanent archival and distribution to users. DAO has suggested that DAS products should be stored in multiple forms to satisfy various user needs. It would be appropriate to use the DAS I/O processor to support this activity. Further study is needed to develop a concept for managing these multiple renditions.
PDPS - DAS Scheduling Facility	PDPS - DAS Local Storage Facility	2591Z+ - As archival file transfers complete, scheduler deletes files not needed for near term expected usage.

6.2 Re-analysis

Re-analysis is the process of running a long-term assimilation using historical data and a "frozen" assimilation system to produce a consistent and useful data set for scientific research. Typically, the DAO will run a complete 20-year Re-analysis about once every 4 years. This will be done to take advantage of advances in the assimilation system including the use of previously unanalyzed data. Other types of Re-analysis will include Platform Re-analysis, Off-line Re≠ analysis, and Pocket Re-analysis.

Table 6.2-1 lists the activities that take place during the running of the Re-analysis PGEs. These PGEs will be scheduled with some notice to the DAAC Operations staff. It is assumed that the input data will all be staged to DAS local storage. These PGEs will have lower priority than the operational jobs but will run with a higher priority than the R&D work.

The Re-analysis software will be configured so that 30 days of data can be assimilated in 24 hours. The Re-analysis will utilize 40% of the available DAS processing resources each day, when it is running. A 20 year Re-analysis will take about 243 days. It is assumed that the remaining time allocated for this purpose is used up with other shorter Re-analysis.

6.3 Research and Development

R&D work will also be run on the DAS machines. This is work that usually would take place at an SCF. This work is the lowest priority work and will not be allowed to interfere with Operations and Re-analysis. The proposed mechanism to accommodate R&D is a special case of the SSI&T Mode. The ECS system has the capability to logically partition the resources into different modes. The 'R&D Mode' for the DAS will have a limited of resources and be prevented from interfering with the other two modes.

R&D jobs will be run by the DAO scientists. These jobs will routinely access the DAAC data sets and the local mass storage system. Some research jobs will also create data that will be archived at the DAAC. These will not generally be made available to the scientific community, however. It estimated that the DAO researchers will access an average of 90 GB/day from the DAAC.

These jobs include: typical software development activities like compilation, linking, and debugging (though this principally is conducted off-line on workstations); check-out assimilation with varying temporal coverage and data staging requirements; and multiyear special assimilation and analysis experiments. There is great variety in the characteristics of individual R&D jobs, but the aggregate load is expected to equal to that of the Re-analysis. Typically the R&D jobs are run interactively.

Table 6.2-1 Reanalysis Scenario

From	To	Action
Reanalysis is a planned activity. Every 24 hours, 30 days are assimilated. 30 1-day assimilation jobs are run in a sequential manner in order to level the loads on data server and communications infrastructure. As this scenario begins, job n-1 has just finished, and job n has just started.		
PDPS - DAS Scheduling Facility	DAS Local Storage Facility	Scheduler deletes job n-1 input files not needed for job n. Scheduler does not delete job n-1 interim files or output files. Scheduler initiates transfers of job n-1 output files for archival.
PDPS - DAS Local Storage Facility	Data Server Subsystem	Appropriate output files copied for permanent archival and distribution to users. DAO has suggested emphatically that DAS products should be stored in multiple forms to satisfy various user needs. It is the responsibility of ECS to manage these multiple renditions. It would be appropriate to use the DAS I/O processor to support this concept.
PDPS - DAS Scheduling Facility	PDPS - DAS Local Storage Facility	As archival file transfers complete, scheduler deletes files not needed for near term expected usage.
PDPS - DAS Scheduling Facility	Data Server Subsystem	Scheduler initiates staging for job n+1 inputs
Data Server Subsystem	PDPS - DAS Local Storage Facility	As data are staged they are copied to the Reanalysis Partition
PDPS - DAS Scheduling Facility	PDPS - DAS I/O Facility	As a data transfer concludes, Scheduler checks to see if it needs to run a DAS data preparation PGE against it. If so, Scheduler allocates a Reanalysis Partition processor and starts job. If a processor is not available, Scheduler queues job and starts it when a processor becomes available. It is assumed that the appropriate binary executable image resides as a permanent file on either the local storage facility or processor attached disc in the I/O facility
PDPS - DAS Local Storage Facility	PDPS - DAS I/O Facility	DAS Data Preparation PGE, if run, reads input, sorts / re-samples to the DAS grid, and writes the output to the Local Storage Facility Reanalysis Partition.
PDPS - DAS I/O Facility	PDPS - DAS Scheduling Facility	Data Preparation job ends with normal completion signal.
PDPS - DAS Scheduling Facility	PDPS - DAS Local Storage Facility	Upon receiving completion signal from job, scheduler cleans up Reanalysis partition by deleting files not needed for anticipated jobs, and keeping those that are.
. . .	. . .	This process is repeated until all files are delivered and all associated data preparation activities are completed. When that happens, job n+1 can be run only when job n finishes and the restart file needed for job n+1 is available.
PDPS - DAS Scheduling Facility	PDPS - DAS Compute Facility	Upon receiving successful completion signal from job n, Scheduler frees associated Reanalysis Partition processors. Scheduler then observes that all necessary inputs for job n+1 exist on Local Storage Facility. Scheduler allocates Reanalysis Partition processors and starts SPMD job n+1. It is assumed that the appropriate binary executable image resides as a permanent file on either the local storage facility or processor attached disc in the compute facility.
As the scenario ends job n is running on the DAS Compute Facility. Things start all over again from Step 1 for de-staging job n outputs, staging and preparing job n+2 inputs.		

7. Science S/W Integration and Test

7.0 Integration and Test Goals

The principal goals of the Integration and Test (I&T) of Science S/W (SSI&T) are:

- To accomplish and verify the integration of the Science S/W with ECS components (e.g., configuration management, planning, processing, data server)
- Test the ability of the Science S/W to run to normal completion repeatedly over the normal range of data inputs and run-time conditions, and
- Verify that the Science S/W executes without interfering with other S/W or DAAC operations.

Other goals of SSI&T include

- Verifying the production resource requirements for the Science S/W such as CPU time, RAM, data access and temporary storage.
- Testing of Science S/W interfaces external to the DAAC such as the input of ancillary data to the Science S/W.

Ordinarily, the SSI&T process is necessary because software is developed at other facilities (e.g., SCF) using different computing hardware, libraries and different operating systems. The goals of SSI&T remain the same for the DAS as for other Science S/W. For the DAS, however, the Science S/W is developed using the same computing hardware used in the production, employing the same versions of operating system and libraries.

7.1 Anomalous Operations and Fault Handling

The SDP Toolkit provides mechanisms to support error handling and logging of error conditions. In most anomalous situations, however, the strategies to permit restart or graceful failure must be incorporated into the Science S/W. Through appropriate use of exit codes, for example, the ECS Processing subsystem can take subsequent actions as designated by the developer of the Science Software.

7.2 DAS Development

The development of new DAS software or enhancements to existing software typically involves stand-alone testing, a period of parallel runs with the output being compared with that of a current operational version, and promotion to operational status if and when improved performance has been demonstrated. Both the stand-alone and the parallel "operations" are included in the 40% of DAS resources devoted to R&D.

7.3 Science Software Integration and Test

7.3.1 Introduction

Unlike all other Science S/W, the DAS S/W will be developed using the specific hardware within the GSFC DAAC which will also support DAS production. There will be no need to transfer the software via an external network or media.

The degree of integration of DAS S/W with ECS increases as the Science S/W is tested under conditions which approach those of the production S/W. In early stages of development, appropriate use of mandatory SDP Toolkit calls and configuration management is sufficient. Later, during parallel operations, the full integration with ECS is necessary, including placing the appropriate information in the planning database and data server.

7.3.2 Operations

Prior to the launch of EOS-AM1, the production Science S/W will need to undergo SSI&T more than once, with the final pre-launch SSI&T effort involving the final DAS Science S/W to support First-look analysis and Ten-day forecast. The production version of the Final-platform analysis S/W must await the validation of the specific EOS instrument data products that will be used, and this will not be completed until several months following launch.

To reduce the volume of EOS instrument data that must be transferred from various production DAACs for use in Final-platform analysis, DAS-provided software may be run at the specific DAACs to form "super-observations" (Super-Obbing) of high resolution data through sub-setting or averaging, for example. This Super-Obbing software needs to be delivered to the appropriate DAAC(s) for SSI&T at the DAAC, with deliveries and SSI&T activities needing to be defined between the DAO and the affected DAAC.

Major upgrades of the production code are expected to be infrequent, perhaps twice a year. Tuning of the DAS software and minor code changes will likely be ongoing, however. Changes in the DAS production S/W and associated documentation need to be carefully managed as part of SSI&T, and a procedures document needs to be developed between the DAO and the GSFC DAAC delineating the respective responsibilities and steps involved. While not a responsibility of SSI&T, it is also vital that procedures be established to alert the users of the DAS products to any and all changes in DAS S/W, their reasons and consequences.

7.3.3 Re-analysis

Each Re-analysis effort involves running a long-term assimilation using historical data using a "frozen" assimilation system. The frozen system will likely be a current or previous production system, or some variant thereof. The level of SSI&T effort would depend on the extent of the changes to the code.

7.3.4 Research and Development

DAS Research S/W will not have gone through the SSI&T process, although the use of the mandatory SDP Toolkit routines is encouraged to simplify later development stages. Since the Research S/W will be running on the same system as the production S/W, strategies need to be employed which will isolate the production S/W from the effects of faults or failures of research S/W. The options might include partitioning of the DAS hardware in some fashion or encapsulation of the R&D work to prevent affecting production.

This page intentionally left blank.

8. Open Issues

8.1 Issues List

8.1.1 H/W Environment

In defining the Release B interfaces, equipment related to DAS was assumed to be resident at the GSFC DAAC. The only exception to this rule is any equipment that may be required at EDC DAAC for the task of Super-Obbing. An alternative approach being considered by DAO involves using the Super Computing facility at Ames for carrying out the computationally intensive tasks. Current modeling has not evaluated the impact to infrastructure from this approach.

This page intentionally left blank.

Appendix A. DAS Data Flows

A.1 Introduction

This Appendix examines the computing load and mass storage requirements for DAS and estimates the impact on the ECS infrastructure resulting from the DAS processing. Rigorous capacity planning methods are used for assessing the capability of the DAS environment to support current and future DAS processing needs without degrading user-specified performance parameters. Capacity planning enables prediction of system saturation point as a function of time and provides an insight into the most cost effective way of extending the useful life span of resources. These capacity planning related questions are being answered in a scientific manner using performance models and simulation of the whole environment. This provides an insight into the behavior of complex computer systems with various interdependencies. In Summary, the inputs for this capacity planning are the workload evolution, system parameters and desired service levels and the outputs are the saturation points and their cost-effective alternatives.

A.1.1 Simulation Methodology

The dynamic simulation model has been developed using the BONeS simulation tool. The model facilitates the study of DAS system performance and can be used to get answers for capacity planning questions. It also enables “What If “ sensitivity analysis using several possible scenarios of system configuration alternatives. The key objectives of DAS performance modeling are:

- Assess DAS Hardware/Software performance
- Study the impact of DAS system on other ECS DAACs
- Assist in DAS system sizing and capacity planning
- Identify critical resources and bottleneck areas
- Propose design alternatives to maximize system performance
- Provide timeline profile of resources
- Furnish detailed modeling results to the DAO
- Update/refine model with workload changes to perform trade studies

A.1.2 Simulation Parameters

The ECS dynamic simulation model for DAS uses the simulation parameters extracted from the ad hoc Working Group on Production (AHWGP) databases. The AHWGP database for DAS operations was updated in February of '96 based on the revised DAS estimates. Preliminary run results are furnished in this appendix. When this simulation run was made, no R&D requirements were specified, hence it was excluded from the current run. The R&D related requirements will be incorporated in future studies. Preliminary results indicate enormous staging and communication requirements at EDC and GSFC DAACs. No Super-Obbing was performed on the MODIS data set for the current simulation run.

For this run, it is assumed that GSFC DAAC will operate and maintain all DAS related equipment including the mass storage and computing resources. Also, the computing resources identified as DAS equipment will be used exclusively for DAS Operations, Re-analysis, and R&D work and will not be shared with other users. Table A.1-1 describes the estimated data input/output and computer load for DAS. In generating this estimate, it is assumed that all data needed by DAS, regardless where they are generated , are archived at the GSFC DAAC.

Table A.1-1. DAS Resource Requirements

Input in GB from GSFC DAAC	DAS Computer Load	Output in GB to GSFC DAAC
24 GB/Day	20 % Operations	28 GB/Day
8 GB/Day	40 % Reanalysis	50 GB/Day
100 GB/Day from GSFC DAAC + 10 GB/Day from DAS Local Storage	40 % Research & Development	50 GB/Day to GSFC DAAC + 10 GB/Day to DAO Local Storage
Operations Break down		
1 GB/Day from GSFC	(1) First Look Analysis	4 GB/Day to GSFC
1 MB/Day from GSFC	(2) 10-Day Forecast	20 GB/Day to GSFC
23 GB/Day from GSFC	(3) Final Platform Analysis	3.5 GB/Day to GSFC
Reanalysis Break Down		
8 GB/Day from GSFC (This estimate is the average data rate for all of the Re-analysis related PGEs)	(1) Platform Reanalysis (2) 20-Year Reanalysis (3) Off-line Reanalysis (4) Off-line Reanalysis	50 GB/Day to GSFC (This estimate is the average data rate for all of the Re-analysis related PGEs)
Research & Development		
100 GB/Day from GSFC + 10 GB/Day from Local Storage	R & D Load Includes any of the above Assimilation	50 GB/Day GSFC + 10 GB/Day to Local storage

A.2 DAS Related Issues

There are four issues currently being worked on in relation to the DAS integration.

1. Process Inputs and Outputs indicate enormous staging and communications requirements
2. Super-Obbing should be considered as a method to reduce communications requirements
3. Research and Development definition needs to be refined

Each issue is discussed in the paragraphs below.

A.2.1 Discussion of the Issues

DAS Reanalysis has several inputs for which the Daily Volume exceeds a tera-byte:

1. ∝MOD34_L3_10DY - MODIS Level 3 10-day Gridded Vegetation Indices (Max NDVI & Integrated MVI)
2. MOD15_L4_10DY - MODIS Level 4 10-day Leaf Area Index & FPAR
3. MOD16_L3_10DY - MOD16 Diurnal Surface Water Lookup Table
4. MOD09_BRDFL3_16DY - MODIS Level 2 Surface Reflectance and Vegetation Indices
5. MOD12_L3_32DY - MODIS Level 3 Monthly Land Cover Database

These are all synoptic MODIS gridded products. They are also indicated as inputs to Final Platform Analysis and Pocket Analysis. These input products are the reason that DAS has staging and communication requirements an order of magnitude above the total of everything else in EOSDIS. The fact that all of these products must be staged from EDC, where they are archived, to GSFC, where the DAS runs, implies excessive inter-center bandwidth.

The tera-input-product enormity inspires several questions:

1. ∝Are these synoptic products appropriate inputs to the DAS? They respectively represent 10 day, 10 day, 10 day, 16 day, and 32 day averages. This seems like a fairly long time scale relative to what the assimilation is doing.
2. ∝Are the products being used properly? The Process Descriptions indicate that for each 'n' days of product coverage, the DAS algorithms are reading 1/n'th of the data per day of assimilation (where 'n' depends on the file type specifically 10, or 16, or 32). Since each n-day product gives one measurement per grid point per n-days, rather than one measurement per grid point per day, this use appears invalid. A possible alternative, is that 1/n'th of the change from the last n-day period is assimilated per day, to avoid excessive assimilation loads. Also, the AHWGP Process Descriptions need to reflect the process.

3. ∝Is it correct to stage each n-day synoptic input once per assimilation day? These files should be staged so that an n-day file supports n days of assimilation, thereby substantially reducing average staging and communication requirements.

A.2.2 Super-Obbing

Super-Obbing could be used at EDC to reduce the communication requirement. The MODIS products have 1 KM^2 resolution. The DAS needs 1.25 deg^2 resolution. This has significant bandwidth reduction potential, perhaps as great as 1000 to 1, maybe more. If, additionally, super≠obbing is performed as MODIS produces the data, rather than as the DAS reads the data, this can also significantly reduce staging requirements.

A.2.3 Research and Development

Some characterization of DAS R&D activity is essential. Roughly 40% of the DAS computing resources are relegated to R&D according to Table A.1-1. Yet the estimates for staging, data storage, and communication related to R&D have low confidence level. ECS has suggested some sort of stochastic model, but this approach has not matured yet. DAO has been looking at ways to improve the quality of the R&D resource estimates. In order to get a handle on the overall infrastructure needs, a valid R&D model needs to be established soon.

A.2.4 Summary

The previous sections identified a number of issues that need to be resolved in order to generate high fidelity simulation estimates for the DAS resource needs. Expeditious resolution of these issues will ensure that ECS design provides adequate interface for DAS. Frequent meetings between DAO and ECS are being planned to resolve these issues.

A.3 Simulated DAS Environment

A.3.1 Resource Parameters

The PDPS and DSS resources and their characteristics used for DAS modeling are summarized below. This simulation covered a period of 21.

1. 40 Processing CPUs at GSFC, each with 50 GFLOPS sustained speed
2. A DSS with 20 Read/Write stations and 6 Robots
3. Tape I/O rate of 11 MB/Sec
4. Processing subsystem with unlimited local Disk storage
5. Max. Robot Delay of 13 seconds
6. Max Tape Mount delay of 17 seconds

7. Max Tape Rewind delay of 85.0 seconds
8. Max Search Delay of 95.0 seconds
9. Max Tape Dismount delay of 10.0 seconds

A.3.2 Input/Process/Output

The DAS Process Flow diagrams (illustrated in Figures A.3-1 through A.3-4) provide details about input/output granules such as granule size, their temporal coverage, and file disposition type (archive, temporary, interim, permanent etc.). Also, this flow diagrams describe PGEs frequency/day and their estimated MFLOPS (Millions of Floating Point Operations/Second) requirements per activation. AHWG database release of February '96 was used to derive these functional diagrams.A.3.3 Simulation Assumptions

The following modeling assumptions and model scenario were used for DAS system capacity planning and performance monitoring.

1. Feb./96 AHWG Technical Baseline used for DAS processing & archiving workload
2. No Super-Obbing was performed on MODIS data sets.
3. No Version 0 migration and user pull load at GSFC site was considered.
4. ∞Processing subsystem CPUs for DAS production processes were performed using 50 GFLOPS sustained speed for DAS processors.
5. DAS resources under study were not constrained for the initial modeling run.
6. ∞DAS PUSH load for year 2000 included operations and reanalysis workload but no R&D load profile was incorporated.
7. ∞Stand alone DAS model run at GSFC site (no load form other ECS instruments at GSFC DAAC or other ECS DAACs were integrated in this run)

Of course, the capacity planning and DAS system performance assessment through simulation and modeling using the dynamic model will be less precise without considering complete service levels and workload evolution. The results presented here are obtained using February '96 technical baseline as input to the model.

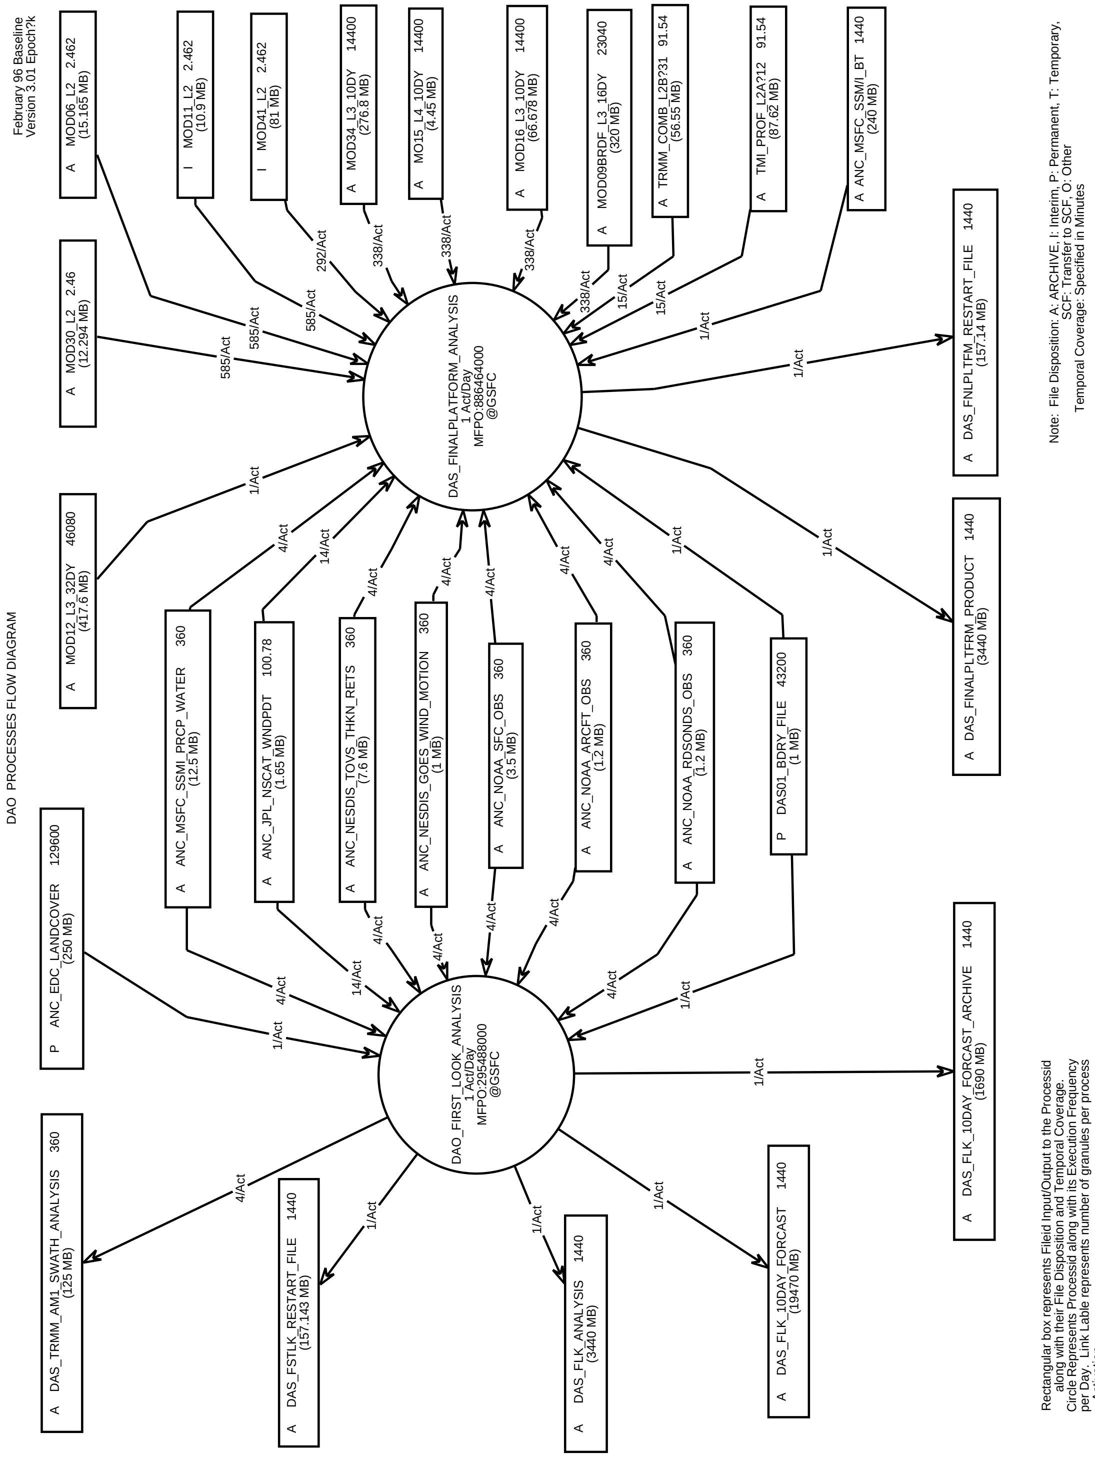

Figure A.3-1. First-look and Final-platform Analysis

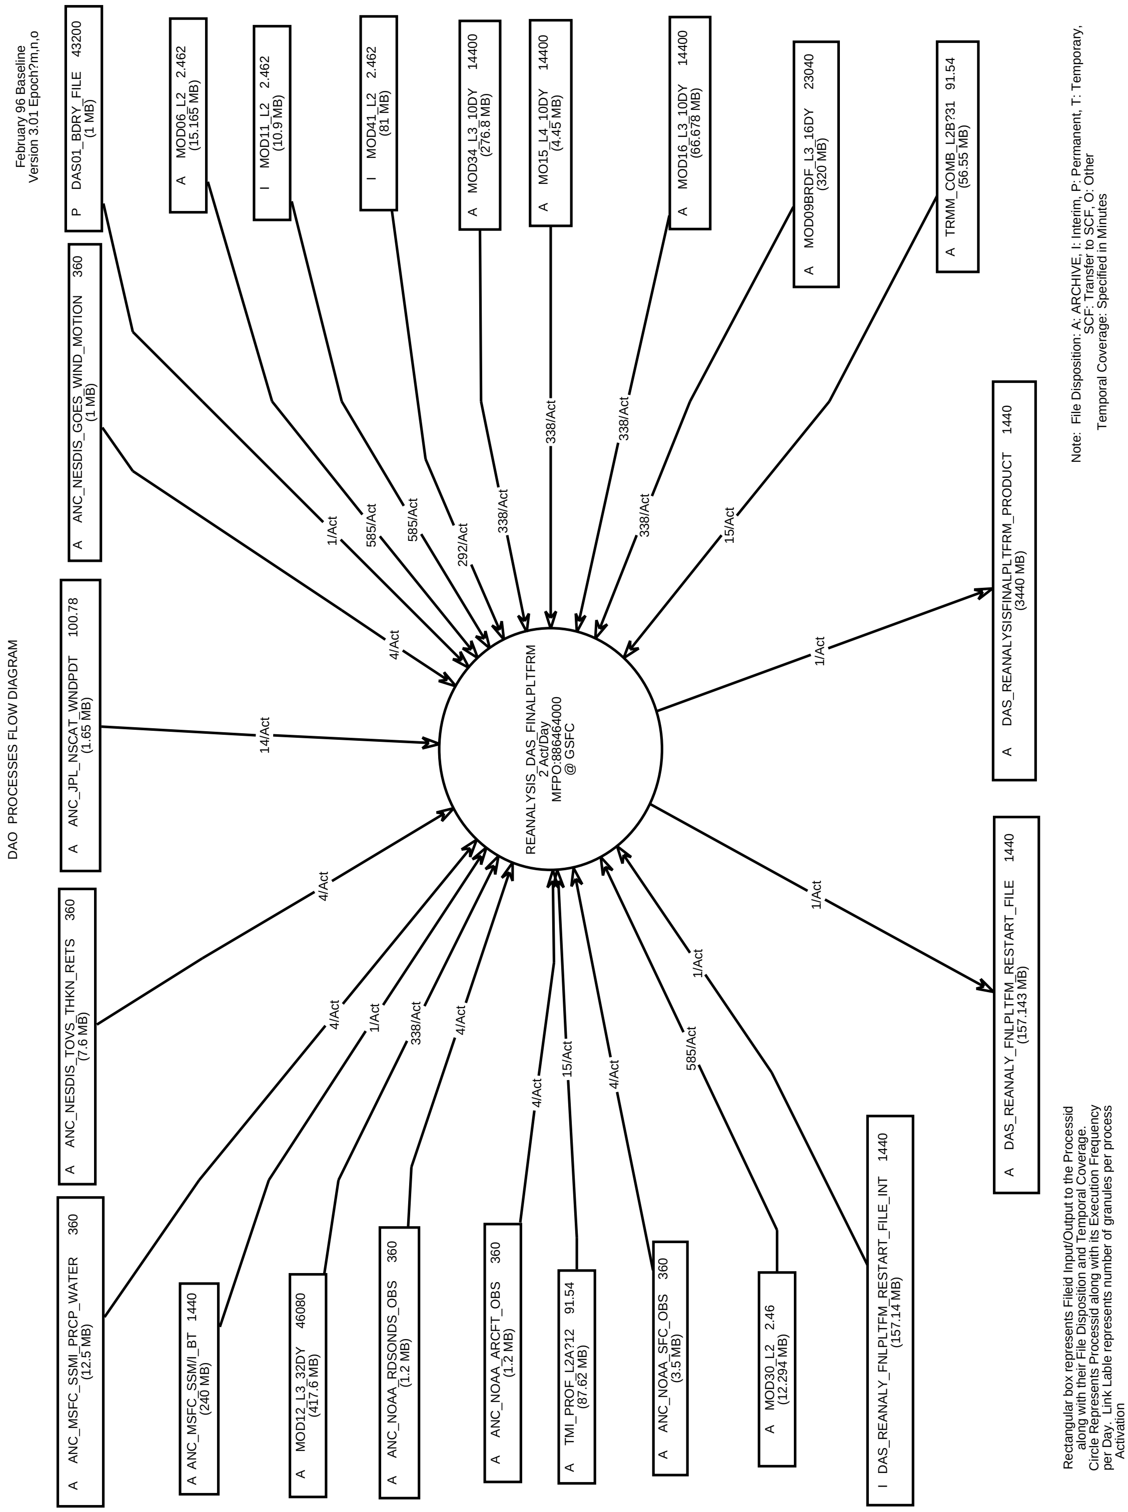

Figure A.3-2. Re-analysis

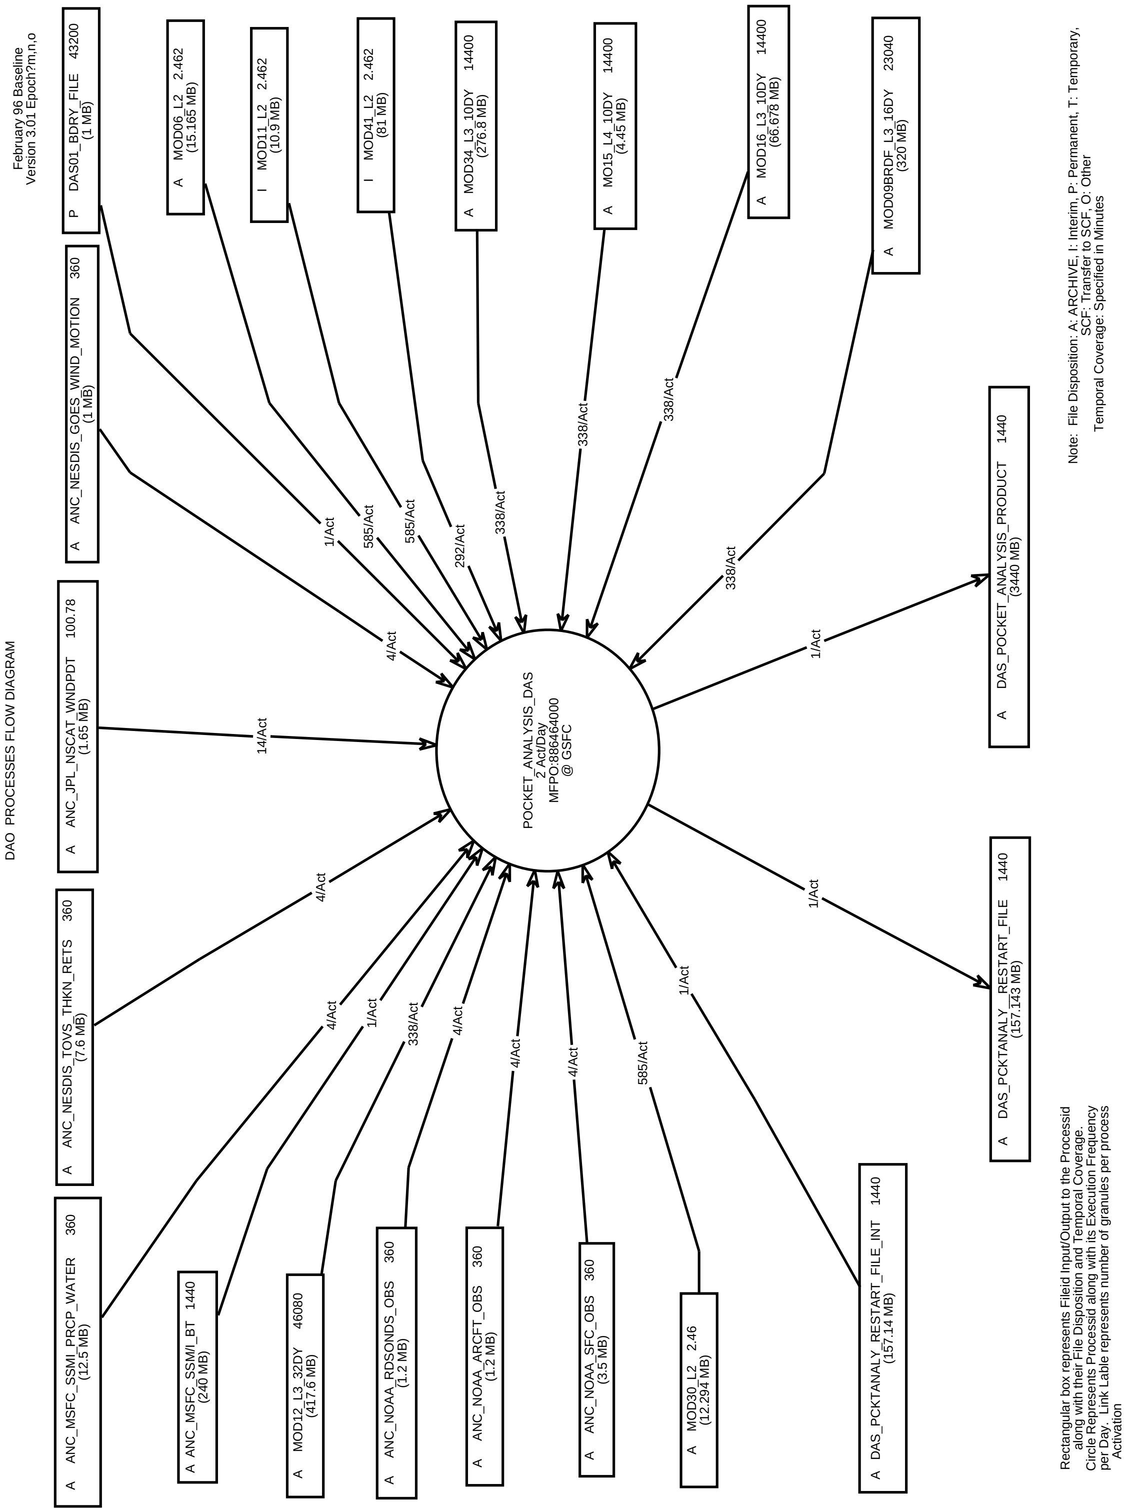

Figure A.3-3. Pocket Analysis

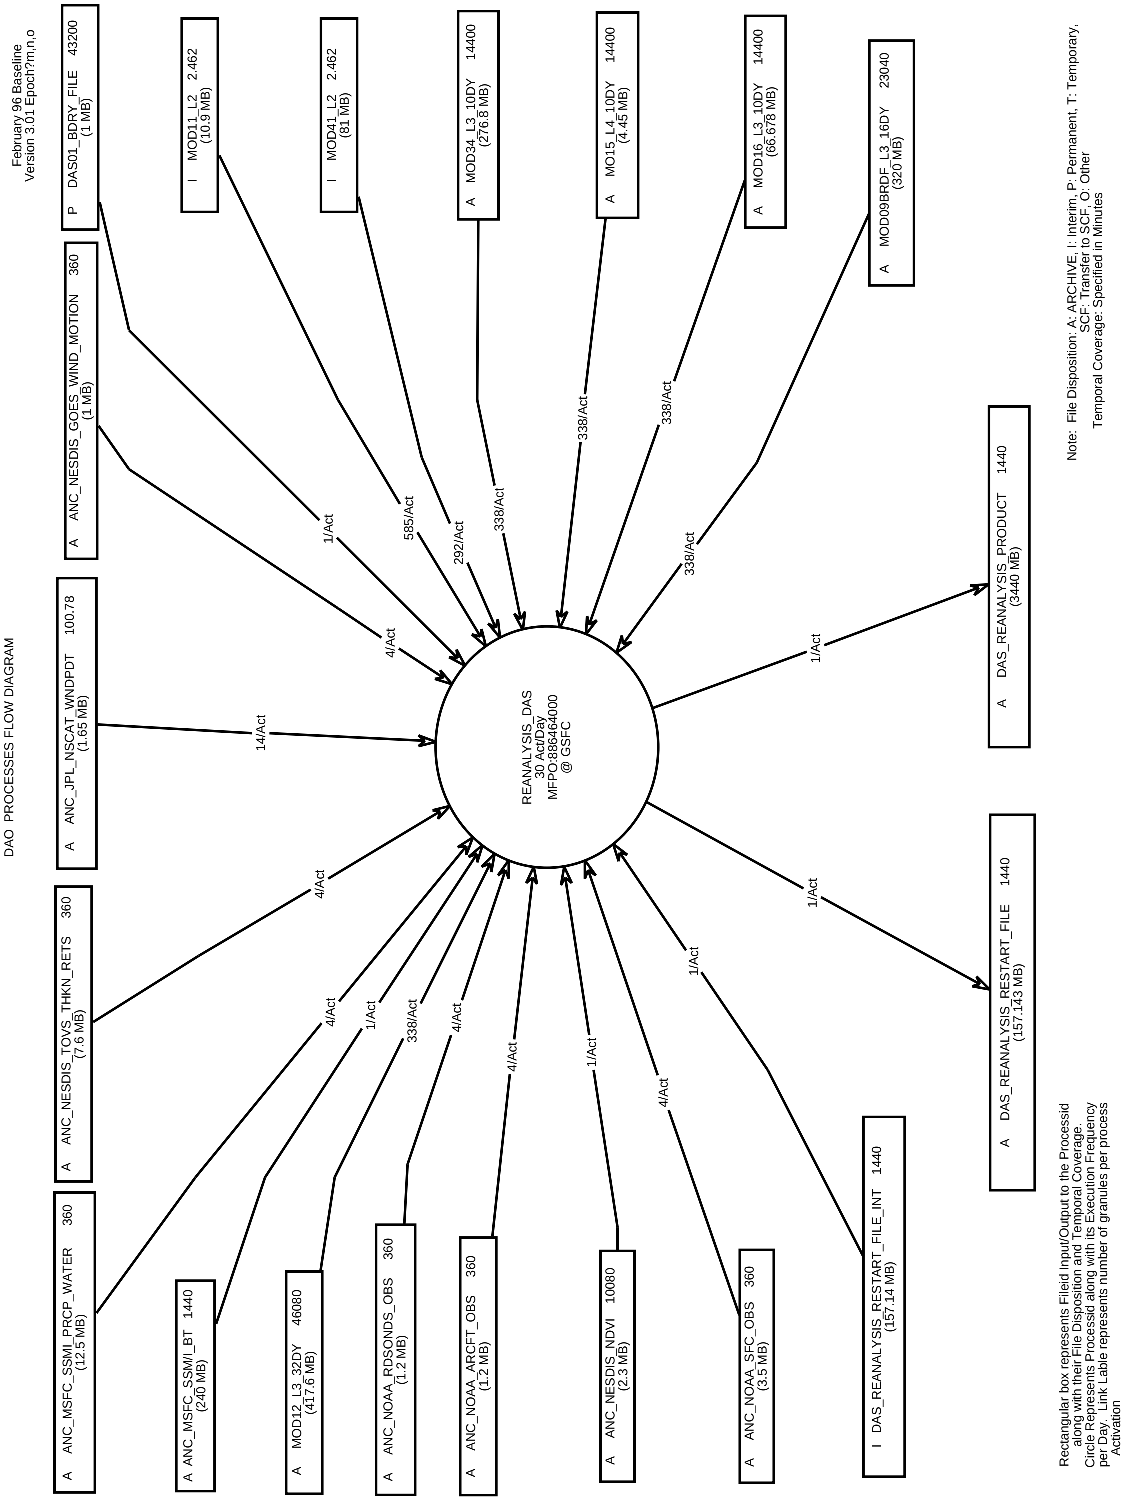

Figure A.3-4. Final-platform Re-analysis

A.4 Simulation Summary

A.4.1 Simulation Results

This section summarizes the modeling results of the DAS system as described by the above scenario. Table 4.1-1 provides resource usage summary in terms of average and maximum for DPS and DSS resources. For this run, resources were not constrained (i.e. more than enough resources were available).

Table A.4.1-1. DAS Resource Utilization Summary

Site ID	PGS Disk Usage (GB)		DH Disk Usage (GB)		Archive R/W Station		Archive Robotics	
	Avg.	Max.	Avg.	Max.	Avg.	Max.	Avg.	Max.
EDC	N/A	N/A	1.15	8.51	4.3	16	0.9015	6
GSFC			3.08	30.35	4.25	20	0.5756	4
GSFC (DAS)	261.29	320.26	N/A	N/A	N/A	N/A	N/A	N/A

Table A.4.1-2. Queuing profile at Data Server subsystem

Site ID	R/W Station Queue		Robotics Queue	
	Avg.	Max.	Avg.	Max.
EDC	60.14	691	0.04	10
GSFC	63.08	941	0.15	16

Table A.4.1-3. Network Throughput (Average)

Source	Destination	Rate in MB/Sec.
GSFC	EDC	1.2446
EDC	GSFC	1.8725
GSFC Data Server	Distribution	1.3665
GSFC Data Server	DAS Processing	3.4031

A.4.2 Planned Simulation Runs

In order to predict a reliable system performance assessment of the DAS system, the modeling assumptions and workload characterization need to be revisited. The following are some of the issues:

R&D workload profile related:

(1) Number of Research Scientist

(2) Mean time between Request/Researchers

(3) Number of jobs/Researcher

(4) Input and Output data volumes and CPU MFLOPS per research job

Super-Obbing related:

1. Data compression ratio (Is it standard for all MODIS data set?)
2. CPU MFLOPS needed to perform Super-Obbing and site where Super-Obbing is performed
3. Location where the Super-Obbed data is archived

A.4.3 Short Term Goals Summary

1. ∞Incorporate R&D workload, MODIS Super-Obbed data sets, revise/update model with workload changes and fully integrate with other ECS instruments load.
2. Provide Results/Analysis to DAS organization for their feedback.
3. ∞Identify critical resources and bottleneck areas and propose design alternatives to maximize DAS system performance.
4. Perform “What If” Sensitivity analysis.

This page intentionally left blank.

Appendix B. Trades and Studies

B.1 Science Data Processing Toolkit:

B.1.1 Introduction

The purpose of the SDP Toolkit is to provide an interface between science processing software and the production system environment. It sets up the context and environment to facilitate portability of code for the execution of production processes and the transfer of data sets and information to those processes. This interface is implemented in the SCF development environment, along with additional utilities that are used to emulate production environment services.

An important goal of the Toolkit is to facilitate the smooth transition and integration of code into the DAAC environment by abstracting out science process dependencies on external system architecture. Another goal is the provision of an interface into which application modules can be incorporated. This may include, for example, math packages; other specialized routines that can be COTS software; freeware; or user supplied modules.

B.1.2 How will the SDP Toolkit meet the needs of the DAS?

The SDP Toolkit will be enhanced in the Release B time frame to provide the functionality required by the DAS. At the present time, no additional user APIs required by the DAS has been identified. It is expected that DAS will be using the following tool groups: Process Control, Status Message Facility, Input/Output, Metadata Access, and Time/Date.

The major additional functionality required by the DAS is support for their selected distributed processing architecture. The effort to incorporate this support will begin immediately after the release of the Release A SCF Toolkit. In addition to modifying the Toolkit code to allow distributed processing, any modifications to simply port the Toolkit to the DAS hardware will also be completed.

The DAS is expected to operate in a distributed processing environment using the Message Passing Interface (MPI). The SDP Toolkit will function in an MPI environment, if each process has access to its own Toolkit library. If however all the processes send some type of remote procedure calls, for example, to one server process in order to access the Toolkit, the Toolkit will then have to be thread safe.

In order to achieve a thread safe Toolkit, the primary areas where the Toolkit will need to be modified are the Process Control and Status Message Facility tools. Although every effort was made to use thread safe functions within these tools, some use of non-thread safe calls was required. There is also extensive use of global static variables which will need redesign. These tools will be redesigned and rewritten to be thread safe. For the case of a parallel shared memory environment, some form of parallel Input/Output capability will be incorporated. Since other Toolkit users have expressed interest in various forms of parallel/concurrent/distributed processing paradigms, much of this work will be done regardless of whether DAS needs it.

ECS has already undertaken a significant amount of prototyping effort using the SDP Toolkit to establish the validity of using parallel and distributed computing to support science algorithms. Information on these prototyping efforts is available on the EDHS Web Site, at the URL: http://ecsinfo.hitc.com/sec6/sdps_intro.html.

Finally, an effort is underway to assist in the development of a 'fly-through' capability which will allow EOS instruments to fly through DAS output data sets. The reuse of some Toolkit geolocation functionality will allow the DAO to develop this capability more rapidly.

B.2 Algorithm Integration and Test

Question: Will DAS S/W undergo SSI&T on the same hardware as it does operations?

Analysis:

There are several aspects of the SSI&T procedures which can be accomplished using hardware not involved in the operations. These include standards checking, checking for correct use of SDP Toolkit mandatory routines, and could include profiling and parallel execution with the operational system.

At some point, however, the DAS Science Software must be integrated with the operational system, and tests will need to be performed to verify that this integration has been successfully accomplished. This integration involves:

- Configuration management procedures.
- Entry of the information necessary for planning and processing into the Plan Database.
- Population of the collection-level metadata into the Science Data Server and the creation of the granule and product-specific metadata tables.
- Insertion of documentation into the Document Data Server
- Establishment/update the advertisements for the DAS products
- Entry/update of the Data Dictionary for DAS products and parameters.
- Placing the scripts and binary executables into the Data Server so that Processing can access these when executing a Data Processing Request (DPR).

B.3 Algorithm Instability

B.3.1 Description of the problem

The DAS is considered to be unstable. Different from PGEs used at most other DAACS, failure can often occur even after a PGE has been thoroughly tested and integrated using test data set. Such failures are most often attributable to input data. This study identifies the features in the design of the ECS that can be used to mitigate problems arising from such failures. This study will not address failures due to coding or operating systems related errors.

B.3.2 Analysis

Statistical estimation techniques, linear and nonlinear Kalman filters, are used heavily in the data assimilation process. The convergence of such recursive estimation techniques depends on the compatibility of the noise model used and the actual noise statistics of the input and already simulated data. The development and tuning of the noise model to be used require the expertise of the scientist developing the algorithm and software. The scientist developing the code could possibly incorporate features to anticipate the occurrence of some failures and identify some of the failure conditions before the estimation begins to diverge. Upon the identification of such failure conditions, the COTS scheduler alarm generation feature could be used to call the attention of the attending operator to intervene and make adjustment before continuing on. The report generation functions of the COTS scheduler Autosys could also be used to maintain reports of intermediate output to assist in the isolation of the problems. To support such an approach, some PGEs, such as the First Look Analysis, may have to be run in the presence of experienced staffs so that the problem can be resolved in a timely manner.

B.4 Research and Development

B.4.1 Statement of the problem

The major functions of most DAACs where the ECS will be put to operations are the production of different types of data products. The other important function is to perform science software integration and testing on the new or upgraded PGEs delivered to a DAAC. Thus all software programs are run under a controlled environment. If a PGE fails under normal operation conditions, it will be sent back to the developer to fix it. At the DAO, the research and development (R&D) needed to produce such PGEs are also carried out on the DAS equipment. In fact much of the resources at DAO are devoted to R&D operations. It is therefore important to investigate if the current ECS design is capable to support such R&D operations.

B.4.2 Analysis

Operations required to perform R&D tasks include code editing, compiling, debugging, data viewing and accessing data residing at the local or remote Data Servers and to perform test runs of the PGEs under development. Facilities to support all these tasks are already available in the SSI&T, PDPS and Client components of the ECS. The first four functions are available in the SSI&T component and the fourth in the Client while the fifth is supported by SSI&T and PDPS. A typical R&D workstation will have an instance of SSI&T and Client running in it. Most of the R&D tasks not requiring high power computing capabilities can be performed at the R&D workstation alone. To support operations requiring high power computing capabilities, Resource Management assigns a set of computing platforms to a group of R&D workstations. Through the Mode management capability of the ECS, the effects of the operations of this group of R&D workstations will be insulated from other operations at the DAAC. Depending on the intended usages of the assigned computation resource, an instance of PDPS can be running at these assigned computing platforms. Using the facilities available in SSI&T, a researcher can register a test PGE in the PDPS data base, generate a production request to execute the test PGE and submit it to the instance of PDPS. Interactive program compilations, executions and debugging will be run under the native operating systems of the assigned computer platforms.

B.5 DAS and the Data Server Sizing

The CDR B design of the Data Server with respect to the DAO at the GSFC DAAC was sized based on the static analysis performed by the Modeling group on the February 1996 Technical Baseline. The "first" and "final" MB/sec flow estimates from the DSS to DAO and from DAO to DSS were taken as 20% of the total data flows to and from DAO. These estimates were used in conjunction with all other data flows to and from DSS, as specified by February 1996 Technical Baseline, to specify the full hardware complement of the DSS. Therefore, based on the current understanding of the DAO requirements, these requirements are fully met within the current DSS hardware architecture. Refinement of the data for DAO requirements could cause the appropriate refinement of the overall design.

Three approaches can be proposed to answer the DAO's concerns about the adequacy of the Data Server design for DAO needs.

1.∞Use the current physical configuration, where DAO equipment is an integral part of the DSS and DAO processing priorities are adjusted within the overall priority scheme.

2.∞Use the current physical configuration, where DAO equipment is an integral part of the DSS and DAO requests to the DSS receive heightened priority within the overall priority scheme.

3. DAO DSS component isolated to dedicated DAO hardware.

Table B.5-1 illustrates the impacts on the technical complexity of the design, schedule and DAO equipment cost that would result if a specific approach is chosen among the above three.

Table B.5-1. Estimated Impact of the DAO Priority Concerns Mitigation Approaches

Option	Technical Complexity	Schedule Impact	Anticipated Cost in DAO H/W	Overall DAAC Impact
1	LOW - no issues within the current baseline assumptions	LOW -no impact	LOW - within current cost estimates	LOW - does not perturb overall operations
2	LOW - within the current design	MED - analysis must be done on the appropriate priority schemes	MED- potential need for additional hardware	HIGH - priorities and ability to service other processing and user requests may be impacted
3	LOW - hardware design can be easily adjusted for adding a separate string; within the current software design	MED - redesign of the GSFC configuration	HIGH-approximate cost of DAO dedicated hardware under the current baseline assumptions is in excess of 5 Million Dollars.	LOW - does not perturb overall operations

B.6 Mass Storage Stress Test

In order to assure that the DAS make efficient use of the computational hardware, a study was conducted to assess the performance of the archive facility. The study evaluated the performance of the Archival Management and Storage System(AMASS) marketed by EMASS. Emphasis was placed on the influence of the file size and system activity on the data transfer rates. This experiment concentrates entirely on stress testing.

B.6.1 System Description

The AMASS system to be tested is connected to an SGI Challenge XL via a SCSI connection. The AMASS consists of one silo or tower which can hold up to 5700 tapes and has one robotics unit and two read/write stations. Each of the tapes that AMASS uses is a 3M 3590E tape that has a 10GB capacity.

B.6.1.1 System Diagram

The AMASS to be tested, has a SCSI connection to a SGI Challenge XL Series which functioned as the FSMS server. Figure B.6.1-2 shows the Goddard DAAC hardware diagram with detail about the AMASS system.

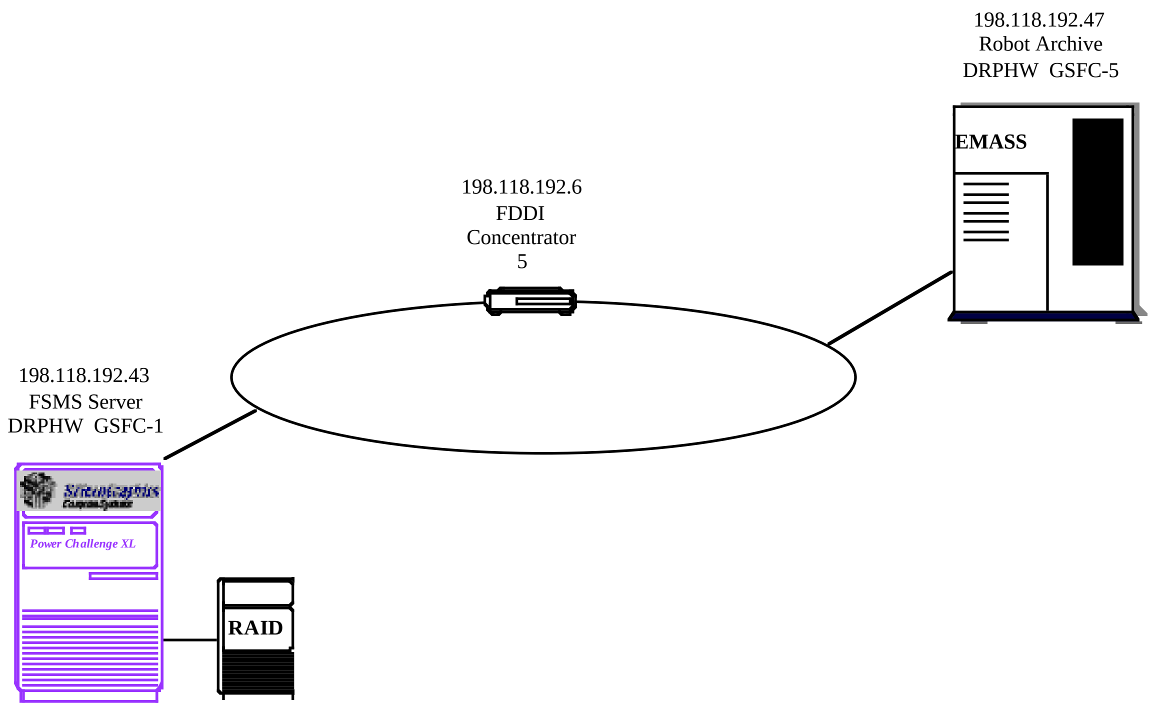

Figure B.6.1-1. System Diagram

B.6.1.2 AMASS Diagram

AMASS is a virtual file system with a linked separate design that allows all UNIX File System (UFS) access methods to be employed. AMASS maintains all index node information in data base files rather than associated disk structures. AMASS utilizes a disk based I/O buffer for communication rate matching between disk and tape resources. Figure B.6.1-2 shows a diagram of the AMASS architecture.

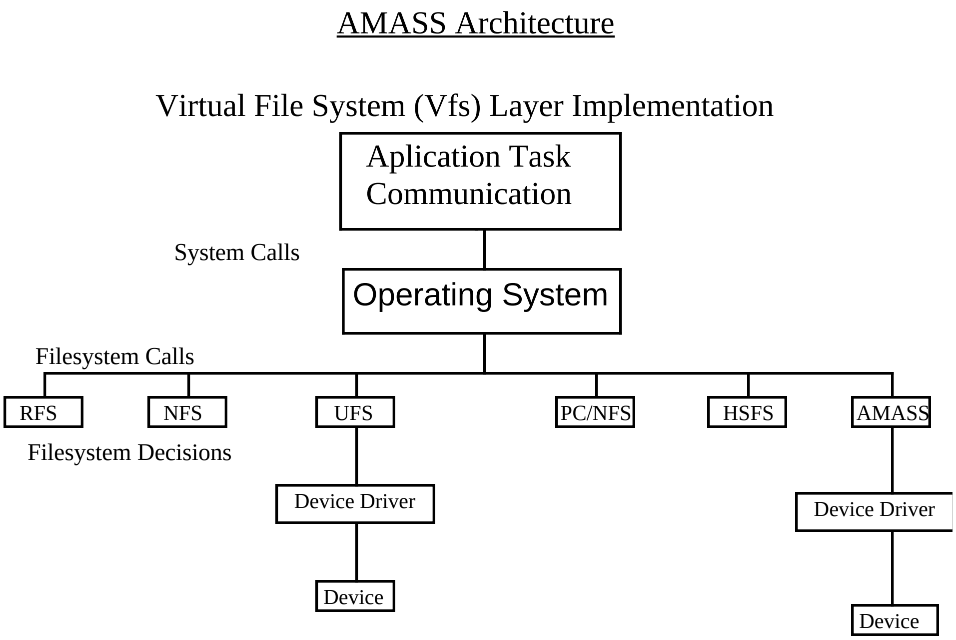

Figure B.6.1-2. AMASS Architecture

B.6.2 Analysis Technique

File throughput performance analysis began with the log file maintained as a measurement record for all the read and write operations that were conducted. A UNIX based log file will contain the records for each event conducted throughout the entire experiment. Each record will contain the following fields:

A. Start - The date and time stamp for the start of the event.

B. Finish - The date and time stamp for the completion of the event.

C. Size - The size of the file used in the event.

The following fields can be calculated or derived from the information that is maintained the records of each event:

1. Throughput Time - File throughput stop time minus start time.
2. Access Rate - The effective throughput rate can be calculated as the ratio of the Size to the Throughput Time, and is presented in MB/Sec.
3. Volume Number - This can be determined from the relationship of the Size and the file name.

Once in a spreadsheet format, the measurement results will viewable in several ways to determine distribution and correlation. This material will be presented in subsequent sections.

B.6.3 Variables

There are a number of variables that will affect performance when accessing files located on the AMASS. This section will identify the major variables that will directly affect the throughput of files residing on the AMASS.

B.6.3.1 File Size

The file size is one of the principle variables that will affect the file access performance. To model this variable the experiment used diagnostic files sized in five classes: 1 MB, 15 MB, 95 MB, 400 MB, and 1 GB.

B.6.3.2 System Activity

Another factor affecting file access performance will be resource competition caused by system activity. No direct method of calibrating or isolating system load was undertaken. Instead, it will be assumed that the system activity will fluctuate itself on access performance as a function of the time of day.

B.6.3.3 Disk Cache

Disk Cache is one of the variables that can be modified on the AMASS. During this testing the disk cache will be set up in such a way to allow a full volume to be copied into the cache memory.

B.6.3.4 Block Size

Block size for this testing has been estimated to be optimal if it can hold up to the average file size. It is very critical to have an optimal blocking size to prevent empty or unused disk space while read and write processes are running and also to not cause a bottle neck during the same type of operations.

B.6.4 Method

Cron jobs will execute C-shell scripts on the SGI Challenge XL. These shell scripts will access single files as well as batches of files from the AMASS. The files accessed will be copied from the archive to an alternate location not on the archival system. The total access time will be measured by bracketing the shell scripts with UNIX date commands. Access rate will be varied throughout the experiment and the transfer rate will be computed as the ratio of file size to the elapsed time.

To best conduct this experiment it was judged that at least five volume groups would be necessary, each consisting of 32 volumes(one on each face/surface of the AMASS silo). Thus, there would be a total of 160 volumes used in the experiment. The distribution of the volume groups was as follows:

Group 1, 2, 3, 5:

12 volumes full

12 volumes half full

8 volumes empty

Group 4:

8 volumes full

8 volumes half full

8 volumes spanned (part of a file on two tapes)

8 volumes empty

It was also established that Group 1 would contain only 1MB files, Group 2 15MB files, Group 3 95MB files, Group 4 400MB files, and Group 5 1GB files.

B.6.5 Overview of the Experiment

In an effort to produce a full characterization of the AMASS hardware it is critical that the experiments conducted closely match the intended use of the system by the DAO. To accomplish this, three major categories of experiments were conducted.

First, there was the calculation of file write speeds for various size files. The focus of this experiment was to have multiple file creation processes running to insure that AMASS was always waiting for a tape transfer not a file creation.

Second, there were scripted file retrievals against full volumes using different files per volume group. The focus of this experiment was to run continuous retrievals over one or more hours to determine a normalized hourly rate per file size/volume group.

Third, there were scripted alternating read and write events of files per volume group on half full volumes. The focus of this experiment was to get a read/write throughput number of a given file size and to run continuous operations over one or more hours to determine a normalized hourly rate.

Other tests that were run include the measurement of the robot exchange rate without tape thread. Essentially, this consists of, Drive 1, find tape, grab tape, load, Drive 2, find tape, grab tape, load, Drive 1, eject tape, grab tape, put away, Drive 2, eject tape, grab tape, put away. This experiment doesn’t include tape thread or rewind in its final calculations to characterize near maximum exchange rate per hour. Also, there were measurements taken that characterize AMASS functions such as VOLSTAT, used to activate and deactivate volumes, and SYSPERF, an AMASS provided function for system performance.

B.6.5.1 Problems Encountered

The major problem encountered during this testing was familiarity with the hardware and network problems. As far as the hardware is concerned all problems were non existent after the first week of testing. Network problems continued throughout the entire testing period. Most problems were due to the active work to configure the Goddard Hub.

B.6.5.2 Statistics

File throughput performance measurements were taken over a 28 day period from Monday 29 April 1996 to Friday 24 May 1996. The throughput results are plotted against the measurement timeline in Figure B.6.5-1. All measurements are in MB/Sec. Figure B.6.5-1 represents read operations only. The rates shown here are effective throughput rates, all forms of delay are included in this measurement. The different forms of delay are network, system, robot, and drive. The first five measurements in Figure B.6.5-1, numbered 1 through 5, read data from cache and therefore, show data rates well above 0.6 MB/Sec. Subsequent experiments forced retrieval of data from the tapes. The effective data rate in these trials were between 0.16 and 0.18 MB/Sec. These observed performance is well above the baseline assumptions that were used for Release B.

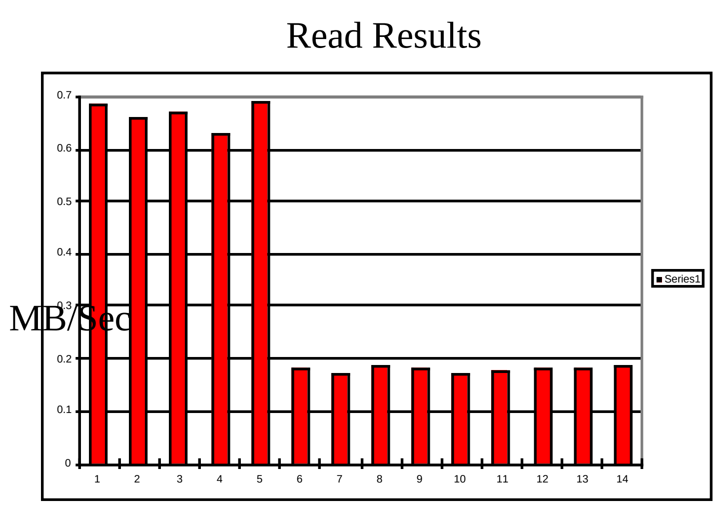

Figure B.6.5-1. Results of Mass Storage Stress Test

The write throughput results are plotted against the measurement timeline in Figure B.6.5-2. All measurements are in MB/Sec. These rates are also effective throughput rates that include all forms of delays as mentioned in the previous paragraph.

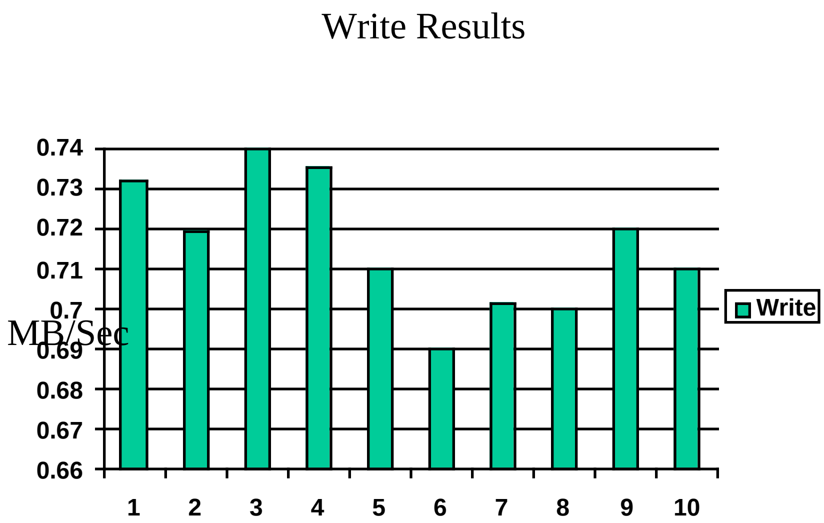

Figure B.6.5-2. Results of Mass Storage Stress Test

B.7 ECS Scheduler Compatibility

B.7.1 ECS Scheduler and DAS Computational Needs

In its current design, the ECS scheduler consists of two major components, Planning (PLANG) and Data Processing (PRONG). PLANG develops the tasks schedule, *the active plan*, for a DAAC for a given time period based on the projected availability of resources and input data. When all the input data for a given PGE in the active plan becomes available, it is released to PRONG for execution. It is one of the responsibilities of PRONG to manage the available resources to support the executions of all released PGEs and to monitor the processing status of the PGEs after they have been executed. In the event of detecting an anomalous processing status for a PGE based on its processing profile, warning signals will be raised by PRONG to alert an operator to the failure. PRONG is also responsible for the real time assignment of resources for input data staging, including predictive staging as needed, and CPUs for PGEs. It also coordinates the executions of all these operations to minimize conflicts and to ensure the delivery of products in a timely manner.

The DAS equipment supports the following types of tasks: First-look analysis, Ten-day forecast, Final-platform analysis, Re-analysis, and R&D. The basic relevant features of these five types of activities have been summarized in section 6 of this white paper. There is an important feature of DAS that makes it different from most other ECS science data processing. It is the need to use parallel processing for most of its computations. The computational requirements are of such a large scale that it will require a super computer in the class of a Cray XMP platform or many SMP platforms such as the SGI/PCA to complete the required tasks in a timely manner. The detailed computational architecture to support such large operations is being developed at DAO and other scientific centers. The percentages of resources at DAO dedicated to each of the five types of operations changes dynamically during a typical day and it also depends on whether re-analysis would is in progress that day.

B.7.2 Trade Alternatives

To accommodate the special needs of DAO, two alternative approaches have been identified:

1. Accommodate the DAS computational needs, including supports for R&D and multiple platform parallel computation, using only the current ECS infrastructures with minimal changes.

2. Modify the current ECS design by introducing a Super Computer Platform (SCP) with a comprehensive operating system that can effectively handle the parallel DAS processing needs. All processing jobs will be submitted to the SCP. The SCP operating system will handle all the resource allocation, actual scheduling. Performance monitoring will be still be done by Data Processing.

Analysis:

In either approach there will be little or no impacts on the operations of PLANG. However, the impacts on PRONG will be different.

B.7.2.1 Approach 1

Projected changes in PRONG:

Some modifications of the PRONG software may be necessary to adapt it to the parallel computation requirements of DAS. Resource management components of PLANG and PRONG can assign enough staging disk and CPU resources for each of the four type of PGEs, First-look analysis, Ten-day forecast, Final-platform analysis, or Re-analysis, before their executions. The detailed management of the resources during execution will be left to the particular PGE in this case. For the purpose of monitoring the operations of a PGE during its execution some messages may have to transmitted to PRONG by the PGE using new functions of the Toolkit.

Although in the current design of ECS, there is no special provision to directly support R&D operations at a DAAC. All the required operations for R&D, such as interactive code editing, compilation, debugging, accessing local or remote data servers and running test PGEs are already available in the SSI&T, Client and PDPS components of the ECS. Details of how these functions can be used to develop R&D facilities are given in Appendix B.4. Resource management may change the resources assigned to production, I&T, or R&D during a given day. During such changes, staff engaging in R&D work usually will not have to log off, but will detect changes in the turn around times for the jobs submitted.

B.7.2.2 Approach 2

Projected changes in PRONG :

1. Develop a comprehensive operating system for the SCP that can effectively handle the parallel DAS processing needs, all the resource allocation and scheduling of the production and R&D operations.
2. Modify PRONG so that it no longer directly manages the memory and computation resources, but will only submit jobs to the SCP and stage and destage data to a buffer.

Using this approach, for each interactive R&D session an SS&T instance and a Client instance will be active. The dynamic changes of resources assigned to each type of operations will be changed by the SPC operating system in a way that will be transparent to the different active R&D sessions, with perhaps detectable changes in the turn around times of submitted jobs.

B.7.2.3 Summary

Table7.2-1 shows a comparison of the two different approaches. Approach 1 is recommended, since it will be more cost effective and will accommodate the DAS parallel processing needs with minimal changes to the software in both the applications and ECS software.

Table B.7.2-1. Comparison of Approaches

	Advantages	Disadvantages
Approach #1	• The changes the ECS design will be minimal. • Reconfiguration is flexible and can be changed dynamically. • Provide physically separated environments for production and R&D. • Taking full advantages of the original design concepts of the ECS system to support many automated features in resource planning/allocation, production planning/real time scheduling.	• Evolutionary changes to DAO needs may not be isolated.
Approach #2	• Evolutionary changes to DAO needs will most likely be isolated to the SCP	• Requires a complete redesign of some of the basic functions of the PDPS subsystem to recapture the original ECS design concepts.

B.8 ECMWF Scheduler

B.8.1 Trade Description

In the current design of ECS, the COTS AutoSys is used for real-time scheduling of all the data processing tasks on the hardware at a given DAAC. AutoSys is developed and supported by PLATINUM Technology, Inc with headquarters at Boulder, Colorado. In the mean time another similar scheduler, currently in use at the European Center for Medium-Range Weather Forecasts (ECMWF), has come to the attention of the DAO staff. It is the purpose of this study to compare these two schedulers and find the one that will be more suitable for the ECS to be installed at DAO. For the purpose of this study the following documents were used:

1. AutoSys Version 3.2 User Manual (August 1995).
2. SMS and CDP User Guide and Reference Manual (August 1992).
3. XCDP User Guide and Reference Manual, of the ECMWF scheduler. (August 1992).

The last two documents were prepared by the ECMWF staff for the scheduler used there.

B.8.1.1 Overview of the two schedulers

Both schedulers are designed to run under the UNIX operating system and schedule jobs over a network using RPCs. They are both general purpose schedulers. AutoSys provides a set of commands, the Job Information Control Language (JIL), written in C that use Remote Procedure Calls, and can be used develop scripts for job scheduling, defining monitors and reports. AutoSys also provides a GUI to perform the same functions. The ECMWF scheduler provides a similar set of commands through its CDP

(Command and Displace Program) text interface, for job scheduling, defining monitors and reports. It also provides a GUI interface, XCdp, that displays graphically the status of the tasks under its control. Both schedulers have the basic functions required for a UNIX scheduling tool, however, Autosys is a more mature product in terms of available features. Table B.8.1-1 summarizes features of Autosys and ECMWF.

Some of the features not available in the ECMWF scheduler are expected to be useful to ECS.

One example of this is the shadow event processor in AutoSys that runs on a different machine from the primary Event Processor. It is normally idle, just listening for regular messages from the primary Event Processor indicating that it is still functioning. The shadow event processor takes over (becomes the primary event processor) when these messages are not received for particular length of time. This allows AutoSys to continue to function in the event that the machine housing the primary Event Processor fails.

AutoSys also makes extensive usage of Sybase. Each instance of AutoSys is supported by three databases, the Event Server, Monitor Server and Alarm Server. The Event Server contains all the information about a particular instance of AutoSys. It is the Event Server that the Event Processor reads to determine what jobs to schedule and execute. As such it also provides a very effective means for all the remote agents residing on different processing machines to communicate with each other. Monitor and Alarm Servers are for maintaining status and reports. Using the built in features of Sybase, it will be easy to maintain backup copies of all the three databases for recovery purpose in the event of system crashes. Such built in features will also expedite the process of analyzing the historic data and obtain other derived information for performance tuning or other purposes.

AutoSys is widely used commercial product that is continuously evolving. As such it will be easier to obtain technical support and respond to new requirements.

Based on the previous analysis, AutoSys is a more suitable scheduler for our present applications.

Table B.8.1-1. Comparison of features between AutoSys and the ECMWF Scheduler

Feature Description	AutoSys	ECMWF
Trigger a process when certain predefined time, resource and/or environment constraints are met and modify such conditions dynamically.	Y	Y
Control the execution of tasks on multiple platforms.	Y	Y
Capability to specify complex triggering conditions of a job or a group of jobs on other processes and to modify such dependencies dynamically.	Y	Y
More than three levels of job hierarchies.	Y	N
Graphical capability to display dependencies of jobs and to monitor execution status.	Y	Y
Capability to restart a task a predetermined number of times before giving up.	Y	Y
Maintain detailed history of all the events that have been triggered by the scheduler and ability to capture the occurrence of other events defined by a user.	Y	Y
Raise alarms to an attending operator upon the occurrence of certain unexpected events.	Y	Y
Repeat the attempt to initiate a task a predefined number of times before giving up completely.	Y	Y
Define waiting time between successive retry attempts.	Y	N
Shadow scheduler or event processor.	Y	N
Use of a data base to maintain schedule events, log file and reports.	Y	N
Features to support load balancing.	Y	N
File watcher to monitor the status prescribed files	Y	N

B.9 Multiple Platform Scheduler

How will the ECS scheduler manage a shared memory distributed computing architecture?

B.9.1 Analysis

Due to the decision that, for PGEs requiring the use of multiple platforms for its execution, the detailed management of the required resources will be handled by the software program of the PGE, multiple platform scheduler is no longer an issue. The scheduling of such PGEs should be no different from a PGE that requires a single CPU. Before such a PGE can be executed all the required CPU and staging disk resources will be allocated. The required input data will be staged to the designated location. After that the PGE will be started, and the PGE will be expected to manage the assigned resources among its multiple processes on its own. There are infrastructures designed to support the normal operations of the ECS system that can be used to support distributed parallel computations over multiple SMP platforms.

B.10 Algorithm Delivery Schedule

The First-Look and Ten-day Forecast software does not rely on other EOS data, and can be tested independently of the testing of the science software for the EOS instruments. It is recommended that the SSI&T procedures be performed on an engineering-version of the DAS software system for these products at least 18 months before launch, although where this can be done is an open question. The mission-ready version of the software for these products should undergo pre-launch SSI&T approximately 9 months prior to the AM-1 launch.

The Final Platform Analysis utilizes other EOS products as inputs and will not be placed into production until after the validation of the dependent EOS instrument data products, which will not be accomplished until many months after AM-launch. It is recommended that a version of the Final Platform Analysis software undergo integration with ECS prior to launch and be run regularly after AM launch, but that the post-launch "commissioning" period be used for testing and tuning. The products generated during this commissioning period need not be permanently archived or generally available, but would allow a more comprehensive test of the Final Platform Analysis utilizing actual data in a production environment.

For Super-Obbing and other DAS PGEs that will run on science data processing hardware located at other DAACs, these PGEs need to be delivered to the affected DAACs and undergo SSI&T just like PGEs from any Instrument Team. The DAO and the relevant DAAC(s) will need to discuss the schedule for delivering and performing SSI&T.

B.11 Communications Fabric

B.11.1 Detailed Architecture

The DAS will interface to ECS via a switched port on the FDDI Switch as shown in figure B-11.1-1.

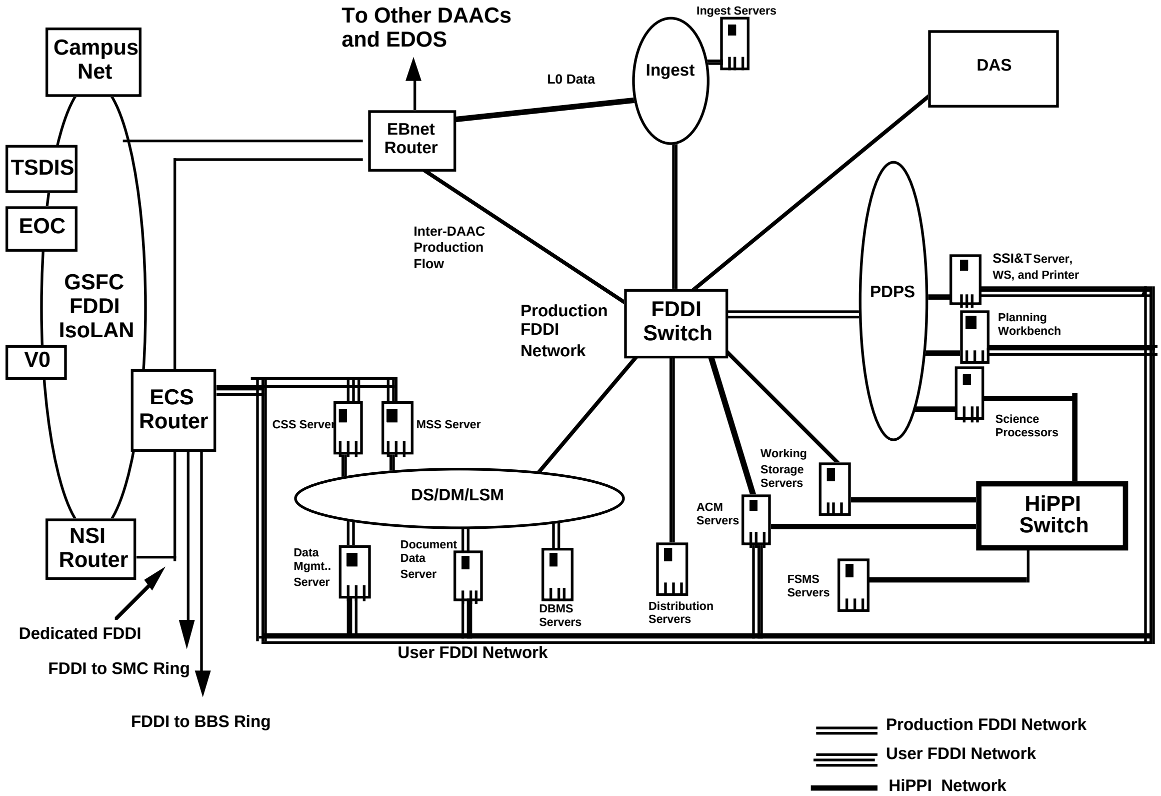

Figure B.11.1-1. FDDI Connection Between DAS and ECS

All of the DAS workstations will connect to the FDDI fabric on their own switched segment. Each workstation will be attached to an FDDI concentrator as shown below. Depending upon the DAS hosts availability requirements, they will be attached to the concentrators as single attached or dual attached (dual homed).

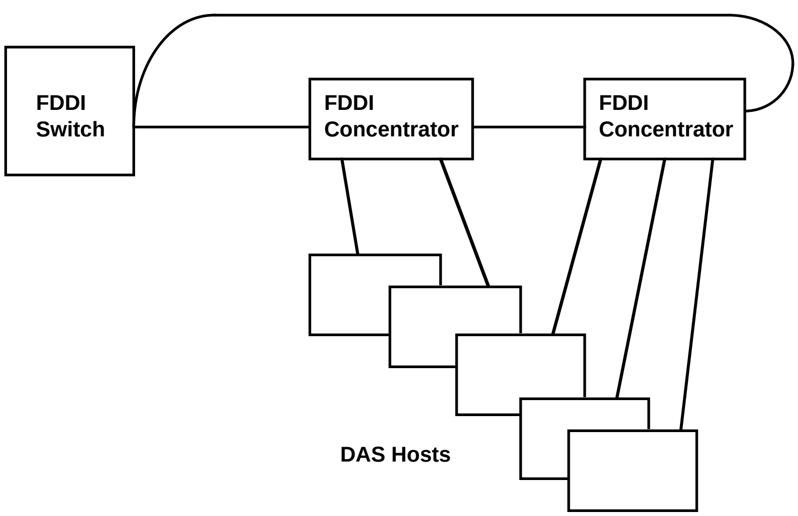

B.11.1-2. DAS FDDI Connectivity

Host addresses will be assigned from the GSFC ECS DAAC address domain.

If further analysis shows that the data flows exceed an FDDI, then selected DAS hosts can be connected to the HiPPI fabric.

B.11.2 Communication Fabric

The following paragraphs describe an optimum way of connecting the DAS processing equipment to the ECS DAAC at GSFC. This approach was developed using the projected data flows for the DAS.

B.11.2.1 Evaluation Criteria

The selection of the appropriate communication fabric is determined by the data flow between DAS and ECS and the current ECS communication fabric design.

B.11.2.2 Evaluation Procedure

Before we can determine how the DAS environment should be connected to ECS, we need to know the data flow between DAS and ECS.

The data flow between DAS and ECS (the Working Storage Servers) were calculated using a dynamic model. This model shows the amount of user data flow between DAS and ECS at the Q4 1999 time frame. To this user data flow, network overhead is added. Based upon the user data flow and added network overhead, an appropriate communication fabric is selected.

B.11.2.3 Analysis

The following data flow information was obtained from the dynamic model run on 3/1/96:

- Data flows from the dynamic model = 13.55 Mbps
- Adding network overhead, the weighted data flow = 38.12 Mbps

Network Overhead Factors:

1. TCP/IP Protocol Overhead: 1.25. Accounts for overhead associated with packet frames, IP, TCP, and other protocols (such as DCE).
2. Maximum Circuit Utilization: 1.25. Accounts for the amount of 100Mbps that is actually usable for sustained data rates.
3. Average-to-peak Conversion: 1.5. This provides elasticity in the network by converting the 24 hour average provided by the model into peaks.
4. Scheduling Contingency: 1.2. This reflects the ability of the network to recover within 24 hours from a 4 hour down-time (24/20 = 1.2).

Based upon the above data flows, an FDDI or 100Base-T switched ethernet interface would provide adequate bandwidth. Both of these interfaces provide a 100Mbps data rate.

Because the current design of ECS uses an FDDI switch as the central connection for the Production Network, using an FDDI interface would be the best solution.

Note: This analysis is for the aggregate data flows only between DAS and ECS. It does not account for any data flows within DAS or additional data flows within ECS.

B.11.2.4 Conclusion

Using the data flows from the dynamic model of 3/1/96, an FDDI interface is the best solution as shown below. If further analysis shows that the data flows exceed the FDDI capacity, then selected DAS hosts can be connected to the HiPPI fabric. A HiPPI interface provides an 800Mbps data rate.

B.12 Multiple Platform Management

The purpose of this study was to determine whether ECS system management concepts and implementation needed to be extended to integrate the existing and planned multiplatform / multiprocessor architecture of the Data Assimilation System (DAS). The results of this study present the DAO with three options and associated tradeoffs for managing the DAS configuration.

B.12.1 ECS Systems Management Overview

ECS system management provides support for the management of performance, faults, security, configuration, and accountability in a distributed client/server environment. The implementation employs a Network Management Framework and an Enterprise Management Framework.

B.12.1.1 Network Management Framework

The Network Management Framework allows for monitoring and controlling network resources and applications from a central management position. The framework is based on the Simple Network Management Protocol (SNMP) standard augmented by secure communication enhancements implemented using the Distributed Object Framework (DOF). This management framework supports the issuing of life-cycle commands controlling managed object operation, the monitoring of status and performance, and the receipt of fault and performance event notifications. The framework has been extended in Release B to manage system resources applied to test and training modes. The HP Openview COTS product has been chosen for the management station. It provides a GUI-based interactive interface allowing M&O personnel to manage network devices and host nodes. It also allows for management of ECS custom applications which have been instrumented to communicate with Openview using the ECS-developed application MIB and the Managed Process Framework. Openview communicates with agents residing on each managed host node. The Peer Network vendors SNMP agent development kit has been chosen as the extensible agent residing at each managed host.

B.12.1.2 Enterprise Management Framework

The Enterprise Management Framework builds on the Network Management Framework supplying higher level management functions, tighter integration, and automation of these functions. It provides for the establishment and integration of management regions, i.e., DAACs and SMC.

The Enterprise framework is implemented using the Tivoli vendor's Tivoli Management Environment (TME). TME provides a platform consisting of a GUI, a Toolkit for integrating COTS and custom applications, and a facility through Tivoli/Plus modules to integrate other management application vendor products. Examples of Tivoli/Plus modules provided include Remedy, Sybase, Networker, and Autosys. The Release B configured TME supports system administration functions such as an interface for Distributed Communication Environment (DCE) administration, system backup/restore, Sybase Database administration, software distribution, enhanced performance monitoring, fault correlation, and automatic initiation of fault corrective action.

B.12.2 DAS Management Configuration Options

The DAS is a distributed processing system composed of supercomputing multiprocessor hosts. Execution of product generation applications is distributed across processors on multiple hosts to maximize throughput and product turnaround. The DAS has an existing complement of product generation applications and an existing means of monitoring and controlling these applications.

The ECS management system accommodates all aspects of DAS operation in its standard form as designed for Release B. However, to take advantage of this management capability, there are costs associated with upgrading the existing DAS hosts and applications with the ECS management frameworks. The following three management configuration options are available to DAO.

B.12.2.1 Option 1

Use only host vendor supplied SNMP Agent/ Host MIB II.

Figure B.12.2-1 is the minimal DAS host configuration option which would allow DAS hosts to be integrated within the ECS framework. Under this option, each DAS host is equipped with the vendor supplied Host MIB II and an agent to communicate with the ECS HP Openview management platform. This configuration is the least cost and allows for the monitoring of host resources and receipt of SNMP faults and alerts generated by the host. This configuration provides the basic network management capabilities but does not support the management of applications on the host nor does it support the higher level enterprise management functions.

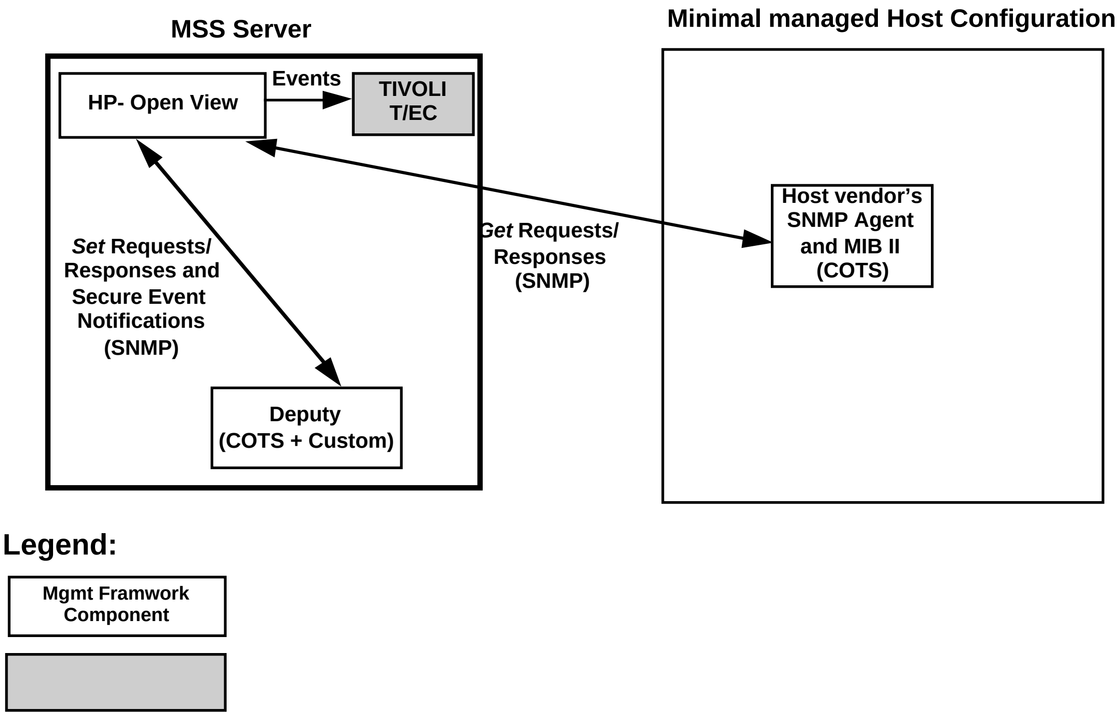

Figure B.12.2-1 - Minimal Management Configuration

B.12.2.2 Option 2

Use Tivoli Sentry on Host plus vendor provided SNMP agent

The second configuration option is to add the Tivoli Sentry module to the configuration of Figure B.12.2-1. This would allow more extensive monitoring of faults and performance from the Tivoli enterprise management position although it still would not support management of DAS applications.

B.12.2.3 Option 3

Use standard ECS management configuration

Figure B.12.2-2 is the fully configured standard ECS host management configuration. It requires the addition of the Peer Network vendor's extensible agent to the option 2 configuration and the integration of DAS applications with the ECS Managed Process

Framework and application MIB. This configuration offers the maximum advantage in terms of ability to evolve within the ECS management standard but does entail the initial cost of upgrading existing DAS applications.

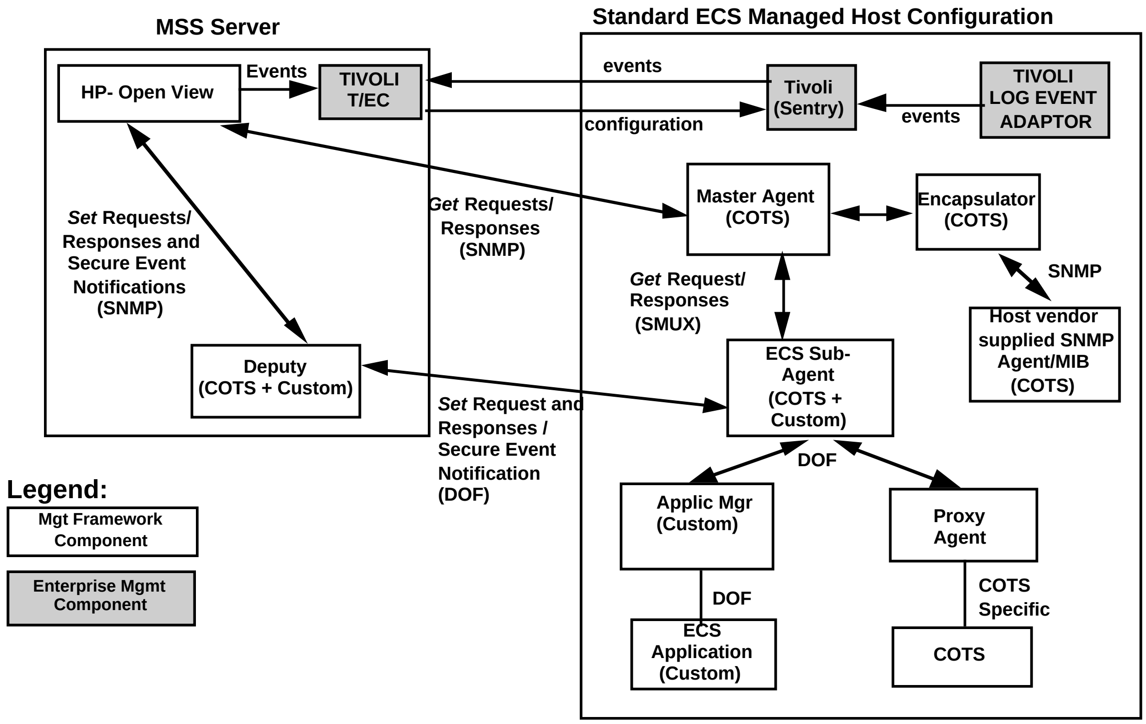

Figure B.12.2-2. Standard Management Configuration

B.12.3 Recommendation

The recommended management configuration for the DAS would be option 3, the ECS standard management configuration. The primary motivation is the ability to evolve and take advantage of the fast advancing improvements in enterprise management capabilities. However the preponderance of existing DAS applications are a consideration in determining the time frame for integrating the full management standard. From a cost and risk mitigation standpoint, it might be advantageous to consider the three management configuration options described above as a phasing plan where the DAO starts with the minimal configuration and progresses to the full ECS configuration incrementally as budget and schedule allow.

B.13 Communication Services

Open System Foundation's (OSF) Distributed Computing Environment (DCE) is a set of integrated services for building distributed applications across multiple systems. DCE is a mature technology and has been chosen as the technology baseline for the ECS communication infrastructure. Built on top of DCE with helps from object-oriented technology (which leads to the object-oriented DCE, i.e. OODCE,) this ECS infrastructure consists of a full set of communication services to support interactions between and within all other subsystems in ECS . ECS communications services are standards-based and inter-operable for applications within or outside of ECS.

Based on scenarios in various DAS operation modes and the current understanding of DAS, the following communications subsystem services are identified as required for the DAS integration with ECS:

- The Messages Exchange Service for session initialization;
- The File Access Service for data transfer;
- The e-mail Service for status report or acknowledge;
- The Security Service for authentication and authorization;

Other related issues have also been analyzed, including the Toolkit usage, platform selection, managed mode environment and threads safe environment. It is found that DCE can provide extra benefits in many cases, especially for secure communications. However, there will be no major breakdown for DAS integration with ECS without DCE at the DAS side.

It is concluded that DCE is not mandatory for the DAS integration at present time. For the DAS operations itself, this is obvious since DCE is not required to run science algorithms on the DAS supercomputer platform. The ECS Toolkit does not need DCE at present. For security purpose, Kerberos protocol can be used for authentication to ECS. In the future time, if DCE based communications with ECS are needed due to incremental development and enhancements, a communications gateway can always be developed.

However, DCE is still strongly recommend for the full utilization of ECS communications services. For now, installing only DCE client at DAS might be the best way. In that case, DAS will contain the minimum DCE intelligence to request ECS services, while keeping its own independence. Figure 13-1 shows a high-level diagram of the DCE installation.

Under this design, authentication to ECS servers and messages exchanges with ECS subsystems like PDPS and MSS will use DCE Remote Procedure Calls (RPC), while data transfer (ingest or distribution) will use socket/ftp directly. Local DAS operations will require no change and use sockets or existing inter-process communications service.

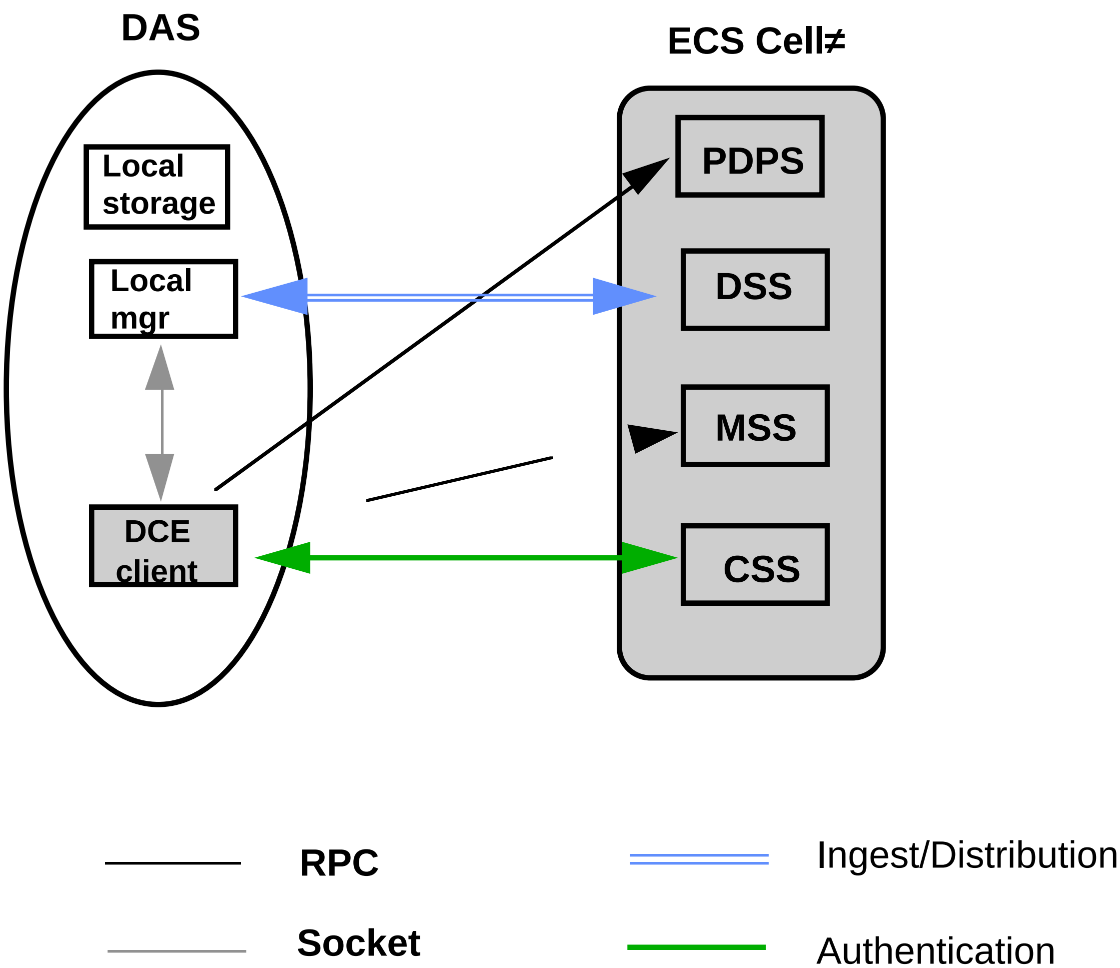

Figure B.13-1. Recommended DCE installation

This page intentionally left blank.

Abbreviations and Acronyms≠

AHWGP	Ad hoc Working Group on Production
BOM	Bill of Materials
CERES	Clouds and Earth’s Radiant Energy System
COTS	Commercial-of-the-Shelf
CPU	Central Processing Unit
CSMS	Communications and System Management Segment
CSS	Communication Subsystem
DAAC	Distributed Active Archive Center
DAO	Data Assimilation Office
DAS	Data Assimilation System
DCE	Distributed Computing Environment
DEC	Digital Equipment Corporation
DH	Data Handler
DPS	Data Processing Subsystem
DPRHW	Deep Repository Hardware
DSS	Data Server Subsystem
ECS	EOSDIS Core System
EDC	EROS Data Center
EGS	ECS Ground Segment
EOSDIS	Earth Observing System Data and Information System
FDDI	Fiber Distributed Data Interface
GDAAC	Goddard DAAC
GFLOPS	Giga Floating point Operations Per Second
GSFC	Goddard Space Flight Center
HiPPI	High Performance Parallel Interface

HP	Hewlett Packard
IBM	International Business Machines
I&T	Integration and Test
I/O	Input/ Output
IV&V	Independent Verification and Validation
LaRC	Langley Research Center
MB	Megabyte
MODIS	Moderate Resolution Imaging Spectroradiometer
MPI	Message Passing Interface
MSS	Management Subsystem
MTPE	Mission to Planet Earth
NCCS	National Center for Computational Science
NMC	National Meteorological Center
NOAA	National Oceanic and Atmospheric Administration
OODCE	Object Oriented Distributed Computing Environment
OSF	Open System Foundation
PGE	Product Generation Executive
PGS	Product Generation System
PLS	Planning Subsystem
PPR	Pre Procurement Review
PVP	Parallel Vector Processor
R&D	Research and Development
RFP	Request for Proposal
SCF	Science Computing Facility
SCP	Super Computing Platform
SDPS	Science Data Processing Segment
SGI	Silicon Graphics Incorporated
SMP	Synchronous Multi Processor

SPMN	Same Program Multiple Nodes
SSI&T	Science Software Integration and Test
WG	Working Group
WKSHW	Working Storage Hardware

This page intentionally left blank.