



HUGHES INFORMATION TECHNOLOGY SYSTEMS

ERRATA NOTICE

EOS Core System (ECS) Project Contract No. NAS5-60000

April 19, 1996

Document No.: 305-CD-030-002

Title: Release B GSFC DAAC Design Specification for the ECS Project

Enclosed please find change pages for the subject document. Please replace the pages as follows:

Remove

3-33 through 3-36

3-39 through 3-44

3-61 through 3-62

3-67 through 3-68

3-105 through 3-118

Insert

3-33 through 3-36

3-39 through 3-44

3-61 through 3-62

3-67 through 3-68

3-105 through 3-118

If you have any questions, please contact our Data Management Office at (301) 925-1037.

305-CD-030-002

EOSDIS Core System Project

Release B GSFC DAAC Design Specification for the ECS Project

March 1996

Hughes Information Technology Systems
Upper Marlboro, Maryland

Release B GSFC DAAC Design Specification for the ECS Project

March 1996

Prepared Under Contract NAS5-60000
CDRL Item #046

APPROVED BY

Rick Kochhar /s/	3/29/96
<hr/>	
Rick Kochhar Release B CCB Chairman EOSDIS Core System Project	Date

Hughes Information Technology Systems
Upper Marlboro, Maryland

305-CD-030-002

This page intentionally left blank.

Preface

This document is one of eighteen comprising the detailed design specifications of the SDPS and CSMS subsystem for Release B of the ECS project. A complete list of the design specification documents is given below. Of particular interest are documents number 305-CD-020, which provides an overview of the subsystems and 305-CD-039, the Data Dictionary, for those reviewing the object models in detail.

The SDPS and CSMS subsystem design specification documents for Release B of the ECS Project include:

- 305-CD-020 Release B Overview of the SDPS and CSMS Segment System Design Specification
- 305-CD-021 Release B SDPS Client Subsystem Design Specification
- 305-CD-022 Release B SDPS Interoperability Subsystem Design Specification
- 305-CD-023 Release B SDPS Data Management Subsystem Design Specification
- 305-CD-024 Release B SDPS Data Server Subsystem Design Specification
- 305-CD-025 Release B SDPS Ingest Subsystem Design Specification
- 305-CD-026 Release B SDPS Planning Subsystem Design Specification
- 305-CD-027 Release B SDPS Data Processing Subsystem Design Specification
- 305-CD-028 Release B CSMS Segment Communications Subsystem Design Specification
- 305-CD-029 Release B CSMS Segment Systems Management Subsystem Design Specification
- 305-CD-030 Release B GSFC Distributed Active Archive Center Design Specification
- 305-CD-031 Release B LaRC Distributed Active Archive Center Design Specification
- 305-CD-033 Release B EDC Distributed Active Archive Center Design Specification
- 305-CD-034 Release B ASF Data Center Distributed Active Archive Center Design Specification
- 305-CD-035 Release B NSIDC Distributed Active Archive Center Design Specification
- 305-CD-036 Release B JPL Distributed Active Archive Center Design Specification
- 305-CD-037 Release B ORNL Distributed Active Archive Center Design Specification
- 305-CD-038 Release B System Monitoring and Coordination Center Design Specification
- 305-CD-039 Release B Data Dictionary for Subsystem Design Specification

This document is a formal contract deliverable with an approval code of 2; as such it requires Government review and approval prior to acceptance and use. This document is under ECS contractor configuration control. Once this document is approved, Contractor approved changes are handled in accordance with Class I and Class II change control requirements described in the EOS Configuration Management Plan, and changes to this document shall be made by document change notice (DCN) or by complete revision.

Any questions should be addressed to:

Data Management Office
The ECS Project Office
Hughes Information Technology Systems
1616 McCormick Drive
Upper Marlboro, MD 20774-5372

Abstract

The Release-B Goddard Space Flight Center (GSFC) Distributed Active Archive Center (DAAC) Design Specification describes the ECS subsystems at the GSFC DAAC. ECS Subsystem-Specific Design Specifications provide detailed design descriptions of the subsystems. This document shows the specific implementation of that design at the GSFC DAAC, including the identification of the specific software, hardware and network configuration for the DAAC.

Keywords: GSFC, DAAC, configuration, design

This page intentionally left blank.

Change Information Page

List of Effective Pages			
Page Number	Issue		
Title iii through xii 1-1 through 1-4 2-1 through 2-4 3-1 through 3-128 4-1 and 4-2 A-1 through A-8 AB-1 through AB-8	Submitted as Final Submitted as Final		
Document History			
Document Number	Status/Issue	Publication Date	CCR Number
305-CD-030-001 305-CD-030-002	Preliminary Submitted as Final	October 1995 March 1996	95-0764 96-0235

This page intentionally left blank.

Contents

Preface

Abstract

1. Introduction

1.1	Identification	1-1
1.2	Scope.....	1-1
1.3	Purpose.....	1-2
1.4	Status and Schedule	1-2
1.5	Document Organization	1-2

2. Related Documentation

2.1	Parent Documents	2-1
2.2	Applicable Documents.....	2-2
2.3	Information Documents Not Referenced	2-3

3. GSFC ECS DAAC Configuration

3.1	Introduction.....	3-1
3.1.1	GSFC DAAC Overview.....	3-1
3.1.2	DAAC-Specific Mission and Operations Activities	3-4
3.2	GSFC External Interfaces	3-9
3.3	Computer Software Component Analysis.....	3-18
3.3.1	Software Subsystem Overview	3-18
3.3.2	Software Subsystem Analysis Summary	3-21

3.4	DAAC Hardware and Network Design	3-33
3.4.1	GSFC DAAC LAN Configuration.....	3-33
3.4.2	DAAC Hardware Configuration.....	3-37
3.5	Software/Hardware Mapping.....	3-115

4. Future Releases

Figures

3-1.1-1.	SDPS Subsystems and Components	3-3
3.1.1-2.	CSMS Subsystems and Components	3-4
3.2-1.	GSFC ECS DAAC External Interfaces.....	3-9
3.4.1-1.	GSFC DAAC LAN Topology.....	3-33
3.4.1-2.	ECS Network Interfaces at GSFC.....	3-35
3.4.2-1.	GSFC ECS DAAC Hardware Configuration Diagram.....	3-38
3.4.2-2.	DMGHW CI RAID Disk Configuration.....	3-64

Tables

3.1.1-1.	Release B Mission Baseline Mapped to GSFC DAAC	3
3.1.2-2.	GSFC SDPS Operations Support Functions	8
3.2-1.	GSFC External Interfaces	12
3.3.2-1.	GSFC DAAC Components Analysis	23
3.4.1-1.	Rel. B Networking Hardware for GSFC DAAC LAN	36
3.4.1.1-1.	Estimated Rel. B Data Flows for the GSFC DAAC	37
3.4.2.2-2.	Release A and B Datasets Held by the Data Server Subsystem	54
3.4.2.2-3.	GSFC DBMS Repository Sizing (GB)	56
3.4.2.3-1.	GSFC User System Accesses per Day for Science Users (Epoch m).....	60
3.4.2.3-2.	GSFC Science User Search Types for Gateway CI Services (Epoch m).....	60
3.4.2.3-3.	GSFC Searches per hour for Science Users (Epoch m).....	60
3.4.2.3-4.	GSFC User System Accesses per Day for Non-Science Users (Epoch m)	61

3.4.2.3-5.	GSFC Non-Science User Search Types for Gateway CI Service (Epoch m).....	61
3.4.2.3-6.	GSFC Science User Transaction Analysis for Gateway CI Service (Epoch m).....	62
3.4.2.3-7.	GSFC Science User Transaction Analysis for Advertising Service CI (Epoch m).....	62
3.4.2.3-8.	GSFC Non-Science User Transaction Analysis for Gateway.....	63
3.4.2.3-9.	GSFC Non-Science User Transaction Analysis for Advertising Service CI (Epoch m).....	63
3.4.2.3-10.	DBMS Transaction Analysis Summary (Epoch m).....	64
3.4.2.3-11.	DBMS Transaction Sensitivity Analysis Results (Epoch m) (loading has been doubled for both search invocations and transaction rates).....	64
3.4.2.3-12.	Vendor Platform Performance Estimates.....	65
3.4.2.3-13.	DMGHW CI Disk Capacity Requirements.....	66
3.4.2.3-14.	GSFC DAAC Hardware Configuration.....	69
3.4.2.4-1.	GSFC Ingest Volumes - Release B.....	70
3.4.2.4-2.	GSFC Ingest Components at Release B.....	71
3.4.2.4-3.	GSFC Ingest Sizing - Release B.....	72
3.4.2.6-1.	Breakdown of instrument production at GSFC.....	78
3.4.2.7.1-1.	GSFC Static Model Summary Results.....	85
3.4.2.7.1-2.	Comparison Of GSFC SPRHW Requirements And Specification.....	97
3.4.2.8.1-1.	GSFC Log Entry Storage Volume - Release B.....	109
3.4.2.8.1-2.	GSFC HP Openview Collection Processing Requirement.....	109
3.4.2.8.1-3.	MDA Data Conversion to Sybase Processing Requirement.....	109
3.4.2.8.1-4.	CSS/MSS Server Configuration - Requirements Estimate.....	111
3.4.2.8.1-5.	DDTS Benchmark Results.....	112
3.4.2.8.1-6.	MSS Workstation #1 Configuration - Requirements.....	114
3.4.2.8.1-7.	MSS Workstation #2 Configuration - Requirements.....	114
3.4.2.8.1-8.	COTS Product Disk Requirements.....	116
3.4.2.8.1-9.	GSFC MSS Release B Storage Requirements.....	117
3.5-1.	GSFC Hardware to Software Component Mapping.....	119

Appendix A

Abbreviations and Acronyms

1. Introduction

1.1 Identification

This Release B GSFC DAAC Design Specification for the ECS Project, Contract Data Requirement List (CDRL) Item 046, with requirements specified in Data Item Description (DID) 305/DV2, is a required deliverable under the Earth Observing System Data and Information System (EOSDIS) Core System (ECS), Contract NAS5-60000.

1.2 Scope

Release B of ECS supports functional capabilities and services required to meet driving requirements and milestones including:

- Functionality/services required to support mission operations for the continuation of TRMM, as well as the initiation of LANDSAT 7, COLOR, ADEOS II, and EOS AM-1. This includes planning and scheduling, command and control, production data processing, data distribution and other ECS functions.
- Functionality/services required to support mission operations for the initiation of SAGE III (METEOR) and ACRIM Flight-Of-Opportunity (FOO). This includes production data processing, data distribution and other ECS functions.
- Provide information management, data distribution and a high level archive for the SAR data from the ERS-1/2, JERS-1 and RADARSAT spacecraft.
- Functionality/services required to support EOS ground system interface testing which includes end-to-end mission simulations, communication services for EBnet, network management services and other ECS services.
- Functionality/services required for V0 Interoperability.
- Functionality/services required for Science Software I&T Support for TRMM, LANDSAT 7, COLOR, ADEOS II, EOS AM-1, SAGE III (METEOR), and ACRIM FOO.

Several of the driving requirements and milestones were initially supported by Release A but are expanded upon for Release B. For example, infrastructure Data Flow and End-to-End Testing and Simulation Readiness Testing are supported early-on by Release A, and will be fully supported by Release B during the final phases of testing. Likewise, V0 interoperability (one way) is supported by Release A for GSFC, LaRC, and EDC DAACs and is expanded to two way interoperability for all DAACs at Release B.

ECS will provide support to eight Distributed Active Archive Centers (DAACs). The DAACs are tasked with generating EOS standard data products and carrying out NASA's responsibilities for data archive, distribution and information management. The DAACs serve as the primary user interface to EOSDIS. These DAACs are located at: Goddard Space Flight Center (GSFC)

Greenbelt, MD; Langley Research Center (LaRC) Hampton, VA; Oak Ridge National Laboratory (ORNL) Oak Ridge, TN; EROS Data Center (EDC) Sioux Falls, SD; National Snow and Ice Data Center (NSIDC) Boulder, CO; Jet Propulsion Laboratory (JPL) Pasadena, CA; the Consortium for International Earth Science Information Network (CIESIN) in University Center, MI; and the Alaska SAR Facility (ASF) at the University of Alaska Fairbanks..

This document is part of a series of documents comprising the Science and Communications Development Office design specification for the Communications and System Management segment (CSMS) and the Science and Data Processing Subsystem (SDPS) for Release B. The series of documents include an overview, a design specification document for each subsystem, and a design implementation document for each DAAC involved in the release, as well as one for the System Monitoring and Control (SMC) center. A design specification for CIESIN, since the ECS contractor is not required to install, operate, or maintain hardware and software at CIESIN.

This document specifically focuses on the GSFC DAAC ECS configuration and capabilities at Release B. It is released, and reviewed at the formal Release B Critical Design Review (CDR). This document reflects the February 14, 1996 Technical Baseline, maintained by the ECS Configuration Control Board in accordance with ECS Technical Direction No. 11 dated December 6, 1994.

This document reflects the incorporation of elements previously identified as part of the Marshall Space Flight Center (MSFC) DAAC. Currently, change request

1.3 Purpose

The purpose of this document is to show the elements of the Release B ECS science data processing and communications design and implementation that will support the ECS portion of the GSFC DAAC in meeting its objectives. The Release B Overview of SDPS and CSMS (305-CD-020-002) provides an overview of the ECS subsystems and should be used by the reader in order to get a basic understanding of ECS design components. The Release Plan Content Description document (222-TP-003-008) provides a detailed mapping of functional capabilities and services that will be available for each release. While some DAAC configurations vary depending on the mission/capability requirements for ECS at their DAAC, the GSFC DAAC at full ECS capability will include all of the ECS science data processing and communications subsystems.

1.4 Status and Schedule

This submittal of DID 305/DV2 meets the milestone specified in the Contract Data Requirements List (CDRL) for Incremental Design Review (pre-CDR) of NASA Contract NAS5-60000. The submittal will be reviewed during the Release B (CDR) and changes to the design which resulted from that review will be reflected in subsequent updates.

1.5 Document Organization

This document is organized to describe the design of ECS portion of the GSFC DAAC as follows:

Section 1 provides information regarding the identification, scope, status and schedule, and organization of this document.

Section 2 provides a listing of the related documents which were used as source information for this document.

Section 3 provides a description of the ECS design at the GSFC DAAC. It includes a description of the DAAC external interfaces, ECS software implementation, including identification of Off the Shelf (OTS) products, hardware configuration and operational activities.

- Subsection 3.1 establishes the context for the technical discussions with an overview of the specific GSFC ECS DAAC mission and GSFC Release B operations. It identifies the key ECS related mission and operations activities that are supported via the ECS functionality at the DAAC.
- Subsection 3.2 addresses the external interfaces of the ECS subsystems as implemented at GSFC ECS DAAC. Major interfaces include TSDIS, users, Version 0 System IMS, ECS Release B DAACs, the MODIS Scientific Computing Facility (SCF), the EOS Operations Center (EOC), and the System Monitoring and Coordination Center (SMC).
- Subsection 3.3 provides a software component analysis. There are 10 ECS data processing and communications subsystems that contain Hardware Configuration Items (HWCI) and Computer Software Configuration Items (CSCI). This section addresses the CSCI and their corresponding lower level Computer Software Components (CSC). The CSCs are described in detail in their respective subsystem design specification documents. In this section, the CSCs are captured in a single table, broken down by Subsystem/CSCI. The table lists the CSCI and the associated CSCs. Notes are provided to expand upon generic explanations from the body of the Subsystem Design Specifications to describe what makes the particular CSC specific to the DAAC. In addition, when a CSC is identified as Off-the-shelf (OTS), the candidate product is identified, if known.
- Subsection 3.4 provides a DAAC specific discussion of the ECS data processing and communications Hardware Configuration Items (HWCI). This section identifies the HWCI components and indicates the specific components and quantities that are resident at the DAAC. It includes the Local area network (LAN) configuration and the rationale for the specific hardware configuration.
- Subsection 3.5 provides a software to hardware configuration mapping.

Section 4 gives a description of what can be expected in the next release of ECS.

Appendix A provides detailed configurations for the Data Processing Subsystem's Science Processing hardware suite.

The section, Abbreviations and Acronyms, contains an alphabetized list of the definitions for abbreviations and acronyms used in this document.

This page intentionally left blank.

2. Related Documentation

2.1 Parent Documents

The parent documents are the documents from which the scope and content of this Release B GSFC DAAC Design Specification is derived.

194-207-SE1-001	System Design Specification for the ECS Project
220-CD-001-004	Communications Requirements for the ECS Project, Revision 1
305-CD-002-002	Science Data Processing Segment (SDPS) Design Specification for the ECS Project
305-CD-003-002	Communications and System Management Segment (CSMS) Design Specification for the ECS Project
305-CD-020-002	Release B SDPS/CSMS Design Specification Overview for the ECS Project
305-CD-021-002	Release B SDPS Client Subsystem Design Specification
305-CD-022-002	Release B SDPS Interoperability Subsystem Design Specification
305-CD-023-002	Release B SDPS Data Management Subsystem Design Specification
305-CD-024-002	Release B SDPS Data Server Subsystem Design Specification
305-CD-025-002	Release B SDPS Ingest Subsystem Design Specification
305-CD-026-002	Release B SDPS Planning Subsystem Design Specification
305-CD-027-002	Release B SDPS Data Processing Subsystem Design for the ECS Project
305-CD-028-002	Release B CSMS Communications Subsystem Design Specification for the ECS Project
305-CD-029-002	Release B CSMS Systems Management Subsystem Design Specification
305-CD-039-002	Release B Data Dictionary for the ECS Project Subsystem Design Specification

2.2 Applicable Documents

The following documents are referenced within this Specification, or are directly applicable, or contain policies or other directive matters that are binding upon the content of this document.

206-CD-001-002	Version 0 Analysis Report for the ECS Project
209-CD-001-003	Interface Control Document Between EOSDIS Core System (ECS) and the NASA Science Internet
209-CD-008-004	Interface Control Document Between EOSDIS Core System (ECS) and the Goddard Space Flight Center (GSFC) Distributed Active Archive Center (DAAC)
209-CD-009-002	Interface Control Document Between EOSDIS Core System (ECS) and the Marshall Space Flight Center (MSFC) Distributed Active Archive Center (DAAC)
209-CD-010-002	Interface Control Document Between EOSDIS Core System (ECS) and the Langley Research Center (LaRC) Distributed Active Archive Center (DAAC)
209-CD-011-004	Interface Control Document Between EOSDIS Core System (ECS) and the Version 0 System
210-CD-001-003	Release B Risk Assessment Report for the ECS Project
305-CD-014-001	Release A GSFC DAAC Design Specification
305-CD-031-002	Release B LaRC Distributed Active Archive Center Design Specification
305-CD-032-001	Release B MSFC Distributed Active Archive Center Design Specification
305-CD-033-002	Release B EROS Data Center Distributed Active Archive Center Design Specification
305-CD-034-002	Release B ASF Distributed Active Archive Center Design Specification
305-CD-035-002	Release B NSIDC Distributed Active Archive Center Design Specification
305-CD-036-002	Release B JPL Distributed Active Archive Center Design Specification
305-CD-037-002	Release B ORNL Distributed Active Archive Center Design Specification
305-CD-038-002	Release B System Monitoring and Coordination Center Design Specification

313-CD-006-002	Release B CSMS/SDPS Internal Interface Control Document for the ECS Project
515-CD-001-003	Availability Model/Prediction for the ECS Project
518-CD-002-001	Release B Maintainability Predictions for the ECS Project
604-CD-001-004	Operations Concept for the ECS Project: Part 1-- ECS Overview
604-CD-002-003	Operations Concept for the ECS Project: Part 2B -- Release B
605-CD-002-001	Release B SDPS/CSMS Operations Scenarios
607-CD-001-002	Maintenance and Operations Manual for the ECS Project
622-CD-002-001	Release B Training Plan for the ECS Project
627-CD-001-002	Security Risk Management Plan for the ECS Project
160-TP-004-001	User Pull Analysis Notebook for the ECS Project
194-TP-312-001	User Characterization and Requirements Analysis
194-TP-569-001	PDPS Prototyping at ECS Science and Technology Laboratory, Progress Report #4
210-TP-001-004	Technical Baseline for the ECS Project, August 1995
210-TP-001-006	Technical Baseline for the ECS Project
222-TP-003-008	Release Plan Content Description for the ECS Project
420-TP-001-005	Proposed ECS Core Metadata Standard, Version 2.0
420-TP-010-001	Transition to Release B, Technical Paper
170-WP-005-001	DAS Integration in the ECS Environment
423-41-03	Goddard Space Flight Center, EOSDIS Core System (ECS) Contract Data Requirements Document

2.3 Information Documents Not Referenced

The following documents, although not referenced herein and/or not directly applicable, do amplify and clarify the information presented in this document. While not binding on the content of this design specification, these and many additional ECS documents are available via the EDHS. EDHS can be accessed via the World Wide Web (WWW) at the following Universal Reference Location: <http://edhs1.gsfc.nasa.gov/>. Please note that Internet links cannot be guaranteed for accuracy or currency.

302-CD-003-001	Release B Facilities Plan for the ECS Project
101-303-DV1-001	Individual Facility Requirements for the ECS Project

333-CD-003-002	SDP Toolkit Users Guide for the ECS Project
601-CD-001-004	Maintenance and Operations Management Plan for the ECS Project
101-620-OP2-001	List of Recommended Maintenance Equipment for the ECS Project
828-RD-001-002	Government Furnished Property for the ECS Project
430-TP-001-001	SDP Toolkit Implementation with Pathfinder SSM/I Precipitation Rate Algorithm, Technical Paper
440-TP-001-001	Science Data Server Architecture Study [for the ECS Project]
423-16-01	Goddard Space Flight Center, Data Production Software and Science Computing Facility (SCF) Standards and Guidelines
423-16-02	Goddard Space Flight Center, PGS Toolkit Requirements Specification for the ECS Project
423-41-02	Goddard Space Flight Center, Functional and Performance Requirements Specification for the Earth Observing System Data and Information System (EOSDIS) Core System
540-022	Goddard Space Flight Center, Earth Observing System (EOS) Communications (Ecom) System Design Specification
560-EDOS-0211.0001	Goddard Space Flight Center, Interface Requirements Document Between EDOS and the EOS Ground System (EGS)

3. GSFC ECS DAAC Configuration

3.1 Introduction

3.1.1 GSFC DAAC Overview

The GSFC Distributed Active Archive Center (GSFC DAAC) is one of the eight DAACs that are part of NASA's Earth Observing System Data and Information System (EOSDIS). These DAACs are generally organized to support specific scientific disciplines. The objective of the GSFC DAAC is to archive science data and provide support services to its users in the discipline areas of upper atmosphere, atmospheric dynamics, global biosphere, and geophysics. In the pre-EOS time frame, the GSFC V0 DAAC system has been developed, and is still being developed, to enhance and improve scientific research and productivity by consolidating access to remote sensor earth science data, as well as by providing services that provide added value to the data stored at the DAAC and help people realize the scientific and educational potential of the global climate data stored at the DAAC.

The GSFC DAAC, with its V0 System and operations staff, provides support for over 150 data sets/products which are developed from heritage data centers concerned with climate, land and ocean data. The DAAC also supports the initial NOAA Advanced Very High-Resolution Radiometer (AVHRR) and TOVS data used in the development efforts of the NASA Pathfinder data sets. The DAAC is also currently preparing to archive and distribute data received from the Sea Viewing Wide-Field-of-View Scanner (SeaWiFS) Project after satellite launch in 1996.

The GSFC DAAC V0 System is constantly adding new data sets, and will continue to archive and distribute a variety of Earth science products before the dedicated EOS AM satellite is launched in 1998. Consequently, for those services/functions not satisfied as part of ECS requirements, it is essential that the ECS design accommodates the GSFC DAAC current and future customer base, who will want the ability to make use of the set of V0 System services/functions currently available and any others developed before the time when ECS is fully operational at the DAAC.

With the advent of ECS, the GSFC DAAC will expand its ability to support its customers. Identified as a prototype system, the GSFC DAAC V0 System provided lessons learned for the ECS design and development activities. This will be reflected in the new capabilities/functions provided by ECS, as well as ECS's interface capabilities which will enable the GSFC DAAC to continue to provide the value added services/functions to its customers once the ECS hardware, software and operations requirements of ECS are fully integrated into the DAAC. In addition, the DAAC inherits the Data Assimilation Office (DAO) Data Assimilation System (DAS) which will be integrated in the ECS production system at the DAAC. The details associated with this

integration process are still being worked. This has been identified as an ECS risk area, thus it is discussed in the Risk Assessment Report for the ECS Project (210-CD-001-003). Reference the whitepaper which discusses the integration into the ECS environment (170-WP-005-001) for the current status of the DAS integration process.

As a result of the changes in the number of DAACs with ECS functionality, the GSFC DAAC will inherit V0 datasets and the support requirements for TRMM instruments previously allocated to the Marshall Space flight Center (MSFC) DAAC. There are still uncertainties associated with this allocation of MSFC data and functionality, consequently, this document will reflect the February 1996 ECS baseline.

Currently identified as the GSFC V0 DAAC, the GSFC DAAC will evolve as an operational DAAC which maintains and operates multiple systems, with the ECS representing the largest component of the DAAC. The GSFC DAAC can be characterized as a site whose growth and development coincide with that of the ECS. Thus, as the Earth Observing System Data Information System (EOSDIS) program prepares for its various missions, the nature of the DAAC changes. For the TRMM Mission, the GSFC DAAC is activated for an early release of ECS (Interim Release 1) which will support early TRMM interface testing. The software and equipment suites installed at the DAAC are then augmented for the implementation of Release A which supports the TRMM Mission, Landsat 7 early interface testing, EOS AM-1 interface testing, science software integration and test, data flow, end-to-end testing, V0 interoperability and simulation readiness testing. This document reflects the GSFC DAAC which will provide support for Release B of ECS.

Table 3.1.1-1 identifies the Release B mission baseline and indicates those missions for which the GSFC DAAC will provide support. Though not technically an EOS mission/instrument, the ECS baseline, for design purposes identifies DAO as a mission. The Release Content Plan document provides a description of the missions and the driving requirements which must be satisfied to support these missions. Figures 3.1.1-1 and 3.1.1-2 illustrate the Release B ECS SDPS and CSMS subsystems, respectively. Subsystem specific design specifications, as identified in Section 2.1 of this document, provide detailed information on the respective subsystems. All, or some portion of the hardware and software components will be integrated into one or more of the DAACs. This document focuses on the selected elements of the ECS design which represents the configuration allocated to the GSFC DAAC at Release B.

Table 3.1.1-1. Release B Mission Baseline Mapped to GSFC DAAC

Mission	Instruments	Launch Date	GSFC Support
TRMM	CERES, LIS, VIRS, PR, TMI	17 August, 1997 (initial support provided by ECS Release A)	yes (LIS, VIRS, PR, TMI)
EOS AM-1	ASTER, CERES, MISR, MODIS, MOPPIT	30 June, 1998	yes (MODIS)
Landsat 7	ETM+	May, 1998	no
FOO	COLOR	July, 1998	yes (under review)
ADEOS II	SeaWinds	February, 1999	no
ALT RADAR (CNES or GFO)	MR, DFA	March, 1999	no
FOO	ACRIM	June, 1999	no
METEOR	SAGE III	August, 1998	no
ERS-1	SAR	July, 1991	no
ERS-2	SAR	April, 1995	no
JERS-1	SAR	February, 1992	no
RADARSAT	SAR	October, 1995	no
DAO	DAS	1997	yes




Figure 3-1.1-1. SDPS Subsystems and Components




Figure 3.1.1-2. CSMS Subsystems and Components

3.1.2 DAAC-Specific Mission and Operations Activities

ECS subsystems provide mission and operations functionality for Release B. Key ECS related mission and operations activities supported by the ECS GSFC DAAC include:

- TRMM mission support
- EOS AM-1 MODIS support
- COLOR support
- DAO-DAS support
- Independent Verification and Validation (IV&V) support
- V0/ADC Interoperability
- Version 0 Data Migration
- DAAC Site Activation

TRMM Mission Support: TRMM (Tropical Rainfall Measurement Mission) is a platform scheduled for launch in August 1997 which relies on ECS to support its mission. The Level 0 data from three TRMM instruments (PR, TMI, and VIRS) will be higher level processed by TSDIS, a production system provided by the TRMM project. The GSFC DAAC will provide data archive, data search, order and distribution services to science users for information derived from these instruments. The interfaces are activated for ECS in Release A and are continued into Release B.

LIS mission support, including data ingest, archive and distribution, is allocated to the GSFC DAAC. Modifications to this Allocation are currently in progress.

EOS AM-1 MODIS Support: AM-1 is scheduled for launch in June 1998. The full set of ECS functionality must be operational to support AM-1 launch. This includes FOS planning, scheduling, command, control and monitoring of the AM-1 spacecraft; SDPS data ingest, production, archive, query and distribution; and CSMS system management and communications infrastructure. The GSFC DAAC provides science software integration and test support, data ingest, data processing, archiving and distribution support for the MODIS instrument.

Color Support: The Color platform is scheduled for launch in October 1998. ECS support for Color (currently under review) is assumed to be similar to the V0 DAAC support for SeaWiFS. Under this assumption, GSFC DAAC will be responsible for receiving higher level (level 1A and above) Color instrument data from a Color production facility. Presumably, this would include associated metadata and browse. GSFC DAAC will be responsible for archiving the data and providing data search, order and distribution services to authorized Color users. In SeaWiFS, the Version 0 GSFC DAAC also provides regular user access and distribution reports back to the production facility. This interface may also be required for Color. These interfaces must be operational in time to support the Color launch. Additionally, user data search, order and distribution services on Color data must be available by launch.

DAO-DAS Support: The GSFC DAAC will operate and maintain the Data Assimilation System of the Data Assimilation Office (DAO) located at GSFC. The DAS, currently and operational system, produces research-quality analyses for general Earth Science applications. The DAS processing environment at the GSFC DAAC will support operational assimilation, reanalysis, and research and development activities. The first two are analogous to first time processing and reprocessing, respectively, for other science algorithms. Research and development includes some of what is Algorithm Integration and Test for other algorithms. GSFC DAAC will house a dedicated equipment suite for the DAS. Categorized as production processes, operational assimilation and reanalysis will be performed using this equipment suite. Similarly, R&D work will also be run on this equipment. The GSFC DAAC will be responsible for ensuring that this does not interfere with operational processing.

The DAS will require NMC observational data and a selected set of data from EOS instruments. This input data will be made available to the DAS. The GSFC DAAC will ingest, archive and distribute products produced by the DAS.

Independent Verification and Validation (IV&V) Support: Prior to the Release Readiness Review (RRR), the IV&V contractor can witness and/or monitor release acceptance testing and document nonconformances. Upon successful completion of the RRR, the IV&V contractor verifies that the ECS release operates correctly within the EOS Ground System (EGS). The ECS contractor, specifically the Independent Acceptance Test Organization (IATO), supports the IV&V contractor in this effort for a period of one month following RRR at the operational sites. The IATO coordinates personnel, facilities, and equipment support in the resolution of ECS nonconformances identified during IV&V testing. The GSFC DAAC ECS Maintenance and Operations personnel will support IV&V activities at the DAAC, as necessary. Specific involvement is still being worked.

V0/ADC Interoperability: Two-way interoperability involves two different capabilities. First, outgoing interoperability allows users to log into the ECS and access ECS services, including the ability to access non-ECS data products from a site external to ECS directly from the ECS user interface. Second, incoming interoperability allows users, who are logged into a non-ECS site, to access ECS data products directly from the non-ECS user interface, using non-ECS IMS services.

Two-way interoperability is viable at the time ECS-unique data holdings are available. This does not occur until ECS integrates with TRMM and other platforms (Release A). Two-way Version 0 interoperability prior to the transition from Version 0 to ECS is required to ease the transition process. One-way interoperability with ADCs (ECS to ADCs) is also required early to ease the Version 0 transition.

Version 0 (V0) Data Migration: Version 0 (V0) data migration includes the ability to transition V0 data sets from V0 to V1; and provide support, data management, search, and access capabilities for these data sets. High priority data sets were defined to be candidate data sets for V0 transition during Release A. Additional data migration takes place during Release B operations.

DAAC Site Activation: The EOSDIS DAACs have the mission of processing, archiving and distributing earth science data. GSFC site activation was initially performed for IR-1, with phased activations required for Release A and Release B. These phased activations are characterized by increases in hardware and software functionality, including added operational responsibilities. This section discusses activities that will take place to facilitate a smooth transition from Release A to Release B.

The ECS contractor will schedule a series of site coordination trips to all DAACs. The objective of these trips is to ensure that the ECS contractor and the DAAC managers are in agreement with all operational issues. When ECS delivers Release B to the sites, ECS will work with the host organizations to ensure that hardware and software installation and segment and system testing all occur in a pre-planned manner that is sensitive to the mission of the host organization. Coordination topics include facility requirements, locations of Release B equipment and personnel, installation and test periods, etc.

The issues of when, and to whom, training is provided on Release B products are critical because of the potential impact on operations and user support. Training on COTS hardware and software, and application software, regardless of the development track, is an absolute necessity. If the site's user services are unable to handle issues about a Release B product, additional demands on developers' time will be made to isolate, remedy, or suggest work-arounds to the issues.

The facility access dates must be at least 2 months prior to the scheduled initial Release B installation date to provide time for site verification inspection, completion of Government facility preparations, and receiving of COTS HW and SW. Installations of HW and SW take between 2 and 6 weeks depending on whether the site is an initial installation (requiring LAN installation) and the quantity and complexity of the configurations to be installed.

After installation, staffing and training of the maintenance and operations (M&O) staff is accomplished. M&O training occurs in conjunction with the 3-month system integration and acceptance testing.

Another key objective is the transition of the Release A GSFC DAAC to Release B. This will involve installing and testing Release B custom software, COTS software and hardware in a currently operational environment. Reference the Transition to Release B Technical Paper (240-TP-010-001) for a more detailed discussion.

ECS subsystems provide mission and operations functionality for Release B. Key ECS-related mission and operations activities supported by the ECS portion of the GSFC DAAC include information management, data distribution and a high level data archive.

In addition to automated support, ECS subsystems provide the capability for the ECS operations staff to perform a number of roles in support of these activities. These operational roles for the GSFC DAAC are identified in Table 3.1.2-2. The table identifies the corresponding SDPS or CSMS subsystem that enables the DAAC ECS operations staff to perform a particular role/function. Detailed descriptions of these activities are captured in the ECS Operations Concept for the ECS Project: Part 2B - ECS Release B (604-CD-002-003) document. The Release B SDPS/CSMS Operations Scenarios document (605-CD-002-001) provides additional detailed scenarios for these activities. The Maintenance and Operations Manual for the ECS Project (607-CD-001-002) provides a detailed discussion of the specific operator positions and their corresponding roles and responsibilities.

Table 3.1.2-2. GSFC SDPS Operations Support Functions

ECS DAAC Operational Roles	ECS Capability
User Services - Support user with data expertise - Generate and maintain data interface	Data Management Subsystem
Data Ingest - Monitor electronic - Handle media	Ingest Subsystem
Production Planning	Planning Subsystem
Resource Planning	Planning Subsystem Systems Management Subsystem
Production Monitoring and control	Processing Subsystem
Archive Management	Data Server Subsystem
Data Distribution - Monitor electronic - Handle media	Data Server Subsystem
Resource Management - Fault Management - Baseline Management	Processing, Ingest, Distribution & Data Server Subsystems in coordination with Systems Management Subsystem
Algorithm Integration Support	Processing Subsystem
Database Maintenance	Data Management Subsystem Data Server Subsystem Application specific (1)
System and Performance Analysis	Systems Management Subsystem
Security Management	Systems Management Subsystem
Accounting and Billing	Systems Management Subsystem
Sustaining Engineering	Office Support Systems Management Subsystem Communication Subsystem
S/W and H/W Maintenance	Office Support Systems Management Subsystem Communication Subsystem
Configuration Management (chg control)	Systems Management Subsystem
Operations Management - Trouble Tickets - Ops policies and procedures	Systems Management Subsystem
Testing, training, property management, integrated logistics support, library administration	Office Support Systems Management Subsystem Communication Subsystem

Note:

1. Included to ensure that the number of small DBMSs throughout the system are not explicitly excluded (e.g., Planning Subsystem has a DBMS)

3.2 GSFC External Interfaces

The GSFC ECS DAAC will interface with multiple entities external to the DAAC. The ECS subsystem-specific DID305 design documents address the interfaces generically in a series of tables supported by textual explanations. For details, the reader is referred to those documents in addition to the various Interface Control Documents (ICDs). Figure 3.2-1 schematically illustrates the interfaces between the ECS subsystems at the GSFC DAAC and its external entities (sinks and sources of data). The figure enumerates data flows which are elaborated upon in Table 3.2-1.




Figure 3.2-1. GSFC ECS DAAC External Interfaces

The following further describes the external entities:

- FDF - The FDF interface provides AM-1 satellite repaired orbit/attitude data during anomalies. This data is required the processing of AM-1 instrument data.
- MODIS SCF - This interface supports the MODIS Instrument Team. MODIS algorithms, metadata, status, quality control products, standard products, calibration data, correlative data, algorithm updates and documentation are example of things that will cross this interface. All ECS-SCF interfaces involve humans at the SCF interacting with

ECS via ECS-provided (Client Subsystem), public domain, or COTS. The interface implementation between the SCF and DAAC operations staff include media, telephonic, WWW, electronic mail, and electronic file transfer. The Team Leader Computing Facility (TLCF) is at GSFC. Minimal communication between the investigators, at their SCF and the GSFC DAAC is expected to occur. The predominant interface will be between the MODIS Science Data Support Team (SDST) which operates in the MODIS TLCF and is the group responsible for providing integrated SCF science software to the DAAC.

- Other ECS DAACs - The GSFC ECS DAAC interfaces with other ECS DAACs to support cross DAAC access to all current data holdings. All data products migrated from the V0 DAACs will be available from these DAACs for access by users. Guide, inventory, standard products and other related information, identified in Table 3.2-1, will flow across this interface. In addition, this interface supports specific interfaces with LaRC, NSIDC and EDC to provide the following data:
 - LaRC: Ancillary data, including NMC, TOMS Ozone, Cloud Imager, and MODIS; VIRS data products.
 - EDC: MODIS Level 2
 - NSIDC: MODIS Level 2
- Version 0 System - This interface to GSFC ECS DAAC supports access to the V0 holdings that have not currently migrated to ECS. It is used to support the interoperability required for cross-DAAC access. The migration of selected datasets from the MSFC V0 DAAC into ECS will also occur via this interface.
- GSFC V0 DAAC - This interface provides access to data or other information that comes into the DAAC via the V0 IMS system but are archived into ECS, or into both ECS and the GSFC V0 archive. The migration of Version 0 datasets into ECS will occur via this interface. This interface also provides access to DAAC-unique components. GSFC candidates currently include:
 - SeaWiFs Interface
 - Anonymous FTP (data and contributed tools)
 - Browse-o-rama (a modified NCSA CGI program to allow users to view HDF files via a WWW interface)
 - CD-ROM Production
 - CD-ROM Ordering and Tracking
 - Elf (a WWW interface that provides search/order of the full inventory, structured access to anonymous FTP, access to server-side browse-o-rama, order form for CD-ROMs, and order form for other pre-packaged products)

- NOAA ADC - This interface supports access to non-EOS data sets to satisfy ECS user queries and to retrieve ancillary data for ECS standard product generation. This interface also supports the transfer of the NMC ancillary data set used to support TRMM and DAS processing.
- GSFC Users - This interface is the mechanism for user community access to ECS data and services. It is the mechanism by which advertisement, user registration, order and product status, desktop object manipulations, and command languages capabilities are utilized.
- EDOS- The science data stream downlinked from the EOS MODIS instruments is routed from EDOS to the GSFC DAAC after L0 processing. This data is ingested and made available for distribution to EDOS. Further information is available from the EDOS ICD.
- SMC - This interface provides the capability for the GSFC DAAC to receive performance information, processing status, scheduling and policy data and user registration information. Policy data includes that established by the ESDIS project. The GSFC DAAC sends system performance and status reports, as well as management summary data and fault and security event notifications to SMC as part of this interface.
- EOC - The EOS operation control center (EOC) interfaces with the GSFC DAAC for the exchange of instrument operations plans and schedules for the AM-1 spacecraft and mission historical information. Further information is available from the EOC ICD.

Note: The format of this data can be read only by EOC.

- TSDIS - This interface primarily supports the transfer of data between the GSFC DAAC and TSDIS for the purpose of archival of VIRS data sets including ancillary data, metadata and documentation. ECS ingests this data product and then on an as required basis, provides the TRMM data products back to TSDIS for reprocessing.
- Non-ECS Service Providers - This interface is required for specialized users who use ECS data to provide and advertise value-added services. These providers include commercial, institution, or other government agencies, as well as IPs, SCFs, and ADCs.
- GCMD - The Global Change Master Directory (GCMD) is a multidisciplinary database of information about data holdings of potential interest to the scientific research community. It contains high level descriptions of data set holdings of various agencies and institutions. It also contains supplementary descriptions about these data centers, as well as scientific campaigns and projects, sources (spacecraft, platforms), and sensors. This interface will allow the GSFC ECS DAAC to import directory level information from the GCMD via GCMD export files and generate ECS data product advertisements. A Data Interface Format (DIF) may be provided to GCMD to identify GSFC DAAC holdings (tbr). This interface is currently available to the V0 DAACs, but is not reflected in the ECS interface requirements.

SDPF - This interface transmits the Level 0 LIS data, metadata, predictive orbit, definitive orbit, backup, attitude, and expedited data. Further details on this interface can be found in the ICD between ECS and TSDIS (209-CD-007-003).

Table 3.2-1. GSFC External Interfaces (1 of 7)

Flow No	Source	Destination	Data Types	Data Volume	Frequency
1	FDF	Ingest	Repaired Orbit Data	low	as required
2	MODIS SCF	Ingest	Status	low	as required
2	MODIS SCF	Ingest	Metadata/updates	low	as required
2	MODIS SCF	Ingest	Documents	low	as required
2	MODIS SCF	Ingest	Algorithms/Updates	medium	as required
2	MODIS SCF	CSS (DAAC Ops via email)	Test Reviews by SCF	low	as required
2	MODIS SCF	CSS (DAAC Ops via email)	Request for Resource Usage	low	as required
2	MODIS SCF	CSS (DAAC Ops via email)	Reprocessing Request	low	as required
2	MODIS SCF	Data Server	QA Data request	low	as required
2	MODIS SCF	Data Server	QA Data Subscription	low	as required
3	Data Server	MODIS SCF	Status	low	as required
3	Data Server	MODIS SCF	Metadata/updates	low	as required
3	Data Server	MODIS SCF	Calibration data	medium	as required
3	Data Server	MODIS SCF	Correlative data	medium	as required
3	Data Server	MODIS SCF	Documents	low	as required
3	Data Server	MODIS SCF	Algorithms/updates	medium	as required
3	Data Server	MODIS SCF	Standard Products	medium	daily as required for QA
3	CSS (DAAC Ops via email, EDHS)	MODIS SCF	Toolkit Delivery and Update Package	low	as required
3	CSS (DAAC Ops via email, kftp)	MODIS SCF	Test Results, QA, and Production History Data	low	as required
3	CSS (DAAC Ops via email, kftp)	MODIS SCF	Resource Usage	low	as required
3	CSS (DAAC Ops via email, kftp)	MODIS SCF	Status	low	as required
4	Other ECS DAACs	Ingest	Ancillary Data	high	as required
4	Other ECS DAACs	Ingest	Correlative Data	high	as required

Table 3.2-1. GSFC External Interfaces (2 of 7)

Flow No	Source	Destination	Data Types	Data Volume	Frequency
4	Other ECS DAACs	Ingest	Calibration Data	medium	as required
4	Other ECS DAACs	Ingest	QA Data	medium	as required
4	Other ECS DAACs	Interoperability	Advertisements	medium	as required
4	Other ECS DAACs	Data Server	Result Sets	medium	as required
4	Other ECS DAACs	Client	Product Results	medium	as required
5	Data Server	Other ECS DAACs	Standard Products	high	as required
5	Data Server	Other ECS DAACs	Metadata	medium-high	as required
5	Data Server	Other ECS DAACs	Ancillary Data	high	as required
5	Data Server	Other ECS DAACs	Correlative Data	high	as required
5	Data Server	Other ECS DAACs	Calibration Data	high	as required
5	Data Server	Other ECS DAACs	Documents	medium	as required
5	Data Server	Other ECS DAACs	Orbit/Attitude Data	medium	as required
5	Data Server	Other ECS DAACs	Data Availability Schedules	medium	as required
5	Data Server	Other ECS DAACs	Algorithms	high	as required
5	Data Server	Other ECS DAACs	Special Products	high	as required
5	Data Server	Other ECS DAACs	L0 Data	high	as required
5	Data Server	Other ECS DAACs	Expedited or Validated Data	medium	as required
5	Data Server	Other ECS DAACs	QA Data	medium	as required
5	Data Server	Other ECS DAACs	Result Sets	medium	as required
5	Interoperability	Other ECS DAACs	Advertisements	medium	as required

Table 3.2-1. GSFC External Interfaces (3 of 7)

Flow No	Source	Destination	Data Types	Data Volume	Frequency
5	Client	Other ECS DAACs	Product Requests	medium	as required
6	Version 0 System	Data Server	Inventory	low	as required
6	Version 0 System	Data Server	Guide	low	as required
6	Version 0 System	Data Server	Browse data	medium	as required
6	Version 0 System	Data Server	Dependent Valid	low	as required
6	Version 0 System	Data Management	V0 Directory search request	low	as requested
6	Version 0 System	Data Management	V0 Inventory search request	low	as requested
6	Version 0 System	Data Management	V0 browse request	low	as requested
6	Version 0 System	Data Management	V0 product order request	low	frequency dependent on user input
6	Version 0 System	Ingest	Ancillary Data for TSDIS	low	as requested
6	Version 0 System	Ingest	Migration Data	high	varies depending on migration strategy
7	Data Mgmt	Version 0 System	V0 Browse Result	low-medium	in response to ECS browse result
7	Data Mgmt	Version 0	V0 inventory result set	low-high	in response to ECS inventory result request
7	Data Mgmt	Version 0	V0 directory search result set	low-high	in response to ECS request
7	Data Mgmt	Version 0 System	V0 product order response	low	in response to product request result
7	Data Server	Version 0 System	Result Sets	medium-high	in response to request
7	Data Server	Version 0 System	Session Mgmt responses	low	in response to request
7	Data Server	Version 0 System	Product Request Status	low	as required
8	GSFC V0 DAAC	Ingest	V0 Data Migration Data sets	medium	varies, depending on migration strategy
8	GSFC V0 DAAC	ECS	DAAC-Unique Component I/F	TBD	TBD

Table 3.2-1. GSFC External Interfaces (4 of 7)

Flow No	Source	Destination	Data Types	Data Volume	Frequency
8	GSFC V0 DAAC	Interoperability	Advertisements	low-medium	as required
8	GSFC V0 DAAC	Interoperability	Subscriptions	low	as required
9	Interoperability	GSFC V0 DAAC	Notifications	low	in response to subscriptions
10	ADC - NOAA	Ingest	NMC Ancillary Data	20 MB/day	frequency dependent on data set
11	Ingest	ADC-NOAA	protocol	low	as required
12	Users	Client	User registration information	low	as requested
12	Users	Client	User login information	low	as requested
12	Users	Client	Search requests	low	as requested
12	Users	Client	Product requests	low	as requested
12	Users	Client	Acquisition requests	low	as requested
12	Users	Client	Desktop manipulate commands	low	as supplied by user
12	Users	Client	Configuration/Profile information	low	as supplied by user
12	Users	Client	Data manipulate requests	low	as requested
12	Users	Client	Command language request	low	as requested
12	Users	Client	Advertisements, Software, and Documents	low	as supplied by user
12	Users	Ingest	User Methods	medium	as required
12	Users	Ingest	Ingest Status Requests	low	as required
13	Data Server	Users	Metadata	low	as requested
13	Data Server	Users	Documents	low	as requested
13	Data Server	Users	Data Products	medium	as requested
13	Data Server	Users	Browse Products	medium	as required
13	Data Server	Users	Product Request Status	low	as requested
13	Data Server	Users	DAR Status	low	as requested
13	Data Server	Users	Schedules	low	as requested
13	Client	Users	Results Set	medium	as requested
13	Client	Users	Application user interfaces	low	as requested

Table 3.2-1. GSFC External Interfaces (5 of 7)

Flow No	Source	Destination	Data Types	Data Volume	Frequency
13	Client	Users	Formatted data	medium	as requested
13	Client	Users	Desktop Objects	low	as requested
13	Client	Users	Advertisement and Software	low	as requested
13	Client	Users	Error and Status information	low	as available
13	Ingest	Users	Ingest Status	low	as requested
14	Ingest	EDOS	Service Requests (Back-up data requests)	low	rare
14	Ingest	EDOS	Fault Report	low	rare
14	Ingest	EDOS	Fault Isolation Request	low	depends on EDOS
14	Ingest	EDOS	L0 data	high	rare
15	EDOS	Ingest	Service Request Disposition	low	as required
15	EDOS	Ingest	PDSs (L0 Data)	high	several times a day
15	EDOS	Ingest	ADSs (Back-up L0 Data)	high	as required
15	EDOS	Ingest	PDS Delivery Record	low	several times a day
15	EDOS	Ingest	ADS Delivery Record	low	as required
15	EDOS	Ingest	Physical Media Unit Delivery Record	low	as required
15	EDOS	Ingest	Undetected Fault Isolation	low	as required
16	SMC	MSS	Policies	low	as required
16	SMC	MSS	Conflict Resolution	low	as required
16	SMC	MSS	Procedures	low	as required
16	SMC	MSS	Directives	low	as required
17	MSS	SMC	Conflict Resolution Request	low	as required
17	MSS	SMC	Status	low	as required
17	MSS	SMC	Performance	low	as required
17	MSS	SMC	Management Summary Data	low	as required
17	MSS	SMC	Security & fault event notification	low	as required
18	EOC	Ingest	Telemetry Data	173 MB/day	Twice per day (12 files each transfer)

Table 3.2-1. GSFC External Interfaces (6 of 7)

Flow No	Source	Destination	Data Types	Data Volume	Frequency
18	EOC	Ingest	Event	12MB/day	Every hour
18	EOC	Ingest	EOC statistics, schedules, reports, etc.	25MB/day	Twice per day (25 files total each day)
19	TSDIS	Ingest	Metadata	1 days worth of products	daily
19	TSDIS	Ingest	Ancillary Data (GV)	1 days worth	daily
19	TSDIS	Ingest	Data Products (VIRS Levels 1A,1B)	4.2 GB/day Processing and Reprocessing	
19	TSDIS	Ingest	Data Products (TMI,PR, combined) Levels 1A,1B,2A,2B, 3A, and 3B	60 (GB/day) Processing and Reprocessing	
19	TSDIS	Ingest	Algorithms	medium	as required
19	TSDIS	Ingest	Documents	low	as required
19	TSDIS	Ingest	Status	low	as required
19	TSDIS	Ingest	Browse Data	high	daily
19	TSDIS	Ingest	Directory	low	as required
19	TSDIS	Ingest	Guide	low	as required
19	TSDIS	Ingest	Schedule Data	low	as required
19	TSDIS	Ingest	Platform Ephemeris		
19	TSDIS	Data Server	Subscription Data	low	as required
20	Data Server	TSDIS	Platform Ephemeris		
20	Data Server	TSDIS	Ancillary Data (NMC)		as required to support reprocessing
20	Data Server	TSDIS	Data Products (VIRS Levels 1A,1B)		as required to support reprocessing
21	Interoperability	Non-ECS Service Providers	Notifications	low	in response to subscriptions
22	Non-ECS Service Providers	Interoperability	Advertisements	low-medium	as required
22	Non-ECS Service Providers	Interoperability	Subscriptions	low	as required
23	GCMD	Interoperability	Advertisements	low-medium	as required

Table 3.2-1. GSFC External Interfaces (7 of 7)

Flow No	Source	Destination	Data Types	Data Volume	Frequency
24	Ingest	SDPF	L0 data	medium	as requested for reprocessing
25	SDPF	Ingest	L0 Data	Instrument Specific (CERES, LIS)	daily
25	SDPF	Ingest	Expedited Data	medium	three times a day
25	SDPF	Ingest	Predictive Orbit Data	medium	daily
25	SDPF	Ingest	Definitive Orbit Data	medium	daily
25	SDPF	Ingest	Back-up Data	medium	as required

In the table, where an exact number is unavailable, the data volume is estimated as low (less than 1 MB), medium (between 1 MB and 1 GB), or high (greater than 1 GB) per use defined in the frequency column .

3.3 Computer Software Component Analysis

3.3.1 Software Subsystem Overview

The 10 ECS software subsystems are described in detail in the ECS Subsystem-specific DID305 documents. This section provides a brief overview description of each of the subsystems.

Client Subsystem (CLS): This software consists of graphic user interface (GUI) programs, tools for viewing and/or manipulating the various kinds of ECS data (e.g., images, documents, tables) and libraries representing the client application program interface (API) of ECS services. The Release A V0 System Client remains as part of this subsystem. The client subsystem components will be available to users for installation on their workstations and will also be deployed on workstations within the DAAC in support of normal operations, including User Services support.

Interoperability Subsystem (IOS): The interoperability subsystem is an advertising service. It maintains a database of information about the services and data offered by ECS, and allows users to search through this database to locate services and data that may be of interest to them. The advertising service will be implemented as a Web server application with a DBMS back-end.

Data Management Subsystem (DMS): This subsystem includes functions which provide uniform access to descriptions of the data and the data elements offered by the ECS repositories and provide a bi-directional gateway between ECS and Version 0. This subsystem also includes distributed search and retrieval functions and corresponding site interfaces.

Data Server Subsystem (DSS): The subsystem provides the physical storage access and management functions for the ECS earth science data repositories. Other subsystems can access it directly or via the data management subsystem (if they need assistance with searches across several of these repositories). The subsystem also includes the capabilities needed to distribute bulk data via electronic file transfer or physical media. Other components include, for example, administrative software to manage the subsystem resources and perform data administration functions (e.g., to maintain the database schema); and data distribution software, e.g., for media handling and format conversions. The main components of the subsystem are the following:

- database management system - uses an off-the-shelf DBMS (Illustra) to manage its earth science data and implement spatial searching, as well as for the more traditional types of data (e.g., system administrative and operational data). It will use a document management system to provide storage and information retrieval for guide documents, scientific articles, and other types of document data.
- file storage management systems - used to provide archival and staging storage for large volumes of data. Provides hierarchical storage support and device/media independence to the remainder of DSS and ECS.
- data type libraries - they will implement functionality of earth science and related data that is unique and not available off-the-shelf (e.g., spatial search algorithms and translations among coordinate systems). The libraries will interface with the data storage facilities, i.e. the database and file storage management systems.

Ingest Subsystem (INS): The subsystem deals with the initial reception of all data received at an ECS facility and triggers subsequent archiving and processing of the data. Given the variety of possible data formats and structures, each external interface, and each ad-hoc ingest task may have unique aspects. Therefore, the ingest subsystem is organized into a collection of software components (e.g., ingest management software, translation tools, media handling software) from which those required in a specific situation can be readily configured. The resultant configuration is called an ingest client. Ingest clients can operate on a continuous basis to serve a routine external interface or they may exist only for the duration of a specific ad-hoc ingest task.

Data Processing Subsystem (DPS): The main components of the data processing subsystem - the science software - will be provided by the science teams. The data processing subsystem will provide the necessary hardware resources, as well as software for queuing, dispatching and managing the execution of the science software in an environment which will eventually be highly distributed and consist of heterogeneous computing platforms. The DPS also interacts with the DSS to cause the staging and de-staging of data resources in synchronization with processing requirements.

Planning Subsystem (PLS): This subsystem provides the functions needed to pre-plan routine data processing, schedule on-demand processing, and dispatch and manage processing requests. The subsystem provides access to the data production schedules at each site, and provides management functions for handling deviations from the schedule to operations and science users.

System Management Subsystem (MSS): The Management Subsystem (MSS) provides enterprise management (network, system and application management) for all ECS resources: commercial hardware (including computers, peripherals, and network routing devices), commercial software, and custom applications. Enterprise management reduces overall development and equipment costs, improves operational robustness, and promotes compatibility with evolving industry and government standards. Consistent with current trends in industry, the MSS thus manages both ECS's network resources per EBNET requirements and ECS's host/application resources per SMC requirements. Additionally MSS also supports many requirements allocated to SDPS and FOS for management data collection and analysis/distribution.

The MSS allocates services to both the system-wide and local levels. With few exceptions, the management services will be fully decentralized and no single point of failure exists which would preclude user access. In principle, every service is distributed unless there is an overriding reason for it to be centralized. MSS has two key specializations: Enterprise Monitoring and Coordination Services, and Local System Management Services.

Communications Subsystem (CSS): The CSS services include Object Services, Distributed Object Framework (DOF) and Common Facility Services. Support in this subsystem area is provided for peer-to-peer, advanced distributed, messaging, management, and event-handling communications facilities. These services typically appear on communicating end-systems across an internetwork and are not layered, but hierarchical in nature. Additionally, services to support communicating entities are provided, including directory, security, time, and other ancillary services. The services of the Communications Subsystem are functionally dependent on the services of the Internetworking Subsystem. The services of the common facility, object and DOF are the fundamental set of interfaces for all management and user access (i.e., pull) domain services.

Internetworking Subsystem (ISS): The Internetworking Subsystem provides for the transfer of data transparently within the DAACs, SMC and EOC, and for providing interfaces between these components and external networks. ECS interfaces with external systems and DAAC-to-DAAC communications are provided by the EOSDIS Backbone Network (EBnet). EBnet's primary function is to transfer data between DAACs, including both product data and inter-DAAC queries and metadata responses. Other networks, such as NSI, will provide wide-area services to ECS. In addition, "Campus" networks, which form the existing networking infrastructure at the ECS locations, will provide connectivity to EOSDIS components such as SCFs and ISTs.

3.3.2 Software Subsystem Analysis Summary

This summary addresses the CSCIs for each subsystem, focusing upon those CSCIs that are specific to the GSFC ECS DAAC. For the most part, the version of software is the same for all ECS DAACs. When differences occur, they are due more to things like the content of databases and schema constructions than to software. The following addresses each subsystem in a somewhat general manner, to point out whether or not there are any GSFC DAAC specific portions.

- Client Subsystem–The client software will have no GSFC DAAC specific portions. This software supports the GSFC ECS DAAC M&O staff (e.g. operations, user services, system administrators) and will be hosted on ECS workstations defined within the system. The Client software may be hosted on existing DAAC workstations to provide additional user access, but the GSFC DAAC will be responsible for providing additional Client host workstations when required by the DAAC's user community.
- Data Server -The software components of the Data Server Subsystem are largely the same for all Data Servers, at all DAACs. The two basic areas in which the Data Server Subsystem software will vary from DAAC to DAAC are configuration and special components.

Data Server software is designed to be highly configurable in order to allow a wide variety of DAAC unique policy implementations. These unique configurations will enable the data server software installations to vary behavior, meeting the DAAC specific needs. Examples of configuration parameters include number of concurrent connections, number of requests per client, inactivity timeout period and allocation of software components to hardware.

Another facet of the Data Server Subsystem software that supports the specific DAAC capabilities is in which actual components are installed at the ECS portion of the DAAC. These opportunities for DAAC specificity are driven by the types of distribution available to the DAACs data server clients and in the types of data (and their data type services) available. There will be portions of the Data Server software specific to the GSFC DAAC that are used to add special distribution devices (e.g. - 3490Tape).

However, the primary portion of the Data Server Subsystem software that will be specific to the GSFC DAAC will be the specific data types supported at the DAAC. These software components are a portion of the Science Data Server (SDSRV) CSCI. The SDSRV is designed to allow complete flexibility in the data types (specifically, Earth Science Data Types, or ESDTs) that offer their services via the Data Server. Examples of services offered by ESDTs include Insert, Acquire, Browse and SpatialSubset. These data types are organized by separate CSCs, generally one per source instrument (i.e., MODIS, DAO, PRI, TMI, VIRS, LIS). Additional CSCs may be required for V0 migrated data.

- Data Management –None of the data management software will be specific to the GSFC DAAC. The V0 Gateway (GTWAY) CSCI will interface with the data servers at each site.

- Ingest–The software portions for ingest at GSFC may differ from those of other DAACs because of dataset dependencies. Data ingestion procedures must match the peculiarities of the ingested data sets.
- Interoperability–There are no GSFC DAAC specific portions of the Interoperability Subsystem.
- Production Planning–There are currently no GSFC DAAC specific portions of the production Planning subsystem. However, the resolution of the DAO-DAS integration issue may result in specific functions at the GSFC DAAC. The configuration and database information will be specific to the DAAC and the production plans developed for each DAAC will be different. Script files, while they may use the same language, will differ for things like fault recovery, fault notification, production management, and exception handling.
- Data Processing–Due to dataset characteristics there may be some software unique to GSFC in the area of Science Data Processing (e.g., compilers for the Science software) and Pre-processing (e.g., DAAC-specific external interface EDOS).
- Communications Subsystem - There are no GSFC ECS DAAC specific portions of this subsystem.
- Systems Management - This subsystem is composed of a variety of management applications, providing services such as fault, performance, security and accountability management for ECS networks, hosts, and applications. Two tiers of "view" (domain of management service interface) provided by the applications in this subsystem. Only the local management view is provided at the GSFC ECS DAAC. There are no GSFC ECS DAAC specific portions of this subsystem.
- Internetworking Subsystem - There are no GSFC ECS DAAC specific portions of this subsystem.

Table 3.3.2-1, lists the ECS subsystems and associated CSCIs and CSCs. For each CSC, there is an indication of the type of component. As defined in the DID 305 subsystem-specific documents, type indicates whether the component is custom developed (DEV), off the shelf (OTS), a CSC reused from another subsystem (reuse), a wrapper (WRP), or a combination of these types. The Use column indicates whether a generic (Gnrc) form of the CSC is implemented or specific (Spfc) tailoring or use is required at a DAAC. The Notes column is included to comment about the characteristics of the system, data, and/or software that makes the CSC specific, as well as to provide any additional information about the generic CSCs. This column also identifies the OTS product.

Table 3.3.2-1. GSFC DAAC Components Analysis (1 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Client	DESKT	Desktop Manager	DEV	Gnrc	
Client	WKBCH	Comment/Survey Tool	OTS/ DEV	Gnrc	WWW Browser
Client	WKBCH	Data Acquisition Request Tool	DEV	Gnrc	
Client	WKBCH	Data Dictionary Tool	DEV	Gnrc	
Client	WKBCH	Document Search Tool	OTS	Reuse	CSS-provided
Client	WKBCH	Earth Science Search Tool	DEV	Gnrc	
Client	WKBCH	E-mailer Tool	OTS	Reuse	CSS-provided
Client	WKBCH	Hypertext Authoring Tool	OTS	Gnrc	MS Office / TBD public domain
Client	WKBCH	Hypertext Viewer	OTS	Gnrc	WWW Browser
Client	WKBCH	Logger/Reviewer Tool	DEV	Gnrc	
Client	WKBCH	News Reader Tool	OTS	Reuse	CSS-provided
Client	WKBCH	Product Request Tool	DEV	Gnrc	
Client	WKBCH	Session Management Tool	DEV	Gnrc	
Client	WKBCH	User Preferences Tool	DEV	Gnrc	
Client	WKBCH	User Registration Tool	DEV	Gnrc	
Client	WKBCH	Visualization Tool	DEV	Gnrc	
Communication	DCCI	Bulletin Board	OTS	Reuse	CSS-provided
Communication	DCCI	Directory/Naming Services	OTS/ DEV	Gnrc	OODCE
Communication	DCCI	Distributed File System (DFS)	OTS	Gnrc	DCE
Communication	DCCI	DOF Services	OTS	Gnrc	OODCE
Communication	DCCI	Electronic Mail Services	OTS/ DEV	Gnrc	native operating system
Communication	DCCI	Event Logger Services	OTS/ DEV	Gnrc	DCE
Communication	DCCI	File Access Services	OTS/ DEV	Gnrc	ftp, kftp, DCE
Communication	DCCI	Life Cycle Services	OTS/ DEV	Gnrc	OODCE
Communication	DCCI	Message Passing Services	OTS/ DEV	Gnrc	Developed with OODCE

Table 3.3.2-1. GSFC DAAC Components Analysis (2 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Communication	DCCI	Security Services	OTS/ DEV	Gnrc	OODCE
Communication	DCCI	Thread Services	OTS	Gnrc	OODCE
Communication	DCCI	Time Services	OTS/ DEV	Gnrc	OODCE
Communication	DCCI	Virtual Terminal Services	OTS	Gnrc	native operating system
Data Management	DDICT	Client Library	DEV	Gnrc	
Data Management	DDICT	Configuration/Setup	DEV	Gnrc	
Data Management	DDICT	DBMS Server	OTS	Gnrc	Sybase DBMS
Data Management	DDICT	Maintenance Tool	DEV	Gnrc	
Data Management	DDICT	Persistent Data	DEV	Gnrc	
Data Management	DDICT	Request Processing	DEV	Gnrc	
Data Management	DDICT	Server	DEV	Gnrc	
Data Management	DIMGR	Configuration/Setup	DEV	Gnrc	
Data Management	DIMGR	Server	DEV	Gnrc	
Data Management	GTWAY	Configuration/Setup	DEV	Gnrc	
Data Management	GTWAY	Server	DEV	Gnrc	
Data Management	GTWAY	V0 Back End	OTS	Gnrc	From V0
Data Management	GTWAY	V0 Client Interface	DEV	Gnrc	
Data Management	GTWAY	V0 External Interface	DEV	Gnrc	
Data Management	GTWAY	V0 Front End	OTS	Gnrc	From V0
Data Management	LIMGR	Client Library	DEV	Gnrc	
Data Management	LIMGR	Configuration/Setup	DEV	Gnrc	
Data Management	LIMGR	Database Interface	OTS	Gnrc	RogueWave DBTools
Data Management	LIMGR	External Interface	DEV	Gnrc	
Data Management	LIMGR	Mapping Layer	DEV	Gnrc	
Data Management	LIMGR	Request Processing	DEV	Gnrc	
Data Management	LIMGR	Server	DEV	Gnrc	
Data Processing	AITTL	Binary File Comparison Utility	DEV	Gnrc	
Data Processing	AITTL	Code Analysis Tools	OTS	Spfc	CASEVision SPARCWorks Specific for science software language.
Data Processing	AITTL	Data Visualization Tools	OTS DEV	Gnrc	IDL
Data Processing	AITTL	Documentation Viewing Tools	OTS	Gnrc	SoftWindows/MS Office Ghostview

Table 3.3.2-1. GSFC DAAC Components Analysis (3 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Data Processing	AITTL	ECS HDF Visualization Tools	DEV	Gnrc	Reused from Client subsystem, WKBCH CSCI, Data Visualization (EOSView) CSC
Data Processing	AITTL	HDF File Comparison Utility	DEV OTS	Gnrc	Custom IDL program
Data Processing	AITTL	PGE Processing GUI	DEV	Gnrc	
Data Processing	AITTL	PGE Registration GUI	DEV	Gnrc	
Data Processing	AITTL	Product Metadata Display Tool	Reuse DEV	Gnrc	Reused from Data Processing subsystem, AITTL CSCI, HDF File Comparison Utility CSC
Data Processing	AITTL	Profiling Tools	OTS	Spfc	CASEVision Specific for science software language.
Data Processing	AITTL	Report Generation Tools	OTS/ DEV	Gnrc	OTS: SoftWindows/MS Office, Dev: SSI&T manager
Data Processing	AITTL	SDP Toolkit-related Tools	DEV	Gnrc	
Data Processing	AITTL	SSAP Processing GUI	DEV	Gnrc	
Data Processing	AITTL	Standards Checkers	OTS/ DEV	Spfc	FORCHECK for Fortran 77; otherwise, native compilers and UNIX lint. Specific for science software language.
Data Processing	AITTL	Update Data Server GUI	DEV	Gnrc	
Data Processing	AITTL	Update PGE Database GUI	DEV	Gnrc	
Data Processing	PRONG	COTS	OTS	Gnrc	AutoSys and AutoXpert.
Data Processing	PRONG	COTS Management	DEV	Gnrc	
Data Processing	PRONG	Data Management	DEV	Gnrc	
Data Processing	PRONG	Data Pre-Processing	DEV	Spfc	Specific based on uniqueness of ancillary data products
Data Processing	PRONG	PGE Execution Management	DEV	Gnrc	
Data Processing	PRONG	Quality Assurance Monitor Interface	DEV/ OTS	Spfc	Specific based on uniqueness of Science software and SCF- DAAC relationship
Data Processing	PRONG	Resource Management	DEV	Gnrc	
Data Processing	SDPTK	Ancillary Data Access	DEV	Gnrc	

Table 3.3.2-1. GSFC DAAC Components Analysis (4 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Data Processing	SDPTK	Celestial Body Position	DEV	Gnrc	
Data Processing	SDPTK	Constant and Unit Conversions	DEV/OTS	Gnrc	UDUNITS freeware
Data Processing	SDPTK	Coordinate System Conversion	DEV	Gnrc	
Data Processing	SDPTK	EOS-HDF	DEV	Gnrc	
Data Processing	SDPTK	Error/Status Handling	DEV	Gnrc	
Data Processing	SDPTK	Geo Coordinate Transformation	DEV/OTS	Gnrc	GCTP (USGS)
Data Processing	SDPTK	Graphics Library	OTS	Gnrc	IDL
Data Processing	SDPTK	File Input/Output Tools	DEV/OTS	Gnrc	NCSA HDF
Data Processing	SDPTK	Math Package (IMSL)	OTS	Gnrc	IMSL
Data Processing	SDPTK	Memory Management	DEV	Gnrc	
Data Processing	SDPTK	Metadata Access	DEV	Gnrc	
Data Processing	SDPTK	Process Control	DEV	Gnrc	
Data Processing	SDPTK	Spacecraft Ephemeris and Attitude Access	DEV	Gnrc	
Data Processing	SDPTK	Time Date Conversion	DEV	Gnrc	
Data Server	DDIST	Distribution Client Interface	DEV	Gnrc	
Data Server	DDIST	Distribution Products	DEV	Gnrc	
Data Server	DDIST	Distribution Request Management	DEV	Gnrc	
Data Server	DDSRV	DDSRV	DEV/OTS	Gnrc	RogueWave class libraries
Data Server	DDSRV	DDSRV Client	DEV/OTS	Gnrc	OODCE; RogueWave class libraries
Data Server	DDSRV	DDSRV CSDT	DEV/OTS	Gnrc	RogueWave class libraries
Data Server	DDSRV	DDSRV ESDT	DEV/OTS/Reuse	Gnrc	Reuse from SDSRV; RogueWave class libraries
Data Server	DDSRV	DDSRV Search Engine	DEV/OTS	Gnrc	Topic and Netscape server
Data Server	DDSRV	DDSRV Server	DEV	Gnrc	Topic API; RogueWave class libraries
Data Server	DDSRV	Gateway	DEV/OTS	Gnrc	Netscape libraries

Table 3.3.2-1. GSFC DAAC Components Analysis (5 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Data Server	SDSRV	Administration and Operations	DEV	Gnrc	RogueWave class libraries
Data Server	SDSRV	Client	DEV/OTS	Gnrc	OODCE; RogueWave class libraries
Data Server	SDSRV	Configuration and Startup	DEV	Gnrc	RogueWave class libraries
Data Server	SDSRV	CSDT	DEV/OTS	Gnrc	HDF-EOS
Data Server	SDSRV	DB Wrappers	DEV/OTS	Gnrc	Illustra DBMS server and API
Data Server	SDSRV	Descriptors	DEV	Gnrc	RogueWave class libraries
Data Server	SDSRV	General ESDT	DEV	Gnrc	RogueWave class libraries
Data Server	SDSRV	Global	DEV/OTS	Gnrc	RogueWave class libraries
Data Server	SDSRV	GUI	DEV/OTS	Gnrc	RogueWave class libraries; X11/Motif
Data Server	SDSRV	Metadata	DEV/Wrpr	Gnrc	Illustra DBMS API
Data Server	SDSRV	Non-Product Science ESDT	DEV	Gnrc	RogueWave class libraries
Data Server	SDSRV	Non-Science ESDT	DEV	Gnrc	RogueWave class libraries
Data Server	SDSRV	MODIS	DEV	Spfc	ESDTs - MODIS
Data Server	SDSRV	VIRS	DEV	Spfc	ESDTs - VIRS
Data Server	SDSRV	PR	DEV	Spfc	ESDTs - PR
Data Server	SDSRV	TMI	DEV	Spfc	ESDTs - TMI
Data Server	SDSRV	LIS	DEV	Spfc	ESDTs - LIS
Data Server	SDSRV	DAS	DEV	Spfc	ESDTs - DAS
Data Server	SDSRV	Server	DEV/OTS	Gnrc	RogueWave class libraries
Data Server	SDSRV	Subscriptions	DEV/OTS	Gnrc	RogueWave class libraries
Data Server	STMGT	Data Storage	DEV/OTS	Gnrc	AMASS File Storage Management System
Data Server	STMGT	File	DEV	Gnrc	
Data Server	STMGT	Peripherals	DEV	Gnrc	This CSC encapsulates the CSS-supplied API which supports the OTS FTP product.

Table 3.3.2-1. GSFC DAAC Components Analysis (6 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Data Server	STMGT	Resource Management	DEV	Gnrc	
Data Server	STMGT	Service Clients	DEV/OTS	Gnrc	CSC encapsulates the AMASS OTS product (Data Server subsystem, STMGT CSCI, Data Storage CSC) . Also Rogue Wave class libraries.
Ingest	INGST	Client	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Configuration/Startup	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	CSDT	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Data Storage	Reuse	Gnrc	Reused from Data Server subsystem, STMGT CSCI
Ingest	INGST	DB Wrappers	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Descriptors	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Distribution Client Interface	Reuse	Gnrc	Reused from Data Server subsystem, DDIST CSCI
Ingest	INGST	Distribution Products	Reuse	Gnrc	Reused from Data Server subsystem, DDIST CSCI
Ingest	INGST	Distribution Request Management	Reuse	Gnrc	Reused from Data Server subsystem, DDIST CSCI
Ingest	INGST	File	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	General ESDT	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Global	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	GUI	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Ingest Administration Data	DEV	Gnrc	
Ingest	INGST	Ingest Data Preprocessing	DEV	Spfc	Specific based on uniqueness of ingested data and preprocessing requirements.
Ingest	INGST	Ingest Data Transfer	DEV	Gnrc	
Ingest	INGST	Ingest DBMS	OTS	Gnrc	Sybase DBMS

Table 3.3.2-1. GSFC DAAC Components Analysis (7 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Ingest	INGST	Ingest Request Processing	DEV	Gnrc	
Ingest	INGST	Ingest Session Manager	DEV	Gnrc	
Ingest	INGST	Metadata	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Non-Product Science ESDTs	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Non-Science ESDTs	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	MODIS	Reuse	Spfc	ESDTs - MODIS; Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	VIRS	Reuse	Spfc	ESDTs - VIRS; Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	PR	Reuse	Spfc	ESDTs - PR; Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	TMI	Reuse	Spfc	ESDTs - TMI; Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	LIS	Reuse	Spfc	ESDTs - LIS; Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	DAS	Reuse	Spfc	ESDTs - DAS; Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Operator Ingest Interfaces	DEV	Gnrc	
Ingest	INGST	Polling Ingest Client Interface	DEV	Gnrc	
Ingest	INGST	Resource Management	Reuse	Gnrc	Reused from Data Server subsystem, STMGT CSCI
Ingest	INGST	Server	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Service Clients	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI
Ingest	INGST	Subscriptions	Reuse	Gnrc	Reused from Data Server subsystem, SDSRV CSCI

Table 3.3.2-1. GSFC DAAC Components Analysis (8 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Ingest	INGST	User Network Ingest Interface	DEV	Gnrc	
Ingest	INGST	Viewing Tools	Reuse	Gnrc	Reused from Client subsystem, WKBCH CSCI, Data Visualization (EOSView) CSC
Interoperability	ADSRV	AdvNavigationServer	OTS	Gnrc	HTTP server
Interoperability	ADSRV	Client Library	DEV	Gnrc	
Interoperability	ADSRV	Core Library	DEV	Gnrc	
Interoperability	ADSRV	HTML Framework	DEV	Gnrc	
Interoperability	ADSRV	HTML Interfaces	DEV	Gnrc	
Interoperability	ADSRV	Installer	DEV	Gnrc	
Interoperability	ADSRV	Persistent Object Framework	DEV	Gnrc	
Internetworking	INCI	Datalink/Physical	OTS	Gnrc	firmware, vendor-supplied with hardware
Management	MACI	Application MIB	DEV	Gnrc	
Management	MACI	ECS Subagent	DEV	Gnrc	
Management	MACI	Encapsulator for non-Peer Agent	OTS/DEV	Gnrc	OptiMate
Management	MACI	Extensible SNMP Master Agent	OTS/DEV	Gnrc	Peer Network's agent, along with its toolkit for Dev
Management	MACI	Instrumentation Class Library	DEV	Gnrc	
Management	MACI	Management Agent Services	OTS/DEV	Gnrc	Peer and Tivoli/Sentry
Management	MACI	Proxy Agent	DEV	Gnrc	
Management	MACI	SNMP Manager's Deputy	DEV	Gnrc	
Management	MCI	Accountability	DEV	Gnrc	
Management	MCI	Application Management	DEV	Gnrc	
Management	MCI	Automatic Actions	DEV	Gnrc	
Management	MCI	Billing and Accounting Management	OTS/DEV	Gnrc	ITS selection in progress
Management	MCI	DCE Cell Management	OTS	Gnrc	HP Account Manager Toolr
Management	MCI	Diagnostic Tests	OTS	Gnrc	vendor-supplied with hardware
Management	MCI	Fault Management	OTS/DEV	Gnrc	Tivoli and HP OpenView

Table 3.3.2-1. GSFC DAAC Components Analysis (9 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Management	MCI	Management Data Access	DEV	Gnrc	
Management	MCI	Management Data Access User Interface	DEV	Gnrc	
Management	MCI	Management Framework	OTS	Gnrc	HP OpenView Network Node Manager
Management	MCI	Management Proxy	DEV	Gnrc	
Management	MCI	Mode Management	DEV	Gnrc	
Management	MCI	Network Manager	OTS	Gnrc	HP OpenView Network Node Manager
Management	MCI	Performance Management	OTS/ DEV	Gnrc	RFP released
Management	MCI	Performance Management Proxy	DEV	Gnrc	
Management	MCI	Performance Test	OTS	Gnrc	vendor-supplied with hardware
Management	MCI	Physical Configuration Management	OTS	Gnrc	Mountain View
Management	MCI	Physical Configuration Proxy Agent	DEV	Gnrc	
Management	MCI	Report Generation	OTS	Gnrc	No decision yet, evaluation in progress
Management	MCI	Report Generation and Distribution	DEV	Gnrc	
Management	MCI	Report Generation Manager	DEV	Gnrc	
Management	MCI	Resource Class Category	DEV	Gnrc	
Management	MCI	Security Databases	OTS	Gnrc	Operating System Password Files, DCE Registry Database, Router Configuration Files, TCP Wrappers configuration files, Operating System Access Control Lists, DCE Access Control Lists
Management	MCI	Security Management	DEV	Gnrc	
Management	MCI	Security Management Proxy	DEV	Gnrc	
Management	MCI	Security Tests	OTS	Gnrc	CRACK, COPS, SATAN, TRIPWIRE
Management	MCI	Trouble Ticketing Management Service	OTS	Gnrc	Remedy Action Request System

Table 3.3.2-1. GSFC DAAC Components Analysis (10 of 10)

Subsystem	CSCI	CSC	TYPE	USE	NOTES
Management	MCI	Trouble Ticketing Proxy Agent	DEV	Gnrc	
Management	MCI	Trouble Ticketing Service Requester	DEV	Gnrc	
Management	MCI	Trouble Ticketing User Interface	DEV	Gnrc	
Management	MCI	User Contact Tool	OTS/DEV	Gnrc	Remedy
Management	MCI	User Profile Server	DEV	Gnrc	
Management	MLCI	Baseline Manager	OTS/DEV	Gnrc	XRP II
Management	MLCI	Configuration Management	OTS	Gnrc	ClearCase
Management	MLCI	Inventory/Logistics/Maintenance (ILM) Manager	OTS/DEV	Gnrc	Vendor evaluation in progress
Management	MLCI	Policies and Procedures Management	DEV	Gnrc	
Management	MLCI	Software Change Manager	OTS/DEV	Gnrc	ClearCase
Management	MLCI	Software Distribution Management Structure	OTS/DEV	Gnrc	ClearCase and Tivoli
Management	MLCI	Software Request Manager	OTS/DEV	Gnrc	DDTS
Management	MLCI	Training Management	DEV	Gnrc	
Planning	PLANG	On Demand Manager	DEV	Gnrc	
Planning	PLANG	PDPS DBMS	OTS/DEV	Gnrc	Sybase DBMS
Planning	PLANG	Planning Object Library	OTS	Gnrc	Delphi C++ class libraries
Planning	PLANG	Production Planning Workbench	DEV	Gnrc	
Planning	PLANG	Production Request Editor	DEV	Gnrc	
Planning	PLANG	Resource Planning Workbench	DEV	Gnrc	
Planning	PLANG	Subscription Editor	DEV	Gnrc	
Planning	PLANG	Subscription Manager	DEV	Gnrc	

3.4 DAAC Hardware and Network Design

This section describes the ECS hardware and local area network design supporting the Release B ECS mission at the GSFC DAAC. Section 3.4.1 contains an overview diagram from the "networks" point of view, and detailed descriptions of the Release B LANs. Section 3.4.2 contains a hardware overview diagram of all of the ECS subsystems at GSFC, followed by detailed descriptions and rationale for each subsystem.

3.4.1 GSFC DAAC LAN Configuration

The Release B GSFC DAAC LAN design builds upon the design for Release A. The major design drivers and rationale for the Release A design are presented in the Release A GSFC DAAC Design Specification (305-CD-014-001), and that document may provide useful background information and context for the current Release B design.

The GSFC DAAC LAN topology is illustrated in Figure 3.4.1-1. The topology consists of a User FDDI Network, a Production FDDI Network, and a HiPPI Network (for processing flows). 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. Each of the networks is discussed in detail below.




Figure 3.4.1-1. GSFC DAAC LAN Topology

The Production Network consists of multiple FDDI rings supporting DAAC subsystems and connections to external production systems (such as EDOS, TSDIS, and other ECS DAACs) via EBnet. The separation and aggregation of hosts and subsystems onto FDDI rings is driven mostly by RMA and data flow requirements. For instance, Ingest is contained on an individual FDDI ring because of the strict RMA requirement for receipt of Level 0 data (0.998 with MDT of 15 minutes). RMA also dictates Ingest's direct connection to the EBnet router. Some Data Server hosts are contained on a dedicated FDDI ring in order to provide adequate bandwidth for the DAAC-to-DAAC and DAO processing flow requirements (see Table 3.4.1.1-1 below). The DM, LSM, and some Data Server hosts are contained on a single ring because their flows are expected to be fairly small given that user traffic will be processed on the separate User Network (see discussion below). Another ring provides access to the EBnet router to handle the DAAC-DAAC production flows. The FDDI Switch is the central device connecting the FDDI rings together, and it provides the necessary routing and filtering control. Note that the device counts in the figure are approximate and (for clarity) workstations and printers are not shown.

Current estimates of data flow volumes between the DAO processing system and the rest of the ECS DAAC (refer to Table 3.4.1.1-1) suggest that the interface to DAO can be achieved via FDDI. This is shown in Figure 3.4.1-1 by the FDDI interface from the Production Network FDDI switch to the DAO Production System. Note that if there is any change in the data flow estimates indicating network flow volumes in the HiPPI range, the ECS-to-DAO interface can be made to be HiPPI. Since the hardware implementation for the DAO processors is still under evaluation by DAO, the network topology interconnecting DAO processing hosts is not yet determined. The network design will be driven by the hardware architecture of DAO hosts (super-computer versus workstation farms, etc.) and the intra-DAO processing flow requirements.

The User Network is an FDDI-based LAN connecting users (via NSI, local campuses, general Internet, etc.) to DAAC hosts responsible for providing user access. It has the main advantage of separating user and production flows. This allows DAAC processing data flows to be unaffected by user demand, so that even unanticipated user pull will not hinder the production network. Basically, the User Network provides access to Data Manager hosts and to a subset of Data Server hosts that interact directly with users. Users will not have access to any other hosts, such as Ingest or Processing devices. The CSS and MSS servers are connected to the User Network but will not allow direct user access. These connections are required for communications with outside networks for such things as name lookups and receipt of Internet mail, as well as communication with and monitoring of the DAAC's interfaces to the user community (such as NSI and the local campus).

The User Network will connect to the GSFC Isolation LAN through an ECS router which will provide the necessary routing and filtering controls. NSI, the local GSFC Campus, and other Internet providers such as SURAnet will also be connected to the GSFC Isolation LAN. The DAAC will also have an additional direct connection to the NSI router because of the large user flows expected at GSFC. The ECS router will also connect to the two SMC FDDI rings (one containing the CSS and MSS Enterprise Servers and the other containing the Bulletin Board

Server) to provide the SMC with connectivity to NSI and the Internet. The GSFC DAAC EBnet router will also connect to the EOC router in order to transfer data between the SMC and the ECS DAACs. The details of the interfaces between ECS facilities and GSFC are shown in more detail in Figure 3.4.1-2.




Figure 3.4.1-2. ECS Network Interfaces at GSFC

Individual FDDI rings for the Production Network will be implemented with FDDI concentrators to provide ease of wiring and central points of management. All Processing Network DAAC hosts will have FDDI interfaces and will be attached directly to the FDDI rings. Workstations will have single-attached FDDI cards, whereas the high-performance servers and processors on the Production Network will have dual-attached FDDI cards to provide redundancy. Dual-attached hosts will be dual-homed to two separate FDDI concentrators to provide an additional level of redundancy in the event of a hub failure. Interfaces of User Network hosts will be single-attached, except for the Data Management Server, which will be dual-attached. Printers and X-terminals will be connected to the DS/DM/LSM FDDI ring via an FDDI-to-Ethernet hub.

The HiPPI Network interconnects Data Server hosts/devices and Science Processors in order to provide a high-speed network to handle the large data transfers between the two subsystems. The HiPPI network will be implemented via a central HiPPI switch with switched 800 Mbps interface ports connected directly to the high-powered processing and storage hosts. The HiPPI

Network shifts the numerous transfers of large volumes of data onto a dedicated high-speed topology, freeing the FDDI-based Production Network to handle control flows and DAAC-DAAC processing flows.

ECS will implement IP over the HiPPI Network. The prototyping effort undertaken by ECS to determine the performance of IP over HiPPI has shown that with proper tuning, the level of throughput needed for processing-to-dataserver flows can be achieved. An IP-based lightweight protocol called Bulk Data Service (BDS) will be used so that processing and dataserver applications can make effective use of the HiPPI fabric. Refer to the Release B CSMS Communication Subsystem Design Specification (305-CD-028-002) for more detail on BDS.

Quantities of networking hardware components to upgrade the GSFC DAAC network from Release A to Release B are presented in Table 3.4.1-1.

Table 3.4.1-1. Rel. B Networking Hardware for GSFC DAAC LAN

Networking Component	Release A Quantity	Release B Additional Quantity	Comments
FDDI Concentrator	10	2	Bay Networks 2914-04 concentrator with 12 M & 1 A/B port
FDDI Cables	89	35	Multimode fiber cables with MIC connectors
FDDI-to-Ethernet Hub	1	1	Cabletron MicroMMAC-22E; used for printers and X-terminals
Ethernet Cables	8	12	10baseT connection to printers and X-terminals
FDDI Switch/Router	1	0	Interconnects Production Network
FDDI Switch/Router Expansion Chassis with 2 FDDI modules	0	1	Expansion for Production Network FDDI Switch/Router
FDDI Router	0	1	Interconnects User Network
HiPPI Switch	0	1	High-speed switch connecting Data Server and Processing
HiPPI Cables	0	26	Copper cable

3.4.1.1 Sizing/Performance Rationale

The data flow estimates used as input to the design process for the NSIDC DAAC LAN topology are contained in Table 3.4.1.1-1. The table, based on static analysis of the February 1996 AHWGP baseline (results for epoch k (3Q99)) and the February 1996 User Pull Baseline (results for the greater of July 1999 and January 2000), is arranged according to the source and sink of the flow. It provides both raw 24-hour average data flows (including any applicable reprocessing) which are the output of ECS models, as well as weighted flows containing all overhead and contingency factors. The "Factors Applied" column shows which factors (listed beneath the table) were applied to each data flow.

Table 3.4.1.1-1. Estimated Rel. B Data Flows for the GSFC DAAC

Major Data Flow Description *	Raw Volume (in Mbps)	Factors Applied	Weighted Volume (in Mbps)
Ingest to Working Storage Server	6.5	2,3,4,5,6	18.2
Working Storage Server to Processing	257.6	1,4,5,6,7,8	584.8
Working Storage Server to/from FSMS Server	408.8	1,4,5,6,7,8	928.0
Working Storage Server to Distribution Server	38.7	2,3,4,5	108.8
Working Storage Server to/from DAO Processing	13.6	2,3,4,5,6	38.2
ACM Server to/from other DAACs (includes TSDIS)	17.0	1,2,3,4,5,6	57.5
User Pull	34.4	2,3,4,5	96.7

* Other flows such as session establishments amongst the subsystem hosts and subsystems to and from MSS are trace amounts and are not included in the table.

Overhead Factors:

- (1) SSI&T Factor: 1.2. Accounts for capacity for integration and test flows. This factor not applied to Ingest or User flows; factor also not applied to DAO flows because DAO estimates already include research and development flows, which are the equivalent of AI&T.
- (2) TCP/IP/FDDI Protocol Overhead: 1.25. This is applied only to the FDDI flows and not to HiPPI flows. Accounts for overhead associated with FDDI, IP, TCP, and other protocols (such as DCE).
- (3) FDDI Maximum Circuit Utilization Factor: 1.25. This is applied only to the FDDI flows and not to HiPPI flows. Accounts for amount of 100 Mbps bandwidth that is actually usable for sustained data rates.
- (4) Average-to-peak Conversion Factor: 1.5. This provides elasticity in the network by converting the 24 hour averages provided by the model into peaks.
- (5) Scheduling Contingency: 1.2. This reflects the ability for the network to recover within 24 hours from a 4 hour down-time ($24/20=1.2$).
- (6) Operational Hours Factor: 1.0 at GSFC. Accounts for percentage of day/week operations are performed. This factor is applied only to production flows, not to user flows.
- (7) HiPPI Protocol Efficiency: 1.05. This accounts for the protocol overhead associated with the HiPPI protocol contained in the packet stream. A factor of "1.05" was chosen because overhead comprised of the (LE and FP protocols) is very low when accessing the HiPPI fabric. This factor is applied only to flows over the HiPPI network.
- (8) HiPPI Maximum Circuit Utilization Factor: 1.0. This factor is applied only to flows over the HiPPI network. A factor of "1.0" was chosen because HiPPI is a dedicated medium and not a shared medium like FDDI. Thus, it is assumed each host will have access to the entire 800 Mbps channel connecting it to the HiPPI switch and that the HiPPI switch is non-blocking. The limiting factor in the HiPPI flows is the host (and peripheral access).

3.4.2 DAAC Hardware Configuration

The ECS DAAC hardware suite at GSFC hosts the Data Server, Data Management, Ingest, Interoperability, Production Planning, Science Data Processing, Management and Communications subsystems. The hardware and COTS software selected for the GSFC configuration is illustrated in Figure 3.4.2-1, GSFC ECS DAAC Hardware Configuration Diagram. These configurations represent the candidate hardware selections which most closely satisfy the processing, storage capacities and communications bandwidth requirements described in the following sections. In some cases the selected configuration appears to significantly exceed the requirements to due to the sizing increments provided by the selected vendor, when in reality, our analysis and selection process has provided cost effective solutions to each problem.

This page intentionally left blank.

GSFC at Rel B (1 of 3)

Added for Release B

Note: All systems come with OS, C, C++, DCE, OODCE, Clearcase agent, SNMP agent (Optima), Sybase client, Tivoli client, CD-ROM, FDDI. RAID quantities are USABLE.

<p>SPRHW-GSFC-1 SGI PC XL 12CPU 2GB RAM; 4GB disk Trade-in 12xR8000 for 18xR10000s; +8GB Disk, +HiPPI</p> <p>IDL, IMSL F77, F90 DCE Dev Kit CASEVision Autosys Agent</p> <p>RAID (replaces Rel A): SPRHW-GSFC-1B/C 10x4.3GB each (Rack) SPRHW-GSFC-1D/E 10x9GB each (Rack) SPRHW-GSFC-1F 10x4.3GB sep. rack RAID Configs: Rack 2: (SPRHW-GSFC-5A/D) 10x9GB; 10x9GB 10x4.3GB; 10x4.3GB Rack 3: (SPRHW-GSFC-6A/D) 10x9GB; 10x9GB 10x4.3GB; 10x4.3GB Rack 4: (SPRHW0GSFC-8A/D) 2*10x9GB; 3*10x4.3GB Rack 5: (SPRHW-GSFC-5E, 6E,8E) 10x4.3GB each Rack 6: SPRHW-GSFC-9A/D: 10x9GB;10x9GB; 10x4.3GB;10x4.3GB</p>	<p>SPRHW-GSFC-2 X-Terminal NCD HMX Pro 16MB</p> <p>SPRHW-GSFC-3 X-Terminal NCD HMX Pro 16 MB</p> <p>SPRHW-GSFC-5 SGI PC XL 18CPU 18GB Disk; 2GB RAM; HiPPI</p> <p>F77, F90 IDL, IMSL</p>	<p>SPRHW-GSFC-11 Science Processor SGI Indy 256MB; 4GB (MSFC Ir1 String)</p> <p>SPRHW-GSFC-10 X-Terminal NCD HMX Pro 16MB (MSFC Ir1 String)</p> <p>SPRHW-GSFC-6 SGI PC XL 18CPU 18GB Internal Disk 2GB RAM: HiPPI</p> <p>F77, F90 IDL, IMSL</p>	<p>SPRHW-GSFC-7A 10GB Disk</p> <p>SPRHW-GSFC-7B 10GB Disk</p> <p>SPRHW-GSFC-7 Queuing Server SUN Ultra 8-slot 512MB RAM; 4GB 4CPUs</p> <p>SPRHW-GSFC-12 X-Terminal NCD HMX Pro 16MB</p> <p>SPRHW-GSFC-13 X-Terminal NCD HMX Pro 16MB</p> <p>SPRHW-GSFC-14 X-Terminal NCD HMX Pro 16MB</p> <p>SPRHW-GSFC-15 X-Terminal NCD HMX Pro 16MB</p>	<p>SPRHW-GSFC-5A 72GB RAID</p> <p>SPRHW-GSFC-5B 72GB RAID</p> <p>SPRHW-GSFC-5C 34GB RAID</p> <p>SPRHW-GSFC-5D 34GB RAID</p> <p>SPRHW-GSFC-5E 34GB RAID</p> <p>SPRHW-GSFC-4A 68GB RAID +Rack</p> <p>SPRHW-GSFC-4 SGI Challenge DM 256MB; 2GB +3 150MHz R4400 +256MB;+4GB Disk (LIS Support)</p> <p>IDL, IMSL F77, F90 DCE Dev Kit CASEVision Autosys Agent</p>	<p>SPRHW-GSFC-6A 72GB RAID</p> <p>SPRHW-GSFC-6B 72GB RAID</p> <p>SPRHW-GSFC-6C 34GB RAID</p> <p>SPRHW-GSFC-6D 34GB RAID</p> <p>SPRHW-GSFC-6E 34GB RAID</p> <p>SPRHW-GSFC-8A 72GB RAID</p> <p>SPRHW-GSFC-8B 72GB RAID</p> <p>SPRHW-GSFC-8C 34GB RAID</p> <p>SPRHW-GSFC-8D 34GB RAID</p> <p>SPRHW-GSFC-8E 34GB RAID</p>	<p>SPRHW-GSFC-9 SGI PC XL 8CPU 18GB Internal Disk 1GB RAM: HiPPI</p> <p>F77, F90 IDL, IMSL</p>	<p>AQAHW-GSFC-1 QA WS SGI Indigo2 Impact 256MB; 4GB</p> <p>AQAHW-GSFC-1A 17GB RAID</p> <p>SPRHW-GSFC-9A 72GB RAID</p> <p>SPRHW-GSFC-9B 72GB RAID</p> <p>SPRHW-GSFC-9C 34GB RAID</p> <p>SPRHW-GSFC-9D 34GB RAID</p>
--	---	--	---	--	---	---	--

<p>AITHW-GSFC-3 Printer HP LaserJet 4M+ 12ppm/14MB RAM</p>	<p>AITHW-GSFC-2 AIT WS/DBMS server SUN Sparc 20/50 128MB, 8GB disk</p> <p>SPARCworks IDL, IMSL F77, F90 SoftPC/MSSoffice Sybase, DCE Dev Kit</p>	<p>AITHW-GSFC-1 AIT WS SUN Sparc 20/50 64MB; 6GB</p> <p>SPARCworks IDL, IMSL F77, F90 SoftPC/MSSoffice Forcheck</p>	<p>AITHW-GSFC-4 Printer HP LaserJet 4M+ 12ppm/14MB RAM (Ir1 - A)</p>	<p>AITHW-GSFC-5 AIT WS/DBMS server SUN Sparc 20/50 128MB, 8 GB disk (Ir1 - A)</p>	<p>AITHW-GSFC-6 AIT WS SUN Sparc 20/50 64MB; 4GB (Ir1 - A)</p>	<p>AITHW-GSFC-7 AIT WS SUN Sparc 20/50 64MB; 4GB (MSFC Ir1 String)</p>	<p>AITHW-GSFC-8 AIT WS SUN Sparc 20/50 64MB; 4GB (MSFC Ir1 String)</p>
---	---	--	---	--	---	---	---

<p>MSS-GSFC-3 MSS WS SUN Ultra 170 128MB;4GB</p> <p>WABI/MSSoffice ZMail, SPARCworks Billing & Acctg Server & Client DBMS Rept Generation Statistical Analysis / Trending SW</p>	<p>MSS-GSFC-2 Printer HP LaserJet 4M+ 12ppm; 14MB</p>	<p>MSS-GSFC-1 MSS WS SUN Ultra 170 128MB;8GB</p> <p>Clearcase Server DDTS (Puresoft) XRP (HTG) ZMail, SPARCworks ILM SW WABI/MSSoffice Billing & Acctg Client SW License Mgmt</p>	<p>MSS-GSFC-4 MSS Server HP J210/1 256MB; 4GB +1CPU;+128MB</p> <p>OpenView (HP) TME (Tivoli) ARS (Remedy) MountainView/PNM WABI/MSSoffice DBTools.h++, Tools.h++ Sybase Server Courier (Tivoli)</p>	<p>MSS-GSFC-4A 20GB RAID +45GB RAID</p>	<p>CSS-GSFC-1 CSS Server HP J210/1 256MB; 4GB +1 CPU;+128MB</p> <p>DCE Name DCE Security Softbench WABI/MSSoffice Mail Server Master Sybase Srvr</p>	<p>PLNHW-GSFC-1 Prod Plan/Mgmt WS SUN Sparc 20/71 384MB, 4GB +128MB RAM;</p> <p>AutoSys SPARCworks Sybase SoftPC/MSSoffice DBTools.h++, Tools.h++</p>	<p>PLNHW-GSFC-3 Prod Plan/Mgmt WS SUN Sparc 20/71 384MB; 4GB +128MB</p> <p>(LIS Support)</p>	<p>PLNHW-GSFC-2 PDPS DBMS Svr SUN Ultra 8-Slot 512MB; 4GB 4CPUs</p> <p>AutoSys SPARCworks Sybase SoftPC/MSSoffice DBTools.h++, Tools.h++</p>	<p>PLNHW-GSFC-2A 10GB Disk</p> <p>PLNHW-GSFC-2B 10GB Disk</p>
---	--	--	--	--	---	--	---	---	---

<p>MSS-GSFC-9 Impact Printer</p>	<p>MSS-GSFC-5 Backup Server SUN Ultra 4-slot 256MB; 4GB Legato Networker SW</p>	<p>MSS-GSFC-5A DLT Stacker</p>	<p>MSS-GSFC-6 Printer HP LaserJet 4M+ 12ppm; 14MB (MSFC Ir1 String)</p>	<p>MSS-GSFC-7 WS SUN Sparc 20/50 64MB; 4GB (MSFC Ir1 String)</p>	<p>MSS-GSFC-4 WS SUN Sparc 20/71 128MB; 8GB (Ir1 - A Transition)</p>	<p>MSS-GSFC-1A WS SUN Sparc 20/71 128MB; 8GB (Ir1 - A Transition)</p>	<p>MSS-GSFC-2A Printer HP LaserJet 4M+ 12ppm; 14MB (LIS Support)</p>	<p>MSS-GSFC-2B Printer HP LaserJet 4M+ 12ppm; 14MB (LIS Support)</p>	<p>MSS-GSFC-8 WS SUN Sparc 20/50 64MB; 4GB (LIS Support)</p>
---	--	---	--	---	---	--	---	---	---


- ISSI**
- FDDI Switch Quantity: 1
 - Alantec Power Hub 7000 w/4 FDDI cards; 3 power supplies
 - FDDI Concentrators: Bay Networks 2914-04 Quantity: 10
 - FDDI Cables: Quantity: 89
 - Ethernet Hub: Quantity: 1
 - Cabletron Micro MAC-22E w/BRIM F6
 - Ethernet Cables: Quantity: 8
 - LAN Analyzer Quantity: 1
 - Comm Cabinets: Quantity: 2
 - FDDI Router Quantity: 1
 - FDDI Concentrator Quantity: 2
 - FDDI Cables Quantity: 35
 - HiPPI Switch Quantity: 1
 - HiPPI Cable: Quantity: 26
 - Ethernet Hub: Quantity: 1
 - Ethernet Cables: Quantity: 12

Hold Page for 11 x 17 (back)

GSFC at Rel B (2 of 3)

Added for Release B

Note: All systems come with OS, C, C++, DCE, OODCE, Clearcase agent, SNMP agent (Optima), Sybase client, Tivoli client, CD-ROM, FDDI. RAID quantities are USABLE.




Hold Page for 11 x 17 (back)

GSFC at Rel B (3 of 3)

Added for Release B





Note: All systems come with OS, C, C++, DCE, OODCE, Clearcase agent, SNMP agent (Optima), Sybase client, Tivoli client, CD-ROM, FDDI. RAID quantities are USABLE.



LaRC at Rel B (1 of 2)

Added for Release B

Note: All systems come with OS, C, C++, DCE, OODCE, Clearcase agent, SNMP agent (Optima), Sybase client, Tivoli client, CD-ROM, FDDI. RAID quantities are USABLE.



This page intentionally left blank.

3.4.2.1 Client Subsystem

There is no dedicated hardware support (HWCI) for the Client Subsystem. The Client software configurations are supported by: (1) non-ECS provided hardware platforms, in the case of Client software utilized by the user community, or (2) ECS provided workstations utilizing Client software in support of operations users (network management, DAAC operations, etc.).

3.4.2.2 Data Server Subsystem

During the Release B time frame at the GSFC DAAC, a Data Server configuration is supplied to support the following data storage and access services: TRMM instruments, AM-1 instruments and data as described in the February 1996 Technical Baseline, as well as V0 migration data. The configuration discussed in the subsections that follow provide a design of the configuration and sizing of the Data Server hardware for Release B. Sizing refinements as compared to the IDR document are based on project level decisions in the areas of reprocessing, backup, user pull sizing, produce versus store policies, and adjustments to the technical baseline driven by processing modifications received from the instrument teams. The reader is encouraged to refer to the Data Server specific volume of DID 305 for more details on the operational and engineering concepts that form the basis for the DSS Release B design and a more detailed discussion of these design variables.

The Data Server configuration has made some assumptions in the area of reprocessing in regards to disk, robotics, tape drives, and general I/O. The dynamic model used for sizing produced peak and average usage statistics for Data Server resources, monitored and calculated over the full processing cycles of the Release B products (i.e., daily, weekly, monthly, etc.). The peaks are caused by different data staging requirements of specific PGEs. The Processing Subsystem's disk usage patterns are consistent with this pattern. The dynamic model has also shown, that Data Server archive hardware components' sizing to two times the average computed requirements will be sufficient to keep up with the overall level of demand. reprocessing load, estimated to be equal to one time the processing load, will be worked off in the off-peak processing times, as shown by the analysis done after IDR B.

The Data Server Subsystem hardware configuration consists of five hardware CIs and is sized for the mission support as described above for a period of one calendar year beyond AM-1 launch date plus V0 data migration for the same time period:

- *Access Control & Management (ACMHW HWCI)* -- The access hardware allows for client access (both the client subsystem and direct "push/pull" user access) to the Data Server subsystem, provides tools and capabilities for system administration, and is broken down into two components; Administration Stations (AS) which consist typically of operations support workstations, and Access/Process Coordinators (APCs) which consist of server class machines with host attached disk pools.

- *Working Storage (WKSHW HWCI)* -- Working Storage (WS) hardware of the Data Server supplies a pool of storage used for temporary file and buffer storage within the Data Server architecture. All data in WS regardless of storage duration is still considered temporary in nature and not part of the permanent archive.
- *Distribution and Ingest Peripheral Management (DIPHW HWCI)* -- The hardware of the Distribution and Ingest Peripheral Management supports the hard media distribution methods for data dissemination from the system, as well as hard media ingest of data into the system. The hardware provided by this HWCI includes a variety of media and media drives, jukeboxes/stackers as necessary, and server hosts and disk storage for network distribution.
- *Data Repository (DRPHW HWCI)* -- Data Repositories (DRs) are the hardware components that store and maintain data permanently. This consists of DBMS repositories and archive tape library based repositories. This HWCI provides the disk, server, archive robotics, media and archive tape drives required to support the permanent storage repositories.
- *Document Data Server (DDSHW HWCI)* -- This HWCI provides the disk and server required to support the Document Data Server portion of the DSS.

3.4.2.2.1 Rationale

The following subsystem-wide assumptions were applied in sizing the Data Server hardware components. Data Server Subsystem is sized for : TRMM instruments, AM-1 instruments, selected backup data from other DAACs, and data as described in the February 1996 Technical Baseline, as well as V0 migration data support for a period of one calendar year beyond AM1 launch date. A combination of static and dynamic modeling was used to size the permanent data repository components, such as the number of robotic arms, tape drives, and production related staging disk. The ECS technical baseline for February 1996 was used both in dynamic modeling as well as in static analysis. User modeling data was used in dynamic models for estimating the user access rates to the system where appropriate, as well as the pattern of user interaction with the system. However, DSS sizing is still being driven by the ECS Project direction to size for a distribution volume two times the site's nominal production volume. The DAAC to DAAC flow estimates are derived from the February 1996 Technical Baseline.

ACMHW Analysis was undertaken primarily for the sizing of the APC server hosts and their attached disk storage. The administration workstations are assumed to continue to be minimally configured workstations designed to perform various operations functions (e.g. DBMS administration, repository administration, etc.). Client desktop services as well as X protocol access to Data Server hosts is the driving sizing factor.

APC Server -- The APC server provides the session establishment point for the client, routes service requests to the other CSCIs and HWCI, provides subsetting and subsampling capabilities in accordance with the Appendix F of DID 304 (Data Type Matrix), and functions as the electronic distribution channel for the DSS user pull and push loads. The APC server host, runs the following software processes and applications: ScienceDataServer Process

ScienceDataServer Administration Process, SubscriptionServer Process, Network ResourceManager Process, PullMonitor Process, CSS DCE client, CSS logging API, and an MSS agent. The internal I/O to be sustained by the processor averaged over 24 hours of operation is not a driving design factor.

The SGI Challenge XL series total I/O bandwidth is 320 MB/sec, which is more than adequate to meet the Release B requirements.

The following is the CPU allocation (CPU loading for peripheral support is based on SGI configuration guidelines):

- 0.5 CPU Monitoring/Managing Tasks (estimated)*
- + *10 CPUs Subsetting/Subsampling (Dynamic Modeling)*
- + *0.07 CPU Subscriptions (Modeling estimates)*
- + *1 CPU Operating System and Applications (estimated)*
- + *2.3 CPUs for client support (session threads)*
- + *0.075 CPU for 1 Differential SCSI -II Interface (1 SCSI [disk] at 5 MB/sec*
*7.5% of a CPU for each 5 MB/sec; 1 * 5 * 0.075/5)*
- + *0.06 CPUs for SCSI RAID controller support (1.5% of a CPU for each*
*5 MB/sec of controller I/O; 18.5*0.015/5)*
- + *0.21 CPUs for FDDI I/O support (7.5% of a CPU for each*
*5 MB/sec of FDDI I/O: 14.2 MB/sec / 5 MB/sec * 0.075)*
- + *1.75 CPUs for HIPPI I/O support.*

This requires a total of 15.97 equivalent CPUs. The SGI Challenge XL server with 16 CPUs and 1 GB of RAM will be selected as the APC server. The server will be equipped with 6 GB of local disk, 1 differential SCSI-II controller for the local disk pool, 24 SCSI RAID Controllers (see *APC Host Storage* below), and 2 FDDI Interfaces. A degraded mode failover will be provided by a second SGI Challenge XL equipped with 8 CPUs, 512 MB of RAM and an identical complement of peripheral interfaces. Here, and elsewhere at the GSFC site all SGI CPUs are 200 MHz R10000.

APC Host Disk -- The APC server host disk is sized for electronic ingest (almost exclusively from sources external to the DAAC), as well as electronic distribution (again, almost exclusively to the recipients external to the DAAC) since this pool of hosts is designed to manage the requests to the Data Server as well as service the response. APC storage must also be sized to support functions such as subsetting/subsampling and storage of user session context, which keeps track of user session interactions that may be suspended and resumed.

For Release B, the electronic distribution, as defined within the ECS Technical Baseline for CDR, is "one times" the total Data Server ingested volume. This "one times" volume is made up of data that is "pushed" and "pulled" by the user. The distribution of this electronic data is assumed to occur over a 48-hour period. It is further assumed that on the average each granule

placed on the distribution volumes are pulled/pushed to two separate users. This assumption is based on the workings of subscriptions, (assuming that a granule is subscribed to more than once), and the fact that interesting data will be requested more than once when inserted while some data may not be requested at all.

Distribution of data to users is estimated at 4.3 MB/s (371.5 GB/day), with 4.9 MB/s (4.4 GB accumulation in 15 minutes disk latency) accounting for data received from users and other DAACs, and data sent to other DAACs. It is assumed, that 100% of DAAC to DAAC transfers are electronic, with the backup data flow (estimated as 2% of production) performed via hard media (DIPHW CI). While both the I/O capacity and the disk allocations take these transfers into account, the rate is not a driving factor in the design. Allowing for two days worth of capacity on a fast access medium, this yields approximately 764 GB of RAID. Rolled into this disk allocation is the 15.5 GB/day estimated by Dynamic Modeling to be required to support subsetting functions, and 1 GB for client sessions.

WKSHW The rationale for allocation of resources to the WKSHW HWCI has changed from that presented in the IDR version of this document for several reasons. The most significant change is the requirement to store interim products. This requirement for temporary storage, has been calculated via the dynamic model using inputs from the February 1996 Technical Baseline. The estimated peak accumulation for interim products at GSFC is 158.5 GB. The overall WKSHW requirement is increased by this amount. Working storage resources had also been assumed at IDR-B to be network attached and largely pooled, allowing sharing of disk resources among servers from the other DSS HWCI's. The technology required to accomplish the pooling of disk resources did not mature as anticipated in some cases, or is only available in a homogeneous vendor solution, and is prohibitively expensive in others. Therefore, the WKSHW HWCI requirement has been expanded to include, at GSFC, hosts in addition to RAID. The WKSHW server host included in the GSFC WKSHW will run the following processes and applications: ResourceManager Process for the staging disk, StagedDataMonitor Process, CSS DCE client, CSS logging API, and an MSS agent.

The following is the total CPU allocation (CPU loading for peripheral support is based on SGI configuration guidelines):

- 0.5 CPU Monitoring/Managing Tasks (estimated)*
- + *1 CPU Operating System and Applications (estimated)*
- + *2.3 CPUs for client support (session threads)*
- + *0.075 CPU for 1 Differential SCSI -II Interface (1 SCSI [disk] at 5 MB/sec*
*7.5% of a CPU for each 5 MB/sec; $1 * 5 * 0.075/5$)*
- + *0.35 CPUs for SCSI RAID controller support (1.5% of a CPU for each*
*5 MB/sec of controller I/O; $116.8 * 0.015/5$)*
- + *0.42 CPUs for FDDI I/O support (7.5% of a CPU for each*
*5 MB/sec of FDDI I/O: $28.1 \text{ MB/sec} / 5 \text{ MB/sec} * 0.075$)*
- + *3.5 CPUs for HIPPI I/O support*

This requires a total of 8.15 equivalent CPUs. The driving factor in the sizing of this host in this case is the large number of SCSI, FDDI, and HIPPI controllers that must be supported. The server will be an SGI Challenge XL configured with 10 CPUs and 1 GB of RAM. In addition, it will be equipped with 6 GB of local disk, 1 differential SCSI-II controllers for interface to local disk, 8 SCSI RAID Controllers, 2 HIPPI interfaces, and 2 FDDI Interfaces. A degraded mode failover will be provided by a second SGI Challenge XL equipped with 6 CPUs and an identical complement of peripheral interfaces. Here and elsewhere at the GSFC site all SGI CPUs contain 200 MHz R10000 processors.

The overall data storage requirement within the WKSHW HWCI is 558.5 GB. The RAID requirement for working storage is 400 GB (assumed Total Volume Groups at GSFC multiplied by the size of a medium multiplied by 2; $4 \times 50 \text{ GB} \times 2 = 400 \text{ GB}$), plus an additional 158.5 GB of disk to support peak interim data accumulation. This provides sufficient storage to support 2 hours worth of data at steady state to and from processing, hard media ingest, electronic V0 migration data ingest, DAAC data ingest, and client support. Working storage RAID is designed to be a temporary staging area only, not a long-term repository, so 2 hours worth of space provides adequate time to either write the data to tape or forward it to its destination. At GSFC the steady state rates alone were deemed to provide an inadequate model for discrete data transfers and a separate calculation was used as stated above. Within the total amount of storage there is a sufficient buffer (185 GB) for a six hours accumulation of the Data Processing to Repository data sub-flow alone. A six hour buffer for the Processing to Repository data sub-flow was estimated by the Dynamic Modeling to provide a reduction of approximately 20% in the number of transactions between the WKSHW and DRPHW.

DIPHW The DIPHW configuration at GSFC includes a host server and disk units to support temporary staging for physical media based production and distribution, (as well as for some physical media Ingest). A number of peripheral devices of varying form factors are provided. Two servers (Sun) are provided for temporary staging and peripheral support in the Release B time frame. The distribution peripherals are spread over both servers due to SCSI I/O limitations on a single server.

Server -- This server is sized to support the media distribution load (as well as a small factor for ingest loads). It is designed to handle the I/O for 1x (one times) distribution of the Data Server ingested volume per twenty-four hours of operation, servicing this load in a nominal 8 hr. total shift period. The platform is sized to handle the network transfer traffic to local disk from the WKSHW Host source, as well as the media preparation and media ingest I/O. The software processes/applications mapped to this server are: DistributionServer process, ResourceManager Processes for CDROM, various tape, and printers, CSS DCE client, CSS logging API, and an MSS agent. Two server hosts are provided in order to comply with RMA requirements for the function of archiving and distributing data with the required availability of 0.98 and the mean down time not to exceed 2 hours.

The disk associated with the hard media distribution platform is sized for hard media distribution of 1x (one times) of the total Data Server ingested volume from local Data Processing per twenty-four hour operations. The same assumption is used here as was used in the APC disk sizing, namely that on the average each granule placed on the distribution volumes are written

twice to two separate users. 24 hours accumulation capacity, plus ten percent, for potential hard media ingest is sized. The ingest activity is assumed to be negligible, since the Ingest Subsystem is sized to handle the high volume electronic ingest loads. The DIPHW RAID is sized at 380 GB, with 6 GB of local disk for each server.

The servers included in the DIPHW must, in aggregate, support 13.2 MB/s of distribution (24 hours of accumulation in an 8 hour shift), .32 MB/s of ingest via hard media, and 0.18 MB/s for backup of all other DAACs data excluding GSFC.

The CPU allocation for each of the two servers is as follows:

- 0.5 CPU Monitoring/Managing Tasks (estimated)*
- + *1 CPU Operating System and Applications (estimated)*
- + *1 CPU for 8.62 MB/s of I/O to the various Differential SCSI -II Interfaces for each peripheral*
- + *0.2 CPUs for Fiber Channel RAID controller support*
- + *0.2 CPUs for FDDI I/O.*

This requires a total of 2.9 equivalent CPUs. The Sun Ultra with 4 CPUs, 6 GB of local disk, and 256 MB of RAM is specified for these servers. A total of 380 GB of RAID will also be allocated across the two servers in this HWCI.

Peripheral Support -- The peripherals supported at the GSFC site for Release B were selected based on Level Four requirements: S-DSS-30440, S-DSS-30470, S-DSS-30480 (reference SDPS Requirements Specification for the ECS Project, 304-CD-002-001). We have provided a single 6250 bpi tape drive. In the case of 6250, V0 operational experience has shown that its use as a heritage device is low enough to make the provision of a spare drive unnecessary. In the event of the drive failure the low workload will allow waiting for the drive repair or replacement without a noticeable impact on operations. The HWCI complement may be easily scaled for both media types as well as capacity. The peripherals supplied here are included in the configuration to primarily support distribution functions. However, the Ingest Subsystem (Ingest Client) residing at GSFC may utilize peripherals to perform some media based ingest, as necessary, based on media form factors received for storage. This applies only to peripherals not already configured into the Ingest complement for performance reasons.

The types of media form factors/formats selected for Release B include:

- 8mm Tape
- 6250 Tape (heritage)
- CD-ROM
- 3480/3490
- 4 mm Tape
- FAX/scanner

The aggregate bandwidth of the 8mm and 4mm peripheral devices will equal the estimated required bandwidth of media based data distribution and ingest in Release B. The quantities of other peripheral devices supplied will be dictated largely by the RMA concerns. This HWCI also provides the network laser printers for the DAAC DS operations.

DRPHW The DRPHW configuration at GSFC includes both archive based as well as DBMS based physical repositories. They are sized as follows:

Archive Repository -- The archive component was sized through a combination of both static analysis as well as dynamic simulation. The principal model employed was a static model that utilized as input both equipment performance parameters and data flow estimations that were output from the dynamic model. The equipment performance parameters are vendor specifications derated based on reported user experience. The dynamic system model was based on a discrete event simulation with constrained resources assumed. Modeling runs based on the February 1996 Technical Baseline were performed. The dynamic modeling results were also used as a check of the static model results. Specific data regarding disk, tape drive and robotic resource utilization was obtained from the static model. This data was coupled with the key driving requirements with respect to distribution (e.g. the 2x distribution cap, the User Model service access predictions for epoch k), flow analysis, and data with respect to hardware and software COTS selections.

AMASS was chosen for the Release B archive FSMS. For the FSMS Manager host server, the platform is selected on the basis of FSMS(AMASS)/platform COTS S/W compatibility. Memory and cache estimates are currently based on vendor recommendations, reference "AMASS Archival Management and Storage System, Installation on Silicon Graphics", EMASS Part Number 600149, AMASS Version 4.2.4, March 1995. Aside from the AMASS FSMS the following processes and applications will run on this server: ResourceManager Process for the staging disk, StagedDataMonitor Process, CSS DCE client, CSS logging API, and an MSS agent.

The DRPHW is required to support and aggregate I/O rate of 51.5 MB/s in the Release B time frame. This requires the use of three Challenge XL Archive Server hosts and the addition of five robotic tape libraries.

The CPU allocation for archive servers 1 & 2 is as follows:

- 0.5 CPU Monitoring/Managing Tasks (estimated)*
- + *1 CPU Operating System and Applications (estimated)*
- + *0.18 CPU FSMS Allocation (22 MB/sec total flow to and from the ATL *0.8% of a CPU per MB of flow)*
- + *1.39 CPU for 33 Differential SCSI -II (2 SCSI [disk] at 5 MB/sec*
7.5% of a CPU for each 5 MB/sec +31 SCSI [tape] at 5 MB/sec - 0.8
*% CPU/ 1 MB/sec; 2 * 5 * 0.075/5 + 31*5*0.008/1)*
- + *0.07 CPUs for 7 SCSI RAID controllers support (1.5% of a CPU for each*
*5 MB/sec of controller I/O; 22*0.015/5)*

- + 0.8 CPUs for FDDI I/O support (*{control flow only: 10% of the total data flow of 51.5 MB/s}/5 MB/sec * 0.075*)
- + 1.75 CPUs for HIPPI I/O support.

This requires a total of 5.69 equivalent CPUs. An SGI Challenge XL server with 6 CPUs will serve as DRPHW archive servers 1 & 2. An AMASS cache of 200 GB, 6 GB of local disk and 512 MB of RAM are also required for each server. It is believed that steady state rates provide an inadequate model for discrete data transfers via the AMASS cache, thus a separate calculation is used. This calculation is the Total Volume Groups (in this case, per DRP Server, so the maximum is 2) at GSFC multiplied by the size of a single archive medium multiplied by 2. The calculation is, (TVG x M x 2); 2 x 50 GB x 2 = 200 GB for the AMASS Cache on DRP Servers 1, 2, & 3.

The CPU allocation for archive server 3 is as follows:

- 0.5 CPU Monitoring/Managing Tasks (estimated)*
- + *1 CPU Operating System and Applications (estimated)*
- + *0.1 CPU FSMS Allocation (10 MB/sec { + 3 MB/sec from the Release A HW} total flow to and from the ATL *0.8% of a CPU per MB of flow)*
- + *0.79 CPU for 18 Differential SCSI -II (2 SCSI [disk] at 5 MB/sec 7.5% of a CPU for each 5 MB/sec +12 SCSI [tape] at 5 MB/sec - 0.8 % CPU/ 1 MB/sec; 2 * 5 * 0.075/5 + 16*5*0.008/1){includes Release A HW}*
- + *0.07 CPUs for 7 SCSI RAID controllers support (1.5% of a CPU for each 5 MB/sec of controller I/O; 22*0.015/5)*
- + *0.8 CPUs for FDDI I/O support ({control flow only: 10% of the total data flow of 51.5 MB/s}/5 MB/sec * 0.075)*
- + *1.75 CPUs for HIPPI I/O support.*

This requires a total of 5.01 equivalent CPUs. An SGI Challenge XL server with 6 CPUs will serve as DRPHW archive server 3. An AMASS cache of 200 GB, 6 GB of local disk and 512 MB of RAM are also required for each server.

The EMASS AML robot selected for Release A will continue in the Release B archive to provide Release A migration, backup storage capacity, and support for special multi-media needs at GSFC. A combination of factors, among these are a change in vendor supported configurations and a dramatic increase in archive transaction requirements, have resulted in the unsuitability of EMASS AML Robots for the Release B time frame.

The STK Powderhorn Archive Tape Library (ATL) Robot was selected for Release B based on its ability to provide high speed transactions, storage density, growth capacity, floor space utilization, compatibility with COTS software and a number of other factors. (This selection process is documented in the SDPS Storage Technology Insertion Plan white paper (June 1995,

#420-WP-003-001) and is not fully discussed here. Given data with respect to the robotic device, and assuming the performance signature of D3 drives in the robotic unit, the analysis resulted in a utilization factor for sizing estimation purposes as given below:

GSFC: Configuration: 5 units

Calculated Tape Drive Resource Utilization: 16 units in each of four ATLs with 12 units in the 5th ATL for a total of 76 units.

The Release B configuration specifies the use of D3 (helical) tape (50 GB/cartridge) technology, yielding a total robotic storage density of approximately 300 TB per ATL (assuming 6000 cartridges/ATL). In Release B the accumulated GSFC data volume is projected to be approximately 330.4 TB. The volume calculations are based on the Technical Baseline of February 1996 for the Release B product data volumes and for the identified V0 migration. It is assumed that the product data accumulates over the 9/97 through 6/99 time period. The Release B volume of 330.4 TB equates to approximately 6,608 50 GB cartridges. No compression is assumed. Backup capacity of 2% (6.61 TB) on-site plus 2% of ASF, EDC, LaRC, JPL, and NSIDC capacity (10 TB) are also supported in the GSFC archive. Browse data capacity of 1MB per granule (16.07 TB), and a spare tape capacity of 10% (34.7 TB) are added. Total volume to be supported in the GSFC DAAC archive is then 397.8 TB. The total calculated media support requirement for Release B and V0 operations is as follows:

Release B ECS = 330,400 GB + 2% GSFC backup capacity + 2% backup of other DAACs + browse data + 10% spare media = 397,780 GB --> 7,956 count of D3 tapes required.

DBMS Repository -- The Data Base Management System (DBMS) Repository component was sized as follows based on static data size analysis as well as transaction based analysis. The transaction analysis is based on both "push" (production metadata update) and "pull" (user access and distribution) loads. Transaction rate was modeled based on the user service request rates as described at PDR time in the User Pull Analysis Notebook, 160-TP-004-001, Question 47 and a cross section of query types derived from the DBMS Benchmark Report, 430-TP-003-001.

The DBMS Server Host was sized based on the transaction analysis mentioned above, as well as platform suitability analysis based on the DBMS COTS software selection for Release B Data Servers (Illustra). Platform suitability is based on the DBMS software manufacturer's compatibility recommendations, benchmark data, and project bench marking activities. Aside from the Data Base engine, the following processes will run on this host: CSS DCE client, CSS logging API, and an MSS agent.

DBMS Server Disk was sized based on the *core metadata* associated with TRMM and AM-1 data as well that associated with the V0 data sets identified for migration within Release B together with metadata associated with other data sets identified for migration with the Release B time frame. This time frame is from 1997 through 1999. Refer to Table 3.4.2.2-3 for DBMS Repository Sizing beyond the Release B time frame. Table 3.4.2.2-2 identifies the Release A and B datasets used for the sizing. The size of the metadata granule for V0 was assumed to be half (0.91 KB) that of the full ECS metadata granule of 1.82 KB. That assumption is based on the 50% mapping of V0 attributes to ECS core attributes across all products throughout the Release A duration.

**Table 3.4.2.2-2. Release A and B Datasets Held by the Data Server Subsystem
(1 of 2)**

Product	Instrument	Product Description	Platform
1A-VIRS	VIRS	VIRS Level 1A Product	TSDIS
1B-01	VIRS	VIRS Radiances	TSDIS
3B-42	Various	Rain estimates	TSDIS
3B-43	Various	Surface rainfall map	TSDIS
1A-TMI	TMI	TMI Level 1A Product	TSDIS
1B-11	TMI	TMI Radiances	TSDIS
2A-12	TMI	Vertical hydrometeor profiles	TSDIS
3A-11	TMI	TMI Emission	TSDIS
1A-PR	PR	PR Level 1A Product	TSDIS
1B-21	PR	PR Total Power	TSDIS
1C-21	PR	Radar reflectivity	TSDIS
2A-21	PR	Surface reflectivity	TSDIS
2A-23	PR	PR Qualitative	TSDIS
2A-25	PR	Hydrometeor profiles	TSDIS
2B-31	Various	Combined radar/radiometer	TSDIS
3B-31	Various	Rainfall combined	TSDIS
3A-25	PR	PR Rainfall	TSDIS
3A-26	PR	Surface rain	TSDIS
1B-51	Radar Tapes	GT Calibration	TSDIS
1C-51	Radar Tapes	OC Reflectivities	TSDIS
2A-52	Radar Tapes	Measurable precipitation	TSDIS
2A-53	Radar Tapes	Rain map	TSDIS
2A-54	Radar Tapes	Convective map	TSDIS
2A-55	Radar Tapes	3-D Reflectivities	TSDIS
3A-53	Radar Tapes	5-day rainfall map	TSDIS
3A-54	Radar Tapes	Rainfall map	TSDIS
3A-55	Radar Tapes	Monthly 3-D Structure	TSDIS
GPI	Ancillary	Global Precipitation Index	GPCP
GPCP	Ancillary	Global Rainingage Data	GPCP
LIS00	LIS	Level 0 Data	TRMM
LIS01	LIS	Level-1A Product	TRMM
LIS02	LIS	Background Images Product	TRMM
LIS03	LIS	Events Product	TRMM
LIS04	LIS	Group Product	TRMM
LIS05	LIS	Flash Product	TRMM
LIS06	LIS	Area Product	TRMM
LIS07	LIS	Orbit Product	TRMM
LIS08	LIS	Vector Product	TRMM

**Table 3.4.2.2-2. Release A and B Datasets Held by the Data Server Subsystem
(2 of 2)**

Product	Instrument	Product Description	Platform
LIS09	LIS	Browse Product	TRMM
LIS10	LIS	Level-4 Product	TRMM
MOD18	MODIS	Water-Leaving Radiance	AM,PM
MOD19	MODIS	Pigment Conc	AM,PM
MOD20	MODIS	Chlorophyll Fluorescence	AM,PM
MOD21	MODIS	Chlorophyll_a Pigment Concentration	AM,PM
MOD22	MODIS	PAR	AM,PM
MOD23	MODIS	Suspended-Solids Conc, Ocean Water	AM,PM
MOD24	MODIS	Organic Matter Concentration	AM,PM
MOD25	MODIS	Coccolith Conc, Detached	AM,PM
MOD26	MODIS	Ocean Water Attenuation Coefficient	AM,PM
MOD27	MODIS	Ocean Productivity	AM,PM
MOD28	MODIS	Sea_sfc Temperature (SST)	AM,PM
MOD31	MODIS	Phycoerythrin Conc.	AM,PM
MOD32	MODIS	Calibration Data, MODIS	AM,PM
MOD33	MODIS	Gridded Snow Cover	AM,PM
MOD35	MODIS	Classification Marks	AM,PM
MOD36	MODIS	Absorption Coef. Total	AM,PM
MOD37	MODIS	Ocean Aerosol Properties	AM,PM
MOD38	MODIS	Water Vapor, Atmospheric (ThermalIR)	AM,PM
MOD39	MODIS	Clear Water Epsilon	AM,PM
DAS_FIRST_LOOK_ANALYSIS	DAO	First Look Analysis	

The key assumptions associated with the DBMS repository sizing are as follows:

- The products lists have been derived from the DAAC instrument teams representatives, the ECS Technical Baseline and in coordination with Science Metadata sizing for the SDSRV DBMS DRPHW CI.
- The period of data capture for Release A products on the TRMM mission is 8/17/97 to 12/31/98.
- The period of data capture for Release B is 1/1/1997 through 12/31/1999.
- All products are assumed to conform to the Proposed ECS Core Metadata Standard v2.0, 420-TP-001-005, Dec. 1994.
- The metadata sizing has been calculated from the Metadata Expected with each granule included in the table on Page 94 in the Proposed ECS Core Metadata Standard v2.0, 420-TP-001-005, Dec. 1994.
 - The calculated size of 1.823 KB per granule has been obtained from this data source.

- An overhead factor of 2.36 for implementation in Illustra has been estimated based on the benchmarking activities as outlined in the DBMS Benchmark Report, 430-TP-003-001. This assumes a high level of query / insertion activity and a low level of update / deletion activity.
- An estimated overhead of 80 MB will be made for the Illustra product code. Note that this includes sizing for the Illustra database product, the 2-D Spatial Data Blade, and the 3-D Spatial Data Blade.
- All instruments on TRMM are considered to exist in the Data Server as a continuation of Release A.
- All instruments on AM-1 were included in the product lists.
- Keyword metadata per document is derived from Release A, phase 2 DDSRV. Average size is 1727 bytes.
- Sizing for document metadata and storage is assumed during initial Release B operations because time phasing of document production is indeterminate.
- Sizing includes requirements for a RAID 5 overhead factor of 1.125.

Table 3.4.2.2-3. GSFC DBMS Repository Sizing (GB)

DAAC	1997	1998	1999	2000	2001	2002	EOC
GSFC	0.6	8.9	24.9	41.0	56.9	72.8	86.0

The calculated disk capacity for the GSFC repository (static analysis) results in a computed requirement of 24.9 GB. Due to the operational experience with the user space requirements, at least 5 GB of disk space must be allocated for a high use data base functioning. Therefore, at least 5 GB beyond the calculated storage requirement will be allocated - 29.9 GB. (The closest available quantity of disk equal or exceeding 29.9 GB will be purchased.) Dual host configuration will allow for failover. Host type will be a SGI class machine.

DOCUMENT DATA SERVER Document handling is handled via a dedicated Data Server implementation, geared to the predicted document ingest and access volumes and the nature of the COTS S/W requirements imposed on the support hardware. The Document Data Server is provided as a simple server configuration with network access. The following assumptions were made in the preparation of the GSFC Document Data Server configuration:

1. Documents and document metadata together have been considered as a basis for the sizing calculations.
2. Totals for MSFC are included in the GSFC totals.
3. Document related data includes: 5 guide document types, algorithm descriptions, production plans and reference papers. These from the DDSRV Detailed Design for ESDTs (DID 305).

4. Guide document sizes (5 guides) from Release A, phase 2 of DDSRV are sized at 1.5 MegaBytes each.
5. Production plans are assumed to be on per DAAC, sized at 1.0 MegaBytes each.
6. Reference papers are sized at 6.5 MegaBytes each; one per instrument per platform.
7. As a typical example of other documents accounted for, the CERES ATBD and LaRC Handbook file sizes were used as the basis for the ALGORITHM DOCUMENTATION sizing for each TRMM instrument. One ATBD is assumed per instrument.
8. The creation of ATBDs is based on the establishment of new data sets, therefore the accumulation of individual granules does not contribute to the number of ATBDs.
9. Sizing includes requirements for a RAID 5 overhead factor of 1.125.
10. The figures are approximations, which will be refined over time. The Document Data Server architecture is scaleable.
11. The Document Data Server continues to exist as a separate server. As the design effort for Release B continues relative to the Sybase/Illustra selection, the Document Data Server could be collapsed into the DBMS Server. This would change the document storage sizing and cause it to be added into the DBMS Server sizing.
12. The entire document Data Server DBMS sizing is assumed to be available at the beginning of the Release B time frame. This is to allow sufficient capacity to be available to handle both historical document conversion and new document requirements. This includes the maximum required by the end of Release B.

A 2 CPU SMP server was selected based upon operational experience with the EDF EDHS. A WAIS-like, full text indexer, an http server, and additional custom developed software will reside on this host. The following processes/applications run on this host: Document Data Server Process, WWW Server Process, Document Repository Process, Client Applications Process, CSS DCE client, CSS logging API, and an MSS agent.

The disk complement was sized to hold the document metadata for the data product collections associated with the TRMM and AM-1 missions and for the V0 data sets identified for migration. Sizing for document metadata was based on available V0 guide document sizing, and the 2.0 Core metadata baseline. The calculated required disk capacity for all document collection is 9.1 GB.

3.4.2.2.2 Configuration

The specific sizing for the Release B GSFC Data Servers, derived from the rationale described above, is synopsized below. Figures (EDS provided figures) provide the a preliminary design for the site's configuration. Additional details on specific component configurations and sizing are provided within the figures.

For the GSFC Science Data Server:

ACMHW

- Admin. Workstations: 2 ea. of SUN Sparc 20/50
- APC Hosts: one 16 CPU SGI Challenge XL, configured with 6 GB local disk and 1 GB RAM, and one 8 CPU Challenge XL, 512MB of RAM
- 764 GB RAID

WKSHW

- WKSHW hosts: 1 SGI Challenge XL with 10 CPUs and 1 SGI Challenge XL with 6 CPUs, 1 GB RAM each, 6 GB local disk each
- 558.5 GB RAID

DIPHW

- Staging Server Host: two 4 CPU SUN Ultra, each with 6 GB local disk, 256 MB RAM, and access to 480 GB RAID.
- Standard Release B Peripheral Set: 8mm Tape drives and stackers, 4mm Tape drives and stackers, 6250 Drive, CD-ROM drive and jukebox, FAX/scanner, 3480/3490 outboard drives.

DRPHW

- FSMS Server Hosts: three 6 CPU SGI Challenge XL with 6 GB local disk, 200 GB of AMASS Cache and 512 MB RAM
- Archive Tape Library Robotics: 5 STK Powderhorn
- Tape Drives: 16 D3 drives per four ATLS 12 D3 drives per one ATL
- Tape Media: 6,942 D3 50 GB cartridges
- DBMS Server: two of 2 CPU SGI Challenge XL, with 30 GB of shared disk

For the GSFC Document Data Server....

- WAIS/http Data Server 2 of 2 CPU SMP Server
- Data Server Disk: 12 GB mirrored in two machines for Release-B

3.4.2.3 Data Management Subsystem

The Data Management Subsystem (DMS) consists of a single Hardware CI (HWCI) that will also support the Release B processing requirements of the Interoperability Subsystem (IOS) at the GSFC DAAC site.

The DMS is responsible for supporting Advertising Service CI, Data Dictionary CI, Gateway CI, Local Information Manger CI and Distributed Information Manager CI processing activities generated directly from user "pull" search invocations. The DMGHW CI consists of three major components: 1) DBMS/Web Server, 2) Database Management Workstation, 3) Data Specialists and User Support Workstations.

The DBMS/Web server is the primary hardware component in the Data Management Subsystem. The server provides DBMS storage, input/output (I/O), and processing resources in support of the Advertising Service CI, Data Dictionary CI, Gateway CI, Local Information Manager CI and Distributed Information Manager CI in Release B.

The DMGHW CI configuration provided in Section 3.4.2.3.2 depicts the migration of the Release A hardware design to the Release B hardware design. The Release B hardware design accommodates Release B platform design issues concerning scalability, RMA and evolvability. The hardware design is tailored to Release B GSFC DAAC specific processing needs in support of Advertising Service CI, Data Dictionary CI, Gateway CI, Local Information Manger CI and Distributed Information Manager CI processing functions. The Release B GSFC DMGHW CI is designed to support TRMM VIRS, AM-1 MODIS and DAO and LIS mission datasets. Section 3.4.2.3.1 provides the rationale behind the recommended Release B hardware configuration and is subject to change as Data Management software CI (under investigation in the incremental development track) prototyping results become available.

3.4.2.3.1 Rationale

The performance impact on the DMGHW CI DBMS/Web server in Release A has been determined to be negligible. In order to avoid unnecessary swap-outs of hardware after Release A, the DMGHW CI has been sized according to Release B processing predictions; therefore, the DMGHW CI Release B design builds on the Release A design without the need for large scale changes. The performance drivers for sizing the DMGHW CI server for Release B are:

- User Characterization analysis of science and non-science user search invocations
- DBMS/CI transaction rate (performance) analysis
- DBMS/CI prototype/benchmark analysis
- Hardware Scalability / RMA / Evolvability Analysis

User Characterization Analysis: User Characterization data provides the projected number of science and non-science users, frequency of search invocations per time period, and the percentage of invocations for different types of searches to be supported in the Release B time frame. In Release B, it is expected that science users will primarily access the DM services (Gateway CI, Data Dictionary CI, Local Information Manager CI, Distributed Information Manager CI) while the bulk of the accesses to the Interoperability service (Advertising Service CI) will originate from within the non-science community.

It is assumed that non-science users will access the Advertising Service CI 86% of the time, and DM CIs 14% of the time as documented in "User Characterization and Requirements Analysis" (19400312TPW). The number of searches per hour being processed by the DM CIs in response to queries by science users is based on the ECS science user scenarios in which users are assumed to be accessing the system through the client. Because of the increasing popularity and ease of use of the WWW, it is also expected that science users will make use of the Advertising Service at a rate equal to 25% of DM searches. Data provided by the User Characterization Team apply to epoch m (first quarter of 2000) since the data are meant to represent maximum usage loading (this will occur at the end of Release B).

Tables 3.4.2.3-1 and 3.4.2.3-2 summarize the number of science user system accesses per day and the fraction of invocations per search type for DM services as documented by the User Characterization Team. Table 3.4.2.3-3 summarizes the total number of searches per hour for the busiest hour of the day.

Table 3.4.2.3-1. GSFC User System Accesses per Day for Science Users (Epoch m)

DAAC	User System Accesses per Day
GSFC	101

Table 3.4.2.3-2. GSFC Science User Search Types for Gateway CI Services (Epoch m)

Search Type	Fraction of total invocations
Simple Search/1 site	.263
Simple Search/multi-site	.279
Match-up Search/1 site	.272
Match-up Search/multi-site	.185
Coincident Search/1 site	0.0
Coincident Search/multi-site	.000374

Table 3.4.2.3-3. GSFC Searches per hour for Science Users (Epoch m)

DAAC	Searches per hour (busiest time of day)
GSFC	105

The data for searches submitted by science users is categorized into six different types (simple/1, simple/multi, match-up/1, match-up/multi, coincidence/1, and coincidence/multi) for DM services; however, there are only three types of searches for the non-science user data. Each of the three search types that exist in the non-science user data was subdivided into one-site vs. multi-site by applying the proportions of one-site vs. multi-site that exist for science users to the non-science user data. For example, the relative proportions of simple search/1 site and simple search/multi-site for science users is 0.263 and 0.279, respectively. The number of simple searches submitted by non-science users was divided into one-site and multi-site using these same proportions. Only one search type (simple/1) pertains to the Advertising Service CI.

Tables 3.4.2.3-4 and 3.4.2.3-5 summarize the number of non-science user system accesses per day and the fraction of search invocations for DM services as documented by the User Characterization Team.

Table 3.4.2.3-4. GSFC User System Accesses per Day for Non-Science Users (Epoch m)

DAAC	User System Accesses per Day
GSFC	5564

Table 3.4.2.3-5. GSFC Non-Science User Search Types for Gateway CI Service (Epoch m)

Search Type	Fraction of total invocations
Simple Search/1 site	.31
Simple Search/multi-site	.29
Match-up Search/1 site	.09
Match-up Search/multi-site	.06
Coincident Search/1 site	0.0
Coincident Search/multi-site	.25

DBMS Transaction Rate Analysis: In order to size the DBMS server it is necessary to estimate the size of the Interoperability (Advertising Service CI) and Data Management (Gateway CI, DDICT CI, LIMGR CI, DIMGR CI) services and then determine the transaction rates, or database throughput that must be provided in support of the "pull" search activities that will be invoked by the user community. The transaction rate analysis is based on assumptions regarding the amount of processing associated with the different types of search requests that pertain to the Interoperability and Data Management software CIs. Release B transaction assumptions were made to define a transaction loading value per search request. Number and type of search request are provided by the User Characterization Team. Depicted transaction loading values are assumptions that are based on search complexity. The loading values for search requests will be refined with actual performance benchmarks as future prototypes are completed. The observed transaction loading from future prototyping/benchmarking activities will be compared to the predicted ones (documented below) and the sizing analysis will be updated as a result (these transaction loading assumptions are defined as "nominal" cases). The transaction data provided is projected for the Release B time-frame.

The processing (transactions per search invocation) assumptions are based on preliminary transaction analysis results for the Advertising Service and DM CIs and will be revised based on future prototyping/benchmarking results as they become available. Tables 3.4.2.3-6 and 2.4.2.3-7 list the predicted transaction load associated with the Gateway CI and Advertising Service CI based on frequency of science user search invocations. Searches/hour are calculated for the busiest time of day at the GSFC DAAC site.

Table 3.4.2.3-6. GSFC Science User Transaction Analysis for Gateway CI Service (Epoch m)

Search Type	Percentage Invoked	Searches/hour	Processing Assumptions (Transactions per Search Type)	Transactions /hour
Simple Search/1 Site	26.3	28	5	140
Simple Search/Multi-Site	27.9	29	10	290
Match-up Search/1 Site	27.2	29	5	145
Match-up Search/Multi-Site	18.5	19	10	190
Coincident Search/1 Site	0.0	0	5	0
Coincident Search/Multi-Site	.0374	.04	25	1

Table 3.4.2.3-7. GSFC Science User Transaction Analysis for Advertising Service CI (Epoch m)

Search Type	Percentage Invoked	Searches/hour	Processing Assumptions (Transactions per Search Type)	Transactions /hour
Simple Search/1 Site	100	26	5	130

Tables 3.4.2.3-8 and 3.4.2.3-9 list the transaction analysis for the Gateway CI and Advertising Service CI based on daily non-science user accesses as depicted in "User Characterization and Requirements Analysis" (19400312TPW). The estimated total non-science user system accesses per day for the Gateway CI and the Advertising Service CI is estimated to be 5564. The percentage that each search type, pertaining to the Gateway CI and Advertising Service CI, is invoked is also taken from the same document. Tables 3.4.2.3-8 and 3.4.2.3-9 are completed with the assumption that there will be at least ten search invocations per non-science user access on average.

Table 3.4.2.3-8. GSFC Non-Science User Transaction Analysis for Gateway CI Service (Epoch m)

Search Type	Percentage Invoked	Searches/hour	Processing Assumptions (Transactions per Search Type)	Transactions /hour
Simple Search/1 Site	31	101	5	505
Simple Search/Multi-Site	29	94	10	940
Match-up Search/1 Site	09	29	5	145
Match-up Search/Multi-Site	06	19	10	190
Coincident Search/1 Site	0.0	0	5	0
Coincident Search/Multi-Site	25	81	25	2025

Table 3.4.2.3-9. GSFC Non-Science User Transaction Analysis for Advertising Service CI (Epoch m)

Search Type	Percentage Invoked	Searches/hour	Processing Assumptions (Transactions per Search Type)	Transactions /hour
Simple Search/1 Site	100	1994	5	9970

Preliminary transaction analysis results for Interoperability and Data Management CI processes are depicted in Table 3.4.2.3-10. At this time the Data Dictionary CI transaction analysis is equated to the transaction analysis of the Advertising Service CI since the type and frequency of transactions are predicted to be very similar. Due to the fact that DIMGR and LIMGR CI processes are perceived as being the most expensive (in terms of cost to the CPU), and that prototypes will not be developed until after Release B CDR, preliminary sizing estimates have been achieved by doubling the transaction load of the Gateway CI and applying the result to LIMGR CI and DIMGR CI processes. Although equating Advertising Service CI and Data Dictionary CI processes, and doubling the Gateway CI transaction load to produce LIMGR CI and DIMGR CI transaction loading does not pin-point the performance cost that will be levied on the DMGHW CI, it does provide preliminary, expected CPU activity in the absence of real-time prototyping/benchmarking results. The preliminary transaction results for the Advertising Service CI, Gateway CI, Data Dictionary CI, LIMGR CI and DIMGR CI will be revised based on future prototyping/benchmarking analysis results as the Incremental Track software design matures.

Table 3.4.2.3-10 summarizes the science, and non-science user transaction loading per hour for the Advertising Service CI, Gateway CI, Data Dictionary CI, LIMGR CI and DIMGR CI services.

Table 3.4.2.3-10. DBMS Transaction Analysis Summary (Epoch m)

DAAC	User Type	Service	Searches/hour	Transactions /hour	TPM
GSFC	Science	Gateway	105	766	13
GSFC	Science	Advertising	26	130	2
GSFC	Science	Data Dictionary	26	130	3
GSFC	Science	LIMGR	210	1532	26
GSFC	Science	DIMGR	210	1532	26
GSFC	Non-Science	Gateway	324	3805	63
GSFC	Non-Science	Advertising	1994	9970	166
GSFC	Non-Science	Data Dictionary	1994	9970	166
GSFC	Non-Science	LIMGR	648	7610	127
GSFC	Non-Science	DIMGR	648	7610	127
Totals:			6193	43055	719

A sensitivity analysis has been performed with larger loading allocations; the results are depicted below in Table 3.4.2.3-11.

**Table 3.4.2.3-11. DBMS Transaction Sensitivity Analysis Results (Epoch m)
(loading has been doubled for both search invocations and transaction rates)**

DAAC	User Type	Service	Searches/hour	Transactions /hour	TPM
GSFC	Science	Gateway	210	3064	51
GSFC	Science	Advertising	52	520	9
GSFC	Science	Data Dictionary	52	520	9
GSFC	Science	LIMGR	420	6128	102
GSFC	Science	DIMGR	420	6128	102
GSFC	Non-Science	Gateway	648	15220	254
GSFC	Non-Science	Advertising	3988	39880	665
GSFC	Non-Science	Data Dictionary	3988	39880	665
GSFC	Non-Science	LIMGR	1296	30440	507
GSFC	Non-Science	DIMGR	1296	30440	507
Totals:			12370	172220	2871

Future Incremental Track Development prototyping/benchmarking activities will provide a more detailed performance analysis of Advertising Service CI, Gateway CI, Data Dictionary CI, Local Information Manager CI and Distributed Information Manager CI processes; therefore, performance transaction analyses will be revised accordingly.

DBMS Prototyping/Benchmarking Analysis: Currently, preliminary Incremental Track Development performance data is being used to size the processing capacity of the DMGHW CI DBMS/Web server. Performance analysis results will be revised as planned prototyping/benchmarking activities are completed. Major prototyping activities that will affect performance estimates for the DMGHW CI include, but are not limited to: 1) Prototype workshop 2, 2) EP7 prototype, July 1996.

DBMS performance estimates provided in "DBMS Benchmark Report" technical paper (430-TP-003-001), show that for multi-user (32 users) queries (20 similar queries accessing different parts of the test database) running concurrently, the test-bed platform's CPU became saturated (SUN SPARCstation 20/50). A vendor supplied TPM benchmark for the selected platform (HP K400) for Release B operations is shown in Table 3.4.2.3-12. As a rule vendor supplied Transaction Per Second/Minute (TPS/TPM) ratings tend to be a maximum, or high-end value and do not take into account processing overhead associated with other system processes. Processes that will run on the DMGHW CI in Release B include DCE client, MSS agent, HTTP server, Sybase SQL Server, Sybase Replication Server, Sybase Backup Server, Operating System Services, Advertising Service Server, Gateway Server, Data Dictionary Server, LIMGR Server and DIMGR Server.

Table 3.4.2.3-12. Vendor Platform Performance Estimates

Platform	TPM	MIPS
HP K400 (SMP) with 1 processor (PA 7200 CPU)	1000	146

Disk Capacity Sizing: Disk storage for the DMGHW CI has been determined for each DAAC site based on preliminary Interoperability and Data Management CI DBMS application sizing estimates plus vendor inputs for the following COTS software: 1) DBMS software, 2) Development software, 3) HTTP server software, 4) Operating System software, 5) Communications and Utilities software. Capacity sizing for the Interoperability and Data Management databases was achieved by multiplying the expected byte size for core and collection specific attribute definitions by the total number of core and collection specific attributes. Temporary workspace has also been allocated for Interoperability CI and Data Management CI services dependent on frequency and variation of queries. For example, the DIMGR service has a larger capacity requirement than the Data Dictionary service because it requires more temporary workspace since it handles a greater number and more varied types of queries. The expected capacity of Interoperability CI, Data Management CI, COTS, Operating System and Communications and Utilities software to be installed on the DMGHW CI at the Release B GSFC DAAC site is depicted in Table 3.4.2.3-13.

Table 3.4.2.3-13 is filled with preliminary disk sizing results for Release B Interoperability and Data Management software CIs, operational databases and COTS software packages that will be installed on the DMGHW CI at the GSFC DAAC site. Some of the results are estimates (such as database sizes) since the DBMS design will mature and impact disk capacity sizing as the Advertising Service CI, Gateway CI, Data Dictionary CI, Local Information Manager CI, and Distributed Information Manager CI under-go future evaluation and prototyping.

Table 3.4.2.3-13. DMGHW CI Disk Capacity Requirements

SW Component	Release B Capacity
COTS Software:	
Sybase System	300 MB
HTTP Server	10 MB
	Total: 310 MB
Databases:	
Sybase Master Database	3 MB
Sybase Tempdb Database	100 MB
Sybase Model Database	2 MB
Advertising Database	150 MB (Estimate)
Advertising DB Workspace	150 MB (Estimate)
Advertising DB Log	100 MB (Estimate)
Advertising HTML Files	100 MB (Estimate)
Data Dictionary Database	400 MB (Estimate)
Data Dictionary DB log	100 MB (Estimate)
DMS Working Store Database	500 MB (Estimate)
DMS Working Store DB log	100 MB (Estimate)
	Total: 1705 MB
Operating System & Utilities:	
Operating System Software	700 MB
Utilities	200 MB (Estimate)
DCE Client	46 MB
	Total: 946 MB
	Total: 2961 MB

3.4.2.3.2 Configuration

The selected DMGHW CI DBMS/Web server to be implemented in Release B is a low-end SMP server (HP K400) that is scaleable from one to four processors. A single physical DBMS/Web server will be implemented at the Release B GSFC DAAC site. At this time a two CPU configuration has been determined to be appropriate for the DMGHW CI DBMS/Web server due to RMA requirements. In Release A the configuration called for a primary, or "active" server and secondary, or "standby" server, each configured with a single CPU. For Release B, the configuration will be changed to a single host server configuration with redundant CPUs (see Figure 3.4.2-2). Note that for Release A an HP K200 low-end SMP server (primary / active) was selected for the DMGHW CI platform. In order to accommodate redundant network and FWD SCSI-2 interface cards and a redundant power supply unit at Release B the HP K200 will be upgraded to an HP K400. The upgrade will be completed to obtain the additional slots for the network and FWD SCSI-2 cards and power supply only, there is no difference in the scalability, or type of CPUs between the K200 and K400; therefore, the upgrade does not consist of "swapping out" hardware; the same physical servers that were used in Release A DM configurations will be upgraded. Also, since a single server is called for in the Release B design and not redundant servers, the redundant (secondary) K200 servers located at Release A DAACs (GSFC, GSFC, EDC) will be upgraded to K400 servers and re-located to Release B DAACs as primary servers.

The Release B HP K400 single server configuration will offer redundancy in the form of: 1) dual processors, 2) dual power supply units, 3) dual FWD SCSI-2 cards, and 4) dual FDDI network cards 5) duplicate OS boot/application disk. The components listed above are hot swap-able units which allow them to be replaced without shutting down the server. Also, the HP-UX operating system features memory page de-allocation which automatically blocks out any portion of memory in which an error has been detected; therefore, a failure to memory will not bring operations to a halt. The single server host configuration will allow applications to be run in parallel across both processors, which enhances load balancing and availability/recovery capabilities. The DMGHW CI DBMS/Web server will automatically reconfigure itself in a single CPU configuration in the event of failure to a single CPU. The dual power supply units, FWD SCSI-2 cards and FDDI network cards will also provide continued availability in the case of failure to a single component (per function). The redundant configuration of the DMGHW CI DBMS/Web server has been analyzed by the ECS Reliability Engineering Group using COTS vendor provided data to ensure that all functional availability requirements are met. Preliminary analysis results revealed that the DMGHW CI DBMS/Web server has an MTBF (mean-time-between-failures) of greater than 20,000 hours which meets all pertinent RMA requirements.

A RAID disk unit will be attached via FWD SCSI-2 (dual ported) to the DMGHW CI DBMS/Web server (see Figure 3.4.2-2). The RAID disk unit will provide operational and mirrored sets of disk devices to the DMGHW CI server in order to provide uninterrupted data availability in the event of a disk failure. The RAID disk configuration will also be implemented such that a failure to a single disk will be recoverable via a "hot-swap" disk capability. The RAID disk unit will be comprised of ten disks: five operational and five mirror disks. Each disk contained in the RAID unit will provide a capacity of 2.1GB.

Since the Release B processing requirements for the Local Information Manager CI and Distributed Information Manager CI are largely unknown at this time, the flexibility of the recommended hardware design assures minimum risk. The design allows for 100% growth in both processing and storage capacity in support of the Advertising Service, Gateway, Data Dictionary, Local Information Manager and Distributed Information Manager CIs in Release B.

The following configuration diagram, Figure 3.4.2-2, depicts the recommended server host / RAID disk unit configuration:




Figure 3.4.2-2. DMGHW CI RAID Disk Configuration

Although the estimates for Advertising Service and Data Management operational databases, COTS software, and operating system and utilities sizing are relatively small, the total disk volume for the DMGHW CI server has been increased in support of Sybase swap (workspace) area, 100% growth capacity for the core operational software, and mirror disk units. The Data Management databases will be replicated (using Sybase replication server) at each DAAC site. The Interoperability and Data Management software CIs will be distributed over multiple disk

drives contained in the RAID disk unit in order to enhance performance (vendor and developer recommended); therefore, the total disk volume is well above the actual capacity needed to support the software alone. Additional disk capacity (internal to the server) in the form of a redundant system/application disk, and a test/development disk have also been added to the configuration in support of RMA, and integration and test requirements.

A single 8mm tape drive unit will be configured on the DMGHW CI server in Release B. The 8mm tape drives will be used to backup Advertising Service and Data Management CIs/databases, as well as perform DBA and routine maintenance operations.

A low-end DBMS/DBA uniprocessor workstation will be used for database/system administration activities. A single 8mm tape drive unit will be configured on the DBMS/DBA workstation in Release B. The 8mm tape drive will perform backup/recovery and routine maintenance operations in support of the DMGHW CI DBMS/Web host servers. A small pool of low-end uniprocessor workstations will support Data Specialist/User Support operations. At a minimum the DBMS/DBA and Data Specialist/User Support workstations will be configured with six gigabytes of local disk each. Workstation disk capacities were sized based on IR-1 workstation installation results which included recommendations for additional space per workstation to accommodate the following: 1) User/Temp workspace, 2) Personal development space, 3) Testing, 4) Software upgrades, 5) Working with large files/datasets, 6) future growth/flexibility. Exact capacities for disk drives are dependent on the procurement process as the type, and size of disks being offered for workstation platforms may fluctuate.

Table 3.4.2.3-14 summarizes the recommended DMGHW CI processing configuration for implementation at the Release B GSFC DAAC.

NOTE: The HP K400 SMP depicted in Table 3.4.2.3-14 is a low-end SMP class server that is scaleable from 1-4 CPUs.

Table 3.4.2.3-14. GSFC DAAC Hardware Configuration

Component	Class/Type	Platform	Qty.	Number of Processors	Memory	Disk Capacity
DBMS/Web Server	SMP	HP K400	1	2	512 MB	27.3 GB
DBA Workstation	Uniprocessor	HP 715/64	1	1	64 MB	6 GB
Data Specialist and User Support Workstations	Uniprocessor	SUN Sparc 20/50	5	1 (each)	64 MB (each)	6 GB (each)

DBMS/Web Server Platform Technical Specifications:

Make: Hewlett Packard

Model: K400 (Low-End SMP class server)

CPU: PA7200 (upgradeable to future PA800 processor)

Clock Frequency: 100 MHz

Number of processors: 1 to 4

MIPS: 146 (1 processor)

TPM: 1000 (1 processor) 3160 (4 processors)

SPECint92: 136 (1 processor)

SPECfp92: 215 (1 processor)

Memory: Expandable to 2GB RAM

Internal Processor-Memory bus bandwidth: 960 MB/sec (peak)

I/O Bandwidth: 128 MB/sec (peak)

3.4.2.4.1 Rationale

The sizing of Ingest Subsystem hardware both from a system level and a component level is based on the 2/7/96 version of the ECS Technical Baseline. Among the information included in the baseline is:

- o data by instrument,
- o average daily data volume by level,
- o and data destination.

Table 3.4.2.4-1 provides a synopsis of the Ingest data volumes required for Release B at GSFC.

Table 3.4.2.4-1. GSFC Ingest Volumes - Release B

Release A DAAC	Daily L0 Volume (GB/day)	Annual L0 Volume (GB/year)	Version 0 Data Migration Volume (GB)
GSFC	70.4	25,696	2,896

The average expected daily and annual data volumes at each site were calculated from this information and used to determine the required ingest hardware capabilities. Ingest client hosts are sized to accommodate the required ingest volumes as well as I/O and CPU capabilities to support internal data transfers associated with metadata validation and extraction and transfer of data to the Data Server or Processing Subsystems. Working storage disks are sized to accommodate the above functions, as well as provide contingency space for the transfer of more than one days worth of data within a 24-hour period. Since high RMA is a driver for the Ingest Subsystem, all critical components also include some type of sparing or redundancy to ensure that availability requirements are met.

Queuing Analysis Model An Ingest Queuing Model (Imodel) was developed to assist in the sizing of Ingest Subsystem components for the Release A GSFC configuration. This model was also used in a preliminary assessment of the Release B volumes. However, the additional interfaces and increased data volumes to be supported in Release B require the use of a dynamic

model to better depict the flow of data into, within, and out of the Ingest Subsystem. A more detailed and comprehensive modeling effort will be completed prior to the Release B CDR. The configuration presented in the following section provides an estimate of the types and classes of hardware required at the DAAC based on the basic paper analysis and queuing model estimates.

3.4.2.4.2 Configuration

The GSFC configuration for Release B is sized for the AM-1 and COLOR (FOO) missions for one year of operations. Brief descriptions of the generic components, provided within the GSFC configuration at Release B are provided in Table 3.4.2.4-3. The results of the system design analysis results in a recommended configuration consisting of mid-level Symmetric Multi-Processing (SMP) 32-bit machines, capable of supporting multiple network (FDDI and/or HiPPI) and direct-connect (SCSI II) devices.

Working storage devices are RAID 5 units with a minimum of two days worth of space allocated to ingest working storage required to support the functions of acquiring, processing, validation, and archiving L0 data. This volume of working storage allows for one days worth of L0 data to be staged for processing, an additional days worth available for subsequent ingest, and an additional 25% available to service additional Ingest Subsystem needs (e.g., retrieval support, pre-processing, quality checking). Additional magnetic disk resources are supplied within the Ingest Subsystem to support items such as:

- o client host operating system
- o application software and L0 archive database directory information

Table 3.4.2.4-2. GSFC Ingest Components at Release B

Component Name	Class/Type	Comments
Client Host	SMP Server W/S	Mid-level SMP servers.
Working Storage	RAID disk	Multiple RAID units, site capacity sized. RAID 5 for random access protocol.
L0 Archive Repository	Archive robotics and tape drives	Single robotics unit with multiple tape drives.

The specific sizing derived from GSFC Release B requirements is synopsisized within Table 3.4.2.4-3 and is highlighted in Section 3.4.2 within the site configuration overview Figure 3.4.2-1. The site overview figure provides additional details on specific component configurations and sizing. The Ingest Subsystem resources at GSFC are sized principally for the operational ingest and storage of Level 0 data sets, and to support the migration of data products from Version 0. Additional available system resources are for contingency.

Table 3.4.2.4-3. GSFC Ingest Sizing - Release B

Ingest Component	Component Class	Quantity	Comments
Client Host (ICLHW)	SGI Challenge L SMP	2	L0 Ingest Client hosts. Hosts are adapted to ECOM I/F and ESN. Host attached disk. SCSI I/Fs to RAID working storage.
Working Storage (ICLHW)	RAID Disk (host attached)	2	Host adapted RAID disk arrays. RAID 5. SCSI II adapted & cross strapped to Ingest Client hosts.
L0 Archive Repository (ICLHW)	Archive robotics	1	Robotics unit sized to store one year's worth of Level 0 data (plus a small amount for long term storage of L0 data without L1A counterparts).
L0 Archive Repository (ICLHW)	Archive tape drives	2	Drives mounted in archive robotics unit to support writing and retrieval of L0 data.
Client Host (ICLHW)	8 mm tape stacker	2	Support for hard media ingest from Version 0.
Client Host (ICLHW)	X-Terminal	1	OPS support for Data Ingest Technician(s).

3.4.2.5 Interoperability Subsystem

For the Release B time frame, the hardware support for the Interoperability Subsystem, particularly the Advertising capabilities are provided by the Data Management HWCI. Please see Section 3.4.2.3 for a complete description of this capability.

3.4.2.6 Production Planning Subsystem

The Release B SDPS Planning Subsystem Design Specification, 305-CD-026-002, Section 5, describes the functionality and architecture of the Planning Subsystem (PLS) Hardware Configuration Item (PLNHW CI). Because of the close coupling of scheduling/queuing functions with planning (i.e., use of the PDPS database), Section 5, also addresses the scheduling/queuing hardware within the Data Processing Subsystem (DPS) Science Processing Hardware Configuration Item (SPRHW CI).

This section describes the sizing of the Planning Subsystem (PLS) Hardware Configuration Item (PLNHW CI) for the GSFC DAAC site. Because of the close coupling of scheduling/queuing functions with planning (i.e., use of the PDPS database), this section also addresses the scheduling/queuing hardware within the Data Processing Subsystem (DPS) Science Processing Hardware Configuration Item (SPRHW CI).

The GSFC PLNHW CI consists of two Production Planning/Management workstations running the Planning Workbench software and one SMP Server running the PDPS database. A Production Planning/Management workstation supports the planning operations staff in performing their routine production planning and management functions. A workstation is provided for operations personnel to access production planning GUIs via the Planning Workbench application. These functions include candidate plan creation, plan activation, entry of production request information and report generation.

The Planning Subsystem is responsible for maintaining the Planning and Data Processing (PDPS) database. This database serves both planning and scheduling. Planning uses the database in support of the Planning Workbench application used for production planning and resource planning activities. Scheduling uses the database in support of the AutoSys application which runs the production schedule.

The GSFC scheduling/queuing hardware in the SPRHW CI consists of a Queuing Server which runs AutoSys and the AutoSys database. Production Monitor Stations run the AutoSys AutoXpert display GUIs.

The scheduling/queuing portion of the DPS is responsible for maintaining the AutoSys database. An active plan is loaded into the AutoSys database from the Planning Subsystem. AutoSys then proceeds to execute this production schedule (the result of activating this plan). AutoSys uses both the PDPS database and its own database during its execution of the production schedule.

Both the PDPS and AutoSys databases are implemented using Sybase. The PDPS database is used for the persistent storage of critical data while the AutoSys database contains the current data for executing a daily production schedule.

3.4.2.6.1 Analysis

Hardware Design Drivers

Drivers for sizing the hardware to support these functions include:

- o database transaction rates,
- o replication of critical data,
- o volume of production status to display,
- o job throughput,
- o the need for a fail-soft science processing architecture.

The primary drivers for sizing the PDPS database are data criticality, support of science processing's fail-soft environment, and the volume of database transactions. The PDPS data is critical because it consists of invaluable PGE information and run-time statistics used for the creation of future plans. The PDPS database supports AutoSys in execution of the production schedule. If the database fails, the launching of production jobs by AutoSys will stop. The PDPS database also needs to support multiple database accesses from both planning and processing.

The primary drivers for sizing the AutoSys database are support of science processing's fail-soft environment, job throughput, and the display of production status. AutoSys launches jobs which stage data, run PGEs, report PGE status, and destage data. If AutoSys fails, no new production jobs will be executed and production will stop. AutoSys must also be able to execute enough jobs (throughput) to process the daily production schedule within the allocated time frame (8 or 24 hours which varies by DAAC). AutoSys' production status also needs to be displayed without saturating the person monitoring production.

Assumptions

The sizing of the planning and scheduling hardware is based on an analysis of the PDPS database transactions and Planning Workbench functions, AutoSys job throughput, and AutoXpert display requirements. The following paragraphs summarize the assumptions used in this analysis.

1. *Technical Baseline* -- The February 1996 technical baseline information was used to derive key parameters that effect the planning and processing workload at each DAAC site. These key parameters include the number of job activations per day for each PGE at each DAAC and the number of PGEs maintained at each DAAC. For example, at GSFC during the Release B period (Epoch k, 3Q99), the following parameters were derived:

- No. PGEs at the DAAC	62
- Number of PGE activations per day. Technical baseline	19,322*
- Number of PGE activations per day after PGE combining	4,004
- Total # of jobs required (tech. baseline) with reprocessing of 100%	154,656*
- Total # of jobs required (PGE combining), reprocessing of 100%	32,032

GSFC hardware is sized to handle 4,004 PGE activations per day. The reduction from the technical baseline's PGE activations per day number of 19,322 was the result of a meeting with the MODIS SDST (Science Data Support Team) which concluded that the number of MODIS PGE execution events could be reduced to less than 4,000 per day. This reduction will be accomplished by:

- a) combining multiple PGEs into one job box; and
- b) scheduling multiple granules to be processed through a set of Level 1 and 2 PGE's as a single job box.

This has the effect of reducing the number of jobs required from 154,656 to 32,032.

2. *OPS Concept* -- These are the operational assumptions.

- a. The production scheduler/planner at each DAAC will prepare and publish a 30 day plan every two weeks. This plan is used to provide some assistance in longer range planning. The 30 day plan is only prepared and published every two weeks - if changes occur after a plan is published, the changes will only be incorporated in the next 30 day plan.
- b. The production scheduler/planner at each DAAC will prepare and publish a 10 day plan every week. This plan will provide a finer grain description of planned activities. Like the 30 day plan, this plan will not be replanned and distributed except on the regular weekly boundaries.
- c. The production scheduler/planner at each DAAC will prepare and activate a daily plan or schedule once per day. This schedule will be replanned as required during the course of the day.

It needs to be noted in this regard, and as described in more detail in the Planning Subsystem design document, that the manner in which each DAAC operations production scheduler/planner decides to conduct planning and scheduling can vary from DAAC to DAAC. The planning and scheduling tools are sufficiently flexible to support a variety of planning and scheduling strategies. The sizing of the storage and the performance of the server is based upon the assumptions given here. It should also be noted that it is not essential that the operations staff conduct replanning for the active schedule if events arise (e.g., processor failures) that would cause the predictions of the active schedule to depart from reality. Processing will continue regardless since jobs will be released for execution as resources (processors and input data) become available.

3. It was assumed that the reprocessing workload is equal to (1x) the standard processing workload in the Epoch k time frame and that the reprocessing workload is two times (2x) the standard processing workload in the Epoch o time frame (near the end of Release B). The hardware selected for GSFC is sized to support 1x reprocessing.
4. Several other parameters were estimated as part of the process of developing a model for the database sizing and processing activity. Some of these parameters and their values are:

- No. of working hours per day	24
- No. of long term (30 day) plans stored	5
- No. of short term plans stored	5
- No. of production requests per PGE	4
- No. of replans (active plans) per day	2
5. *Planning Activities* -- The design for the Planning Subsystem was represented as a collection of several activities or functions (e.g., Subscription Manager, Planning Workbench: Candidate Plan Creation, etc. See Planning Subsystem Design). Functions that were activated only infrequently (e.g., Subscription Submittal) were ignored. The activation frequency of these subsystem activities was then identified, where possible from the February 1996 technical baseline summary (e.g., number of jobs activations per

day). Subactivities within these functions were identified and an estimate of the processing load was made based upon the complexity of the activity. These processing load estimates were made in terms of the number of transactions (updates, inserts, and reads) made to the PDPS database. These load estimates were then totaled across the complete set of Planning Subsystem activities for the specified frequency of activation for each DAAC site. The PDPS DBMS server sizing was then estimated by scaling upwards from the base processor size.

PDPS Transaction Rate Analysis

Database transaction rates were modeled for the PDPS database by estimating the number of updates, inserts, and reads made to the database when performing planning and scheduling functions. These functions include:

- o Creating Data Processing Requests to be added to AutoSys database (DPRs),
- o Release of DPR jobs for execution,
- o Updating Data Processing Request Job,
- o Getting DPR Job Status,
- o Data Initialization,
- o Local Data Management,
- o Data Staging,
- o Data Destaging,
- o Initiate PGE execution,
- o Monitor Execution,
- o Execute Post Processing,
- o Plan Creation/Activation,
- o Subscription Notifications, and
- o On-Demand Production Requests.

For each of the activities listed above, the number of average daily transactions (updates, inserts, and reads) made to the PDPS database was estimated based on the number of PGE activations per day. The average daily transaction load at GSFC based on 4,004 PGE activations per day is 6,820,000 (including Sybase replication). This would require 79 transactions per second (TPS). The peak transaction load is 9,009,000 worst case, for a 30-day plan. The transaction rate of a 167 MHz 4-CPU Sun Server is estimated at 1,360 TPS. It would therefore take approximately 111 minutes to create a 30-day plan containing 4,004 PGE executions at GSFC. A 10 day plan containing 4,004 PGE executions would take approximately 25 minutes and a 1 day plan would take approximately 8 minutes.

PDPS Database Sizing

The PDPS DBMS server supports the PDPS database that contains all the information central to the functioning of the Planning Subsystem. To size the PDPS DBMS server it was necessary to estimate the number of inserts made into the database in order to determine how quickly and by how much the database would grow each day.

The database classes, associated with the PDPS Transaction Rate Analysis functions above, are:

- o DpPrExecutable,
- o DpPrDataMap,
- o DpPrPGE,
- o PlProductionRequest,
- o PIDpr,
- o PlDataGranule,
- o PIODProdRequest.

The number of inserts made to the database are from the PDPS Transaction Rate Analysis. The number of inserts (additional rows added to the Sybase tables) were then input into a Sybase system procedure called sp_estspace. This procedure returns the estimated megabytes of disk space required to store a particular number of rows for a given class (table) above.

The number of PGE activations for GSFC is 4,004. The disk capacity required to store 90 days worth of PDPS database growth is estimated to be 4 GB

Planning hardware

The Production Planning Subsystem is made up of two 20/71 workstations and one Sun UltraServer with 4 CPUs. The UltraServer is attached to two 10.5 GB Sun SPARCstorage MultiPacks (external disk) via fast wide SCSI controllers. Each workstation will run the Planning Workbench software. Both workstations can use the Planning Workbench, one for Production Planning, the other for Resource Planning. The UltraServer will run Sybase (the PDPS DBMS). Sybase will run on this machine while Sybase replication is run on the Queuing Server, thus creating a copy of the PDPS database and allowing automatic switch over to the replicated database upon a failure of the primary PDPS database. Sybase replication ensures the integrity of the critical data that the PDPS database contains.

The Sun operating system and applications will reside on disk internal to the machines. The Sybase log, index, and data files will be on external disks. A set of five 2.1 GB external disks will be accessible via 2 fast wide SCSI controller cards.

Each workstation can accommodate 2 CPUs and use faster emerging hyperSPARC technology CPUs running at 150 MHz. The UltraServer can accommodate up to 12 CPUs. Future experience with the Planning Workbench software and the PDPS Database will determine if this extra capacity will be needed.

Each SPARCstorage MultiPack disk can be increased to a maximum of 12 disks for a total of 25.2 GB if required.

AutoXpert Display Sizing

Three pairs of X-terminals will be used to show production status by using the AutoSys AutoCons' display and the AutoXpert JobScape GUI. Additional AutoXpert GUIs: TimeScape and HostScape are also available. If an operator wants to view the progress of MODIS production, he would modify his view of the entire active queue to display only those jobs associated with MODIS. To properly do this requires two displays, one for AutoCons and another for JobScape. The other two pair of displays can be used to display TimeScape and HostScape for the same instrument's production status or another instrument's production could be viewed or maybe the progress of reprocessing jobs could be displayed. Table 3.4.2.6-1 shows the contribution that each instrument's production (PGE activations) makes on total daily production at GSFC. Reprocessing at Epoch k (3Q99) will be approximately 100% of normal production growing to 200% of normal production by the end of Release B.

Table 3.4.2.6-1. Breakdown of instrument production at GSFC

Instrument	# of instrument PGE activations per day / total # of all PGE activations per day
DAO	.05%
LIS	.05%
MODIS	99.9%

The RAM requirements for each of the AutoXpert and AutoCons GUIs can be calculated using the formula:

$$20,000 \text{ bytes} + (5.5 \text{ bytes} \times \# \text{ of jobs loaded}) = \text{amount of RAM required.}$$

To open all 4 GUIs when the entire days' processing schedule is loaded would require:

$$[20,000 \text{ kbytes} + (5.5 \text{ kbytes} \times 32,032 \text{ jobs})] \times 4 \text{ GUIs} = 785 \text{ MB of RAM minimum.}$$

The 32,032 jobs is derived by multiplying the number of PGE activations at GSFC for a 24 hour day (4,004) by the number of jobs it takes AutoSys to activate a PGE (4) plus an additional 100% increase in job activations due to reprocessing (x2),

$$4,004 \text{ PGE activations} \times 4 \text{ jobs per PGE} \times 2 = 32,032 \text{ job activations per day.}$$

AutoSys Database Sizing

The size of the AutoSys Database is dependent upon the number of jobs loaded into AutoSys at one time. The database, when loaded with a test of 20,000 jobs, consumed 6 MB of disk space. This number would vary as a function of the complexity of job dependencies.

The AutoSys events database will also require disk space. The events database saves information (starting and run times, success or failure of job, alarms, etc.) about the job activations launched by AutoSys. Estimates show that an additional 180 MB of disk space would be needed to archive 14 days worth of event information for GSFC.

AutoSys Job Throughput Analysis

Experience by other AutoSys customers has indicated that AutoSys can sustain 45 jobs per minute while running on a 75 MHz Sun SPARC 20/712 2-CPU machine. Operating in our high availability configuration reduces performance by approximately 35%. Therefore, AutoSys could execute a sustained 29.25 jobs per minute when operating in its Dual Data Server Mode (high availability option). It has been estimated that a 4-CPU Sun UltraServer machine could launch a maximum of 187,200 jobs for processing within a 24 hour period.

The predicted number of jobs for one 24 hour period at GSFC for Epoch k (3Q99) including an additional 100% increase in job activations due to reprocessing is 32,032 based on the February 1996 technical baseline.

Scheduling hardware

The scheduling function of the Data Processing Subsystem is done by a Sun UltraServer with 4 CPUs. The UltraServer is attached to two 10.5 GB Sun SPARCstorage MultiPacks (external disk) via fast wide SCSI controllers. Three pairs of X-terminals will be used to show production status by using the AutoSys AutoCons' display and the AutoXpert GUIs. AutoSys software runs its Event Processor (using Sybase) on the Sun UltraServer machine while the AutoSys Shadow Event Processor runs on the PDPS DBMS Server. When the Shadow Event Processor detects that the Event Processor has malfunctioned, the Shadow Event Processor takes charge of the execution of the schedule, thereby allowing the continued launching of production jobs. AutoSys will also be run in Dual Data Server Mode which allows AutoSys to replicate its database on separate disks. When a disk or database failure occurs on the primary database, AutoSys continues to operate, using the backup database.

The Sun operating system and applications will reside on disk internal to the machine. The Sybase log, index, and data files will be on external disks. A set of five 2.1 GB external disks will be accessible via 2 fast wide SCSI controller cards.

The UltraServer can accommodate up to 12 CPUs. Future experience with AutoSys and its Sybase database will determine if this extra capacity will be needed.

Each SPARCstorage MultiPack disk can be increased to a maximum of 12 disks for a total of 25.2 GB if required.

3.4.2.6.2 Configuration

- PDPS DBMS Server:
 - Sun UltraServer
 - Clock: 167 MHz
 - CPUs: 4 of 12
 - Disk: 4.2 GB
 - RAM: 512 MB per CPU
 - Fast wide SCSI control: 2
- Production Planning/Management Workstation: Quantity: 2

- Sun SPARC 20/71
- Clock: 75 MHz
- CPUs: 1
- Disk: 4.2 GB
- RAM: 512 MB
- SPECint92 = 125, SPECfp92 = 121
- Planning Disk: Sun SPARCstorage Multipack Quantity: 2
 - 10.5 GB (5, 2.1 GB disks)
 - Hot swappable drives with multiple active spindles
 - Maximum of 12 drives for 25.2 GB
- Queuing Server:
 - Sun UltraServer
 - Clock: 167 MHz
 - CPUs: 4 of 12
 - Disk: 4.2 GB
 - RAM: 512 MB per CPU
 - Fast wide SCSI control: 2
- Queuing Disk: Sun SPARCstorage Multipack Quantity: 2
 - 10.5 GB (5, 2.1 GB disks)
 - Hot swappable drives with multiple active spindles
 - Maximum of 12 drives for 25.2 GB
- Production Monitor: NCD HMX Pro Quantity: 4
 - X Terminal, color
 - RAM: 16 MB
 - Size: 21"

3.4.2.7 Data Processing Subsystem

The Data Processing Subsystem (DPS) includes three hardware CIs:

- (1) *Science Processing (SPRHW)* -- The Science Processing HWCI (SPRHW) is the primary HWCI in the Data Processing Subsystem and contains processing resources (processors, memory, disk storage, and input/out subsystems) necessary to perform first-time processing, reprocessing, and Algorithm Integration & Test (AI&T). SPRHW also provides the hardware resources necessary to support management of the science processing, in the form of a Queuing Server and Production Planner Stations. However,

the queue management function in science processing is closely coupled to the ECS planning function, and the Queuing Server and Production Planner Stations are therefore discussed in the Planning Subsystem Design Specification (305-CD-026-002). This chapter describes only the hardware resources required to execute the science software.

- (2) *Quality Assessment and Monitoring (AQAHW)* -- The Algorithm Quality Assurance HWCI (AQAHW) provides hardware resources to support DAAC operations personnel in performing planned, routine, non-science QA of product data.
- (3) *Algorithm Integration and Test (AITHW)* -- The Algorithm Integration & Test (AI&T) HWCI (AITHW) provides hardware resources to support the integration and test of science software at the DAAC, and system level validation, integration, and test. It is important to note that this HWCI provides workstations and tools for software integration and test, but does not provide the compute environment or compute capacity required for science software test. This integration and test compute capacity is included in SPRHW.

These HWCI are specified below for GSFC. This specification does not include equipment to satisfy the DPS requirements for the Data Assimilation System (DAS) of the Data Assimilation Office (DAO). While DAS will be an integral part of ECS at the Goddard DAAC, its science processing requirements are somewhat different from those of the other instruments supported at the DAAC; DAS intends to use a distributed memory software and hardware architecture. The DAO processing resource will be dedicated to DAO processing, and the requirements for this processing power do not have the same procurement timelines, phasing, or allocation requirements as the other GSFC instruments. ECS has prepared a white paper that describes the deployment plan for the DAS environment (170-WP-005-001). The DAO hardware architecture is still under evaluation.

3.4.2.7.1 SPRHW

The design specification for SPRHW is derived by first reviewing the requirements for SPRHW in Section 3.4.2.7.1.1. The requirements are traced from the first principles established in the system requirements, through the detailed processing requirements levied by the instrument teams, and through the models built to simulate system behavior under the load defined by the instrument teams. The technology assessments performed to support SPRHW specification are reviewed in Section 3.4.2.7.1.2. The conclusions of the technology assessments are then applied to the derived, detailed system requirements to produce the SPRHW specification, provided in Section 3.4.2.7.1.3. Details of the specification are discussed in Section 3.4.2.7.1.4 to show how the system requirements are met.

3.4.2.7.1.1 SPRHW Requirements Analysis

The specification of SPRHW must satisfy sizing, expandability, reliability, maintainability, availability, compatibility, and interoperability requirements. These are discussed in the sections below.

SPRHW SIZING AND EXPANDABILITY. The SPRHW sizing and expandability requirements are derived from several sets of fundamental system requirements established in the ECS Functional and Performance Requirements Specification (F&PRS). Specific numerical detail is provided for these requirements by the processing plans provided by the instrument team inputs to the Ad Hoc Working Group on Production (AHWGP). Because the processing plans are quite complex, with many inter-relationships between products, significant analysis is required to understand the processing plans in order to reach a processing hardware design. This analysis has been performed using static and dynamic modeling. The modeling and design efforts have identified some design parameters that are not addressed by the AHWGP inputs; to address these design parameters, additional information concerning the processing plans has been solicited from the instrument teams in the form of a survey. The information derived at each of these steps has been used to determine the design parameters (requirements) for the science processing hardware. The following sections describe that process.

Functional And Performance Requirements. The F&PRS provides several requirements specifying the throughput and timeliness of ECS science processing.

Requirements EOSD1050, EOSD1060, EOSD1070, and EOSD1080 state the timeliness requirements for producing products for first-time processing of data. Level 1, Level 2, and Level 3 products must be produced within 24 hours of the time that ECS receives all of the data required to produce the product. Level 4 products must be produced within seven days of the time that ECS receives all of the data required to produce the product.

Requirement EOSD1040 states that ECS shall have the capacity to perform reprocessing at twice the rate of first-time processing. Requirement PGS-1300 states that ECS shall have the capacity to perform AI&T, production of prototype products, ad hoc processing for "dynamic browse" or new search and access techniques, and additional loads due to spacecraft overlap at one times the rate of first-time processing.

Requirement PGS-1301 states the requirement that vendor peak processing estimates for processors be derated by a factor of four for design purposes.

Requirement PGS-1270 states the requirement that the processing design and implementation have expandability by a factor of up to three without design change, and by a factor of up to ten without major design change.

The Technical Baseline And The AHWGP Inputs. Appendix C of the F&PRS is included by reference in the F&PRS requirements. This appendix contains estimates of the volume of processing to be performed by ECS. Because these estimates have changed over time and become significantly more detailed as instrument team plans have matured, the material in Appendix C of the F&PRS has been replaced by a group of separately maintained documents referred to as the ECS Technical Baseline.

The Technical Baseline contains material that defines the level of processing required by each instrument for first-time processing. Much of this data has been provided by the instrument teams in the form of inputs to the AHWGP. These inputs are updated approximately every six months; the most recent updates to the AHWGP inputs and the ECS Technical Baseline were baselined in February 1996. The February 1996 Technical Baseline is, therefore, the basis for the CDR design for SPRHW.

AHWGP processing requirements are expressed in terms of Product Generation Executives (PGEs), which are the smallest scheduled increment of science processing. For each PGE, the AHWGP input identifies the DAAC at which the PGE will be executed, and provides the execution frequency, the list of input data requirements, the list of output data products, and an estimate of the number of millions of floating point instructions (MFPI) performed with each execution. For each input file, an estimate is given of the fraction of the file that is read by the PGE. The baseline also identifies the calendar quarters during which the PGE will be executed; this allows the instrument teams to identify how they will phase in their processing after the instruments are put on orbit. As a result, the processing demand for each instrument may vary by calendar quarter.

The Technical Baseline documents program plans and directives that affect the SPRHW specification. The hours of operation for each DAAC are provided. For DAACs having less than 24 hour per day, seven day per week operations, the assumption is made that all processing must be accomplished during DAAC operating hours.

The Technical Baseline also documents the program directive to phase the implementation of processing capacity for each instrument relative to the date of the instrument's launch. The purpose of the phasing factors is to provide sufficient early processing resources to support AI&T, without purchasing the required full-up capacity before it is needed. With **X** defined as the processing resources required to do first-time processing for an instrument and **L** defined as the launch date of the instrument, the phasing factors are defined below:

- o **0.3X for $L-2 < t < L-1$.** Pre-launch AI&T requires 0.3X during the period from one to two years before launch.
- o **1.2X for $L-1 < t < L+1$.** Pre-launch AI&T and system I&T requires 1.2X during the year before launch. Standard instrument processing requirements (X) begin from the launch date and last for the remainder of the life of the instrument.
- o **2.2X for $L+1 < t < L+2$.** Post-launch AI&T, standard processing, and reprocessing of data require 2.2X starting at launch plus one year.
- o **4.2X for $t > L+2$.** Post-launch AI&T, standard processing, and reprocessing of data require 4.2X starting at launch plus two years.

The launch dates for instruments supported by ECS are documented in the Technical Baseline. Because the PGE processing plans are expressed on a quarterly basis, for the purpose of applying the phasing factors the launch dates are represented as the first day of the appropriate calendar quarter. The GSFC DAAC provides ECS processing resources for three instruments: LIS, MODIS, and COLOR. LIS will be carried by the TRMM platform and MODIS by the AM-1 platform; COLOR is scheduled on a flight of opportunity (FOO) in July of 1998. TRMM is scheduled for launch during the third quarter of (calendar year) 1997 (3Q97), and AM-1 is scheduled for launch during the third quarter of (calendar year) 1998 (3Q98). Therefore the phasing factors for SPRHW are assessed at 3Q96, 3Q97, 3Q98, 3Q99, and 3Q00. However, Release A of ECS has already put resources in place to cover the 3Q96 requirements; hence the focus of Release B design is for the latter four dates.

Specifications for SPRHW have been designed for each of the four Release B dates identified above. However, for the purpose of discussing the design, the 3Q99 period (also referred to within ECS as Epoch k) has been selected as a common frame of reference across the entire set of Release B Design Specifications. At Epoch k, TRMM instruments are supported at the 4.2X level of processing, and AM-1 instruments are supported at the 2.2X level of processing.

Static Modeling. The AHWGP inputs define over two hundred PGEs, executed at five DAACS for eleven instrument teams. The PGEs create over three hundred products, and require over one hundred ancillary and Level 0 input data sets. The interdependencies of the products and PGEs creates a network so complex that it cannot be understood by hand. The inputs are therefore analyzed using models, to reduce the volume of data into aggregates for each instrument and DAAC, and to understand the dynamic behavior of the interdependencies. The first step in this analysis is to use a static model to analyze the PGEs.

The static model represents each PGE within ECS on a single line of a spreadsheet. The columns of the spreadsheet contain the attributes for the PGE: its process identifier, instrument, processing site, frequency of initiation, number of executions per initiation, total input file size, total output file size, number of input files, number of output files, volume of files staged per execution, volume of files destaged per execution, and number of MFPI per execution.

Most PGEs are executed once per set of inputs, and those inputs are produced (by the instrument and/or by other PGEs) periodically. However, some PGEs may be executed several times on a given set of inputs that are periodically produced. An example is a MODIS PGE that produces products defined in terms of geographic tiles; for each set of input data covering the whole earth, the PGE may be executed 355 times, once for each MODIS tile. Hence it is necessary to track both the frequency of initiation of the PGE (the periodicity of the inputs) and the number of times the PGE is executed for a given number of inputs.

The numbers and sizes of input and output files are evaluated by summing the inputs provided for the PGE by the instrument team. The staging and destaging inputs to the static model are derived from the AHWGP inputs. Staging refers to the process of moving data to the science processing resources, and destaging refers to moving data away from the science processing resources. SPRHW does not archive data; it retains over extended periods (more than 24 hours) only a very limited amount of ancillary data identified by the instrument teams as permanent. All other data required by or produced by processing must be staged or destaged to the ECS archive. The staging and destaging entries in the static model spreadsheet are calculated by subtracting the permanent ancillary file volumes from the total file volumes.

These inputs are then used to calculate average resource usage levels for processing, network input/output (I/O), archive I/O, and disk I/O. The processing usage is calculated as the number of millions of floating point operations per second (MF) required to meet the PGE's requirements based on the PGE's frequency of initiation, number of executions per initiation, and MFPIs per execution. Similarly, the data flow to the archive (due to destaging), from the archive (due to staging), over the network (due to staging and destaging), and over the disk I/O channels (due to processing, staging, and destaging) are calculated. These values are then aggregated (summed) for each instrument and DAAC.

One method of static model analysis is to calculate the average level of resource usage required to support each PGE, instrument, and DAAC, assuming that the processing is spread smoothly over time. This analysis ignores the timeliness requirements for products, but it produces a baseline estimate of the minimum resources required to keep up with the processing requirements. It is this analysis that defines the value of X used for phasing requirements.

Another approach for analysis is to manipulate the formulas used to calculate the derived values (the usage rates) to take into account the timeliness requirements; further, it is assumed that in the worst case each periodic product is executed at least once on the worst case day. This produces a (generally) worst case estimate of the resources required to meet the timeliness requirements for first-time processing. (Because the static model cannot take into account the dynamic interactions between PGEs, it is theoretically possible that the worst case for the system is actually exceeded by this analysis, but the dynamic model has not shown this to be the case.)

Table 3.4.2.7.1-1, *GSFC Static Model Summary Results*, shows summary static model results for processing at GSFC for Epoch k (3Q99). The values in the table represent the static model estimates for resource requirements by instrument for *first-time processing only*. I/O estimates for COLOR are not available, but will be small relative to the MODIS requirements. These results represent a worst case analysis; that is, the results reflect the execution of all of the weekly, monthly, and yearly processes on the same day. The results also reflect completion of these processes within 24 hours, in accordance with the requirements. However, the static model analysis for the average case, in which all processing is spread smoothly over time, shows only modestly lower requirements: 4,819 MF processing, 97.2 MB/s processing I/O, 76.4 MB/s staging I/O, and 13.8 MB/s destaging I/O. Hence the periodic processes have almost no impact on processor load, and only moderate impact on the I/O requirements.

Table 3.4.2.7.1-1. GSFC Static Model Summary Results

Instrument	Processing (MF)	Processing I/O (MB/s)	Staging I/O (MB/s)	Destaging I/O (MB/s)
LIS	2	0.3	0.3	0.0
MODIS	4,838	120.8	97.4	16.3
COLOR	36	N/A	N/A	N/A
DAAC Total	4,876	121.1	97.7	16.3

Dynamic Modeling. The static model has significant limitations. It cannot take into account the dynamic interdependencies of the PGEs or the system's computing resources. As a result, it cannot accurately predict the end to end clock time required to produce a stream of products once Level 0 data has been received. It also provides extremely inflated estimates of the system staging requirements, because it must assume that every (non-permanent) file required for each execution of each PGE is staged for that execution; it cannot take into account that the file may

already have been staged to SPRHW in support of some other PGE. On the other hand, it assumes uniform utilization of the SPRHW resources, not accounting for the fact that one type of resource (such as a processor) may be idle while the system uses another type of resource (a network).

To overcome these limitations, a dynamic model of a subset of ECS has been implemented. The dynamic model is implemented using the Block Oriented Network Simulation (BONeS) tool. BONeS is a discrete-event simulation tool for analysis and design of communication networks and distributed processing systems. Components of a distributed processing system are represented by nodes, which have resources associated with them that get allocated as events request them. First-time production of products by DPS is simulated by the Processing module, in conjunction with the Data Handler module, the Event Driven Scheduler, and the Ingest module. The Data Handler simulates the behavior of the ECS Data Server. The Data Handler stores and retrieves data from the permanent archive, routes data to the requesting subsystems, and manages tiered storage and staging resources. The Event Driven Scheduler monitors the availability of data, requests data to be staged from the Data Handler to Processing, routes newly created data to the appropriate Data Handler or Processing resource, and initiates execution of a process when all required inputs are present. It should be noted that the Event Driven Scheduler is not intended to serve as a simulator of the Planning Subsystem; rather, it simulates the effects of data arrival at ECS and the actions of the Planning Subsystem. The Ingest module simulates the Ingest subsystem, which accepts data from external systems and users and contains rolling storage of Level 0 instrument data.

The dynamic model is run by creating an input event stream representing the arrival times of Level 0 and ancillary data files. The PGEs are defined to the model in terms of their resource requirements and file dependencies. The system design is also represented as an input to the model, in terms of the number and capacity of system resources -- for example, the number of SPRHW processors and their throughput in MF. The Event Driven Scheduler then executes the event stream, which is driven by the external events (receipt of Level 0 and ancillary files) and by derived events (the availability of the data and computing resources required to execute a PGE). The simulation is run over an extended simulation period -- typically 21 days -- to examine the effects of products produced periodically. The simulation is started with an empty ECS system; there is therefore a starting transient in the simulation's output results. Because of this, the results of the first five days of the simulation are generally discarded.

The results of the dynamic model are analyzed to identify peak and average resource requirements, and average and maximum system response characteristics. Graphs showing resource utilization over time during the period of the simulation can also be generated.

The model explicitly represents each processor within SPRHW, and constrains the simultaneous execution of PGEs to the number of processors in the design, operating at the vendor's peak MF rating derated by a factor of four per the requirements. Average and peak processor utilization and average and peak processor queue lengths are reported, as well as processor utilization curves over the period of the simulation. These results identify whether processing backlogs are occurring, and whether there is excess processing capacity in the design. In this respect the

dynamic modeling is an iterative process; an initial, generous configuration of resources is fed to the model, and the results are examined. Design choices are made, and a refined configuration of resources is fed to the model. This process continues until all of the design decisions have been made, and the model indicates that the configuration will satisfy the system's requirements.

For disk storage at the science processors, the model does not constrain the available space; instead, it implements an algorithm to depict the purging of unneeded data from the SPRHW storage array, and monitors the demand for storage. By plotting the storage demand over the period of the simulation, it can be determined when the peak occurs, and for what duration. This information is used to size the SPRHW storage arrays.

Dynamic model results for first-time processing from the August 1995 Technical Baseline indicate that LIS will require 21 GB of disk space and MODIS will require 483 GB. These results also show that LIS disk space requirements do not grow significantly when reprocessing load is added, but that MODIS disk space requirements double when a one X reprocessing load is added. COLOR is not represented in the dynamic model because detailed information about its PGEs is not available.

Network and disk I/O are averaged over time by the dynamic model to provide average utilization rates. The model represents these as constrained resources; hence, the model results provide assurance that a given level of network and disk I/O throughput will satisfy the system's requirements, and provide an approximation of the fraction of these resources that is used on average. Based upon dynamic model results from the August 1995 Technical Baseline, the average staging requirement is 0.1 MB per second for LIS and 6.1 MB per second for MODIS. These staging estimates are much less than those predicted by the dynamic model.

Memory Usage And I/O Survey. The inputs to the AHWGP do not provide any information as to how the PGEs will use processing memory, or about the detailed characteristics of their disk I/O.

Although all of the potential platforms considered for SPRHW make use of virtual memory, the ratio of physical memory (RAM) to virtual memory required during the operation of a system is important for performance reasons. The rate at which memory is accessed by a PGE and the pattern of these accesses can also have important performance impacts. For the SPRHW specification, these requirements affect the amount of RAM, number of memory channels, amount of disk (for virtual memory), and number of disk channels (for virtual memory access). Because the SPRHW processing requirements are very demanding and generally require sophisticated and expensive processors, it is desirable to get as much throughput as possible from the processors by keeping as much of a PGE in RAM as practical.

Similarly, the size and pattern of disk I/O requests strongly affects the performance that can be expected from a disk I/O subsystem. Because of various overheads associated with each disk access, throughput from a disk subsystem increases strongly with the average size of the disk request, up to maximum limits imposed by the disk, interface, and controller technologies. While vendors usually quote I/O subsystem performance for large request sizes (hence peak performance values), it is important to know whether these request sizes are typical for ECS applications; and if not, to derate the vendor performance specifications accordingly.

The type of memory usage and file I/O characteristics that are needed to determine the system requirements for these resources can generally be determined either by code inspection or by measurement during execution. ECS solicited this information from the Release B instrument teams in January of 1996. Unfortunately, it is still too early for most teams to provide accurate assessments of their requirements, and little data has been received to date in response to the survey. The integration and test of beta science software for Ir1 during the first half of 1996 may provide requirements in time to support the procurement of the first increment of Release B hardware.

For purposes of the generation of the Release B CDR design, assumptions have been made in the absence of requirements data in these areas. It is assumed that each PGE will require approximately 128 megabytes (MB) of RAM. This figure is viewed as a reasonable upper limit for most algorithms, especially if the algorithms are designed with consideration of their memory requirements. It is recognized that some algorithms may require large data structures to iteratively refine their data using models, and will have correspondingly larger memory requirements. These uneven needs between algorithms are mitigated in part by the Symmetric Multiprocessor (SMP) architecture that has been selected for SPRHW.

It is also assumed that ECS SPRHW file I/O will be buffered by the platform operating system, with the average request size characterized by the operating system I/O buffer size. This assumption is made because it is believed that this correctly characterizes the use of HDF, which is expected to be the dominant form of disk access on the science processors. The buffer size varies by operating system, but is typically small (64 kilobytes or less) relative to request sizes used for direct file I/O (256 KB to 4 MB). This means that the vendor disk I/O performance ratings will be derated to reflect small request sizes.

SPRHW RELIABILITY, MAINTAINABILITY, AND AVAILABILITY REQUIREMENTS. The product generation function of DPS has an availability requirement of 0.96 (96%) and a mean down time requirement of not to exceed four hours. Implicit in these requirements is also the requirement that SPRHW be able to achieve its timeliness and throughput requirements while allowing for 4% non-availability and component outages of up to four hours.

SPRHW COMPATIBILITY AND INTEROPERABILITY REQUIREMENTS. The Release B ECS software must satisfy ECS requirements for the use of open systems technology, and must be compatible with other Release B subsystems and with hardware provided to the DAACs for Release A. Specifically, the SPRHW hardware suite must be supported by OODCE. SPRHW must also support the network interfaces to be implemented by ECS. For GSFC, these interfaces must support data exchange between DPS and DSS at aggregate rates (in Epoch 0) of on the order of 50 MB per second.

3.4.2.7.1.2 Technology Assessments

The requirements identified in Section 3.4.2.7.1.1 can be compared to the capabilities of current and future technologies to identify candidate technical solutions. This type of evaluation can be done on paper, in the form of trade studies or technology evaluations, or it can be done in the

laboratory, through prototyping and benchmarking. The paragraphs below discuss the technology assessment activities that have been performed to support the Release B design process, with emphasis on the activities performed since IDR-B.

TECHNOLOGY EVALUATIONS. A number of technology evaluations have been performed under ECS to identify the best technical approaches and products for the SPRHW suite. Three of these evaluations are discussed below.

Production Platform Families. The requirements for first-time processing for instruments at the DAACs range from a few MF for the less compute-intensive instruments, up to 4.8 GF for MODIS at GSFC. Computing requirements in the GF range have historically been the turf for vector supercomputers, but advances in massively parallel processors, symmetric multiprocessors, and workstation farms have pushed these technologies into this range of performance. An evaluation was performed for Release A and Release B of the following candidate platform family architectures for SPRHW:

- A. Single processor workstations;
- B. Farms of workstations connected via local area networks;
- C. Symmetric Multiprocessors (SMPs);
- D. Massively Parallel Processors (MPPS); and
- E. Vector Supercomputers.

The key criteria used to evaluate the platform families were lifecycle cost per unit of processing power, expandability, ease of science software development, and flexibility. The evaluation found that the requirements for the less compute-intensive instruments could be met by single processor workstations, which are the least expensive and most flexible resources. For the compute-intensive instruments, however, only the SMPs, MPPs, and vector supercomputers offered the expandability required. Among these three, the vector supercomputers were found to have significant cost disadvantages. The SMPs were found to have significant advantages over the MPPs in the ease with which applications can be developed, and the evaluation concluded that the SMP platform class was best suited for ECS's high-end SPRHW requirements.

This evaluation is documented in *Platform Families for the ECS Project*, 440-TP-007-001.

Distributed And Parallel Processing. Parallel processing in both shared memory and distributed memory architectures is viewed as an emerging technique for solving large processing problems. An analysis has been performed to examine the benefits of distributed and parallel computing for ECS. The analysis studies various processing alternatives for ECS science algorithms and provides information on processing technologies. The analysis examines the applicability of using the OSF/Distributed Computing Environment (DCE), symmetric multiprocessing with shared memory, distributed multiprocessing (including workstation clusters), and massively parallel processing for ECS science software.

The SMP architecture is found to provide significant flexibility: it fully accommodates non-parallelized code, supports easy migration of non-parallel code to shared memory parallel processing, and can also support distributed memory parallel processing. This flexibility is viewed as giving SMP systems a significant advantage for ECS processing. Although a need to

parallelize particular PGEs has not been identified to date, the ability in the future to parallelize ECS science software is an important consideration, especially with the large processing requirements of Release B and beyond.

This evaluation is documented in *Distributed and Parallel Processing For ECS Science Algorithms: A Trade Analysis*, 440-TP-008-001. ECS Science and Technology Lab prototyping efforts that have demonstrated parallel program development, based primarily on parallelization tools, are documented in 194-00569TPW and 194-430-TPW-001.

Production Topologies. A technical analysis has been performed examining the advantages and disadvantages of distributing processing tasks from one or more instruments across one or more processing clusters. The purpose of the study was to examine how the interactions between the processing requirements of different instruments might increase or decrease the total hardware resources required to perform the aggregate load. The results of this analysis are reflected in the SPRHW design.

LIS is a TRMM instrument; therefore it requires that a production platform be in place at GSFC beginning with ECS Release A. To avoid disruption of LIS processing at Release B, it is recommended that this platform remain in place, and continue to support LIS throughout its production life. Because LIS processing requirements are so small, its platform will be able to support the COLOR processing requirements as well.

MODIS processing at GSFC will require the resources of more than one large platform. Because MODIS requirements at GSFC dwarf LIS and COLOR requirements, mixing the processing from these instruments on one platform is not recommended; modest variations in MODIS requirements might severely impact processing for the less compute-intensive instruments.

This analysis is documented in *Production Topologies: A Trade-Off Study Analysis for the ECS Project*, 440-TP-006-001.

PROTOTYPES AND BENCHMARKS. Prototypes and benchmarks are used to assess the applicability of a technology or product to a set of requirements. Two sets of recently acquired benchmarks have been used in refining the ECS SPRHW design for Release B. The first set of benchmarks, measuring the performance of IP over HiPPI, was performed in the ECS development facility. The second set of benchmarks, measuring I/O subsystem performance on the SGI platforms, was performed at the University of Minnesota and by SGI.

Network Benchmarking. Benchmarking tests of the performance of IP over HiPPI have been performed in the ECS development facility. The purpose of these tests was to measure the throughput that can be achieved using IP over HiPPI, to determine whether IP introduces a significant overhead in HiPPI communications.

The tests used two SGI Challenge XL systems, equipped SGI HIO HiPPI adapters. These systems were running version 5.3 of the Irix operating system. Two different benchmarking tools were used to send TCP streams from the memory of one system to the memory of the second system. Throughput rates were measured as the configuration parameters for TCP/IP were varied.

The maximum transfer rates observed in the test were approximately 55 MB per second, just over half of the 100 MB per second theoretically possible with HiPPI. SGI has stated that using version 6.2 of Irix, rates of up to 90 MB per second have been observed in their laboratories. This indicates that IP does not burden HiPPI with a substantial overhead, and that IP-based protocols can be used over HiPPI in ECS. This eliminates the need to develop custom protocols for ECS using raw HiPPI interfaces.

I/O Subsystem Benchmarking. The Laboratory for Computer Science and Engineering (LCSE) at the University of Minnesota, in affiliation with the Army High Performance Computing Research Center (AHPCRC), has performed a number of benchmarking tests on I/O subsystem performance on SGI platforms. These tests have included performing high speed data transfers over HiPPI using a lightweight TCP/IP protocol, and performing high data rate disk transfers. These tests have demonstrated that the SGI architecture can sustain I/O rates in excess of 150 MB per second between machines, and over 500 MB per second to locally attached file systems.

The HiPPI experiments performed at LCSE demonstrated the ability of the SGI platforms to support data visualization. In these experiments, a single Challenge server, configured with four HiPPI connections, was used to transfer data from its file system to the graphics processors of two SGI Onyx systems. A lightweight TCP/IP protocol referred to as "NFS-Bypass" was used to perform the transfer. This protocol, which provides an NFS-like interface, transfers data between systems using a socket to socket connection, bypassing most of the overhead normally imposed by NFS. Transfer rates for a single HiPPI connection, from disk to graphics processor, of 65 MB per second have been observed. Using multiple HiPPI connections transfer rates from a single machine of over 150 MB per second have been observed.

Research at LCSE has also been performed to evaluate the maximum transfer rates available from various storage configurations on SGI platforms. These experiments have focused on the use of striped, SCSI-2 based RAID file systems using SGI's XFS file system and direct I/O. The LCSE researchers have built file systems capable of sustaining over 500 MB per second transfer rates on a Challenge system. A key aspect of their findings in the building of fast file systems is the need to use large block request sizes. They found that block request sizes associated with buffered I/O -- I/O buffered by the operating system's I/O cache -- are 64 KB or less. Disk subsystem performance at these request sizes is substantially below peak levels; at this request size, throughput of less than 10 MB per second is observed on SCSI-2 channels having a theoretical bandwidth of about 20 MB per second. Using direct I/O, which is not buffered by the operating system's cache, and with request sizes of one MB or more, throughput rates rapidly approach the per channel limitation imposed by SCSI-2.

SGI has further tuned the NFS Bypass software to make efficient use of the Challenge memory subsystem, and is productizing the software. The product, Bulk Data Service (BDS), will be offered as an extension to Irix 6.2 in the second quarter of 1996. Significantly, since BDS is built on top of TCP/IP, it is not specific to HiPPI; BDS may be used to implement highspeed file transfers over any TCP/IP connection.

SGI has also recently performed benchmarking tests on their newest generation of RAID controllers. These tests indicated sustained transfer rates of eight MB per second per controller are achieved when small request sizes (64 KB) are used. This benchmark is used as a basis for sizing the disk I/O subsystems for SPRHW in Release B.

3.4.2.7.1.3 SPRHW Specification

SPRHW has three top-level component types: Science Processors, Queuing Servers, and Production Planner Stations. As discussed earlier, the Queuing Server and Production Planner Station components are specified in the Planning Subsystem Design Specification.

Appendix A provides detailed configuration information for the GSFC SPRHW suite at Release A and Epochs c, g, k, and o. The intended function and the processing, memory, I/O subsystem, disk channel, storage device, network, enclosure, and console configuration are specified for each system in the SPRHW suite. Upgrades to systems can be identified by comparing the configurations for a system at different Epochs.

The Science Processors are characterized by their vendor, enclosure, processors, memory, I/O subsystems, internal disk drives, external disk systems, backup and update devices, network interfaces, and monitors. The sections below provide further explanation for the specification provided in Appendix A. The system designators defined in Appendix A (e.g., SPRHW-GSFC-1, or just -1) are used in the text below.

SPRHW VENDOR. The Release B SPRHW components, with the exception of tape libraries provided by Exabyte, will be provided by Silicon Graphics (SGI). The architecture trade-offs show that SMP systems provide clear advantages for ECS science processing, and SGI is a leading vendor of SMP systems. When considerations are taken into account for cost, availability of required ECS software capabilities (such as OODCE), availability of highspeed network components (particularly HiPPI), and ease of re-use of Release A hardware, the logical choice for Release B hardware is SGI.

Final vendor selection has not been made for the external disk systems required by SPRHW. A baseline design is presented here which features RAID-5 devices using SCSI-2 interfaces. This technology is mature and its performance is fairly well understood. However new technologies, particularly fibre channel based RAID arrays, are making their way into the market. If these products are found to be stable and cost effective before the first Release B procurement, they may be selected as the Release B baseline.

PROCESSORS. For new equipment purchased for Release B, SPRHW will use the MIPS R10000 processor. This chip has a superscalar 64 bit architecture, capable of performing two floating point operations per clock cycle. SGI has announced that the chip will be offered in a 275 MHz implementation in the second half of 1996. Using the derating factor defined in the F&PRS, this processor is estimated to provide 137.5 MF.

GSFC will retain one system, purchased for Release A and designated SPRHW-GSFC-4, that has a 150 MHz R4400 processor. This processor is estimated to provide 18.75 derated MF. Additional processors will be purchased for this system, which will be used to support LIS (beginning in Release A) and COLOR.

A second system purchased for Ir1, SPRHW-GSFC-1, will be upgraded from 75 MHz R8000 processors to 275 MHz R10000 processors.

Decisions about how many processors to configure in a single system were driven by the processing requirements for the instruments and for AI&T.

The decision to retain a system with R4400 processors is considered a cost-effective means of providing for the requirements of LIS and COLOR. This machine (-4) is already in place for Release A, and LIS will begin production operations on it. It can be upgraded through Epoch o to provide for 3X processing for both LIS and COLOR.

First-time processing requirements for MODIS require approximately 48 R10000 processors. These could be housed in two enclosures, with 24 processors each; however, the I/O requirements for these two systems would be high. Since additional systems are required at Epoch k to handle the initial reprocessing load, splitting the load at Epochs c and g across three systems (vice two) does not increase the ultimate platform count. The specification for Epochs c and g therefore uses three systems, each with 16 processors. This also reduces the number of processors required for the AI&T system to act as a failover replacement for one of the MODIS production systems, and leaves more room for growth (via upgrades in place) if it is needed before Epoch k.

The requirement for AI&T processing power for TRMM instruments is 0.2X at Epochs c and g, and 1.2X at Epochs k and o. The requirement for AI&T processing power for AM-1 instruments is 0.2X at Epochs c, g, and k, and 1.2X at Epoch o. The value of X for each instrument is defined above in Table 3.4.2.7.1-1, as determined by the static model. System SPRHW-GSFC-1 is configured to meet the AI&T requirements at Epochs c and g. At Epoch k, this system is switched to MODIS production, and a new system (SPRHW-GSFC-9) provides GSFC AI&T support. Note that for Epochs c and g, SPRHW-GSFC-1 is configured with significantly more processing power than strictly required for AI&T; this is so that this system can be used as a failover system, should any of the production machines go out of service. At Epoch k, the SPRHW-GSFC-9 system is configured only to provide AI&T support; MODIS failover at this point can be handled within the MODIS suite. At Epoch o, an additional system is added (SPRHW-GSFC-10) to satisfy the AI&T requirements.

MEMORY. It was noted above that firm data for memory requirements are not available yet, and thus it would be assumed that each processor requires approximately 128 MB of RAM.

For SMP configurations, the SGI memory architecture features leaf controllers, which process requests to a subset of the system memory. Each memory board can have one or two leaf controllers. A system can have up to eight memory boards (subject to backplane limitations), but only up to eight memory controllers. The I/O benchmark results by the group at the University of Minnesota have shown that memory interleaving is important for I/O performance. It is also expected to be important for science algorithm performance on systems with many processors. Therefore, machines with more than one GB of RAM and machines performing high rates of I/O will be configured with eight memory controllers; lesser machines will generally be configured

with four memory controllers. The system used by LIS and COLOR (SPRHW-GSFC-4) has significantly less demand placed on it, and is configured with two-way memory interleaving.

SGI offers RAM at two chip densities. With options for up to eight memory boards, there are a large number of memory configurations offered, ranging in total size up to 16 GB.

I/O SUBSYSTEMS. For SMP systems, the SGI architecture provides configurable I/O subsystems which attach to the backplane; each subsystem occupies one backplane slot and provides up to 320 MB per second of bandwidth. The subsystem card is referred to as a PowerChannel 2 or IO4 card. A system may have up to six IO4 cards, subject to backplane limitations.

The first IO4 card in an SMP system provides console and ethernet connections. Each IO4 provides serial and parallel connections, two fast wide differential (FWD) SCSI-2 channels, and space for two HIO controller cards. HIO controller card offerings include a HiPPI card, a FDDI card, and a card supporting three SCSI-2 channels.

The number of IO4 cards specified for each SMP is determined by allocating HIO slots to the FDDI and HiPPI interfaces, and counting the number of SCSI-2 interfaces required. The number of SCSI-2 interfaces required is determined by the number of internal and external SCSI-2 devices supported by the system. In general, it is assumed that the internal slow SCSI-2 devices (CD-ROMs, floppy disk drives, and tape drives) will be aggregated on the first SCSI-2 channel. Internal disk drives will be allocated to the second SCSI-2 channel. External disk arrays will be allocated to subsequent SCSI-2 channels; the number of channels required is based on the desired throughput of the external file systems.

The Release A GSFC systems (-1 and -4) have VME FDDI cards which attach directly to the backplane via special VME slots. These controllers will be retained in Release B, and reduce the number of HIO slots needed on these machines.

The table in Appendix A shows how each IO4 HIO slot and each SCSI-2 channel is allocated on the GSFC systems.

INTERNAL DISK DRIVES. The bulk of storage for the science processors will be provided by external storage arrays. Internal disk drives will only be used to provide swap space for the operating system, and to provide file system space for the operating system and applications. Applications in this context does not include the PGE executables or any temporary file space required by the PGEs; rather it refers to file system space requirements for the Autosys client software and ECS custom code. The file system requirement for the operating system and applications is not expected to exceed two GB.

The allocation for swap space is estimated at four times the size of the physical memory (RAM). This exceeds what is generally configured for servers, and probably represents an upper bound; if it is found that systems require virtual memory larger than four times their physical memory, it is likely that physical memory will also have to be upgraded to reduce paging.

SGI currently offers 2 GB and 4.3 GB internal disk drives and will soon offer 9 GB internal disk drives. A combination of these drives has been selected for each system to satisfy the estimated space requirement. These drives will be allocated to a single SCSI-2 channel.

EXTERNAL DISK ARRAYS. The disk size and throughput requirements for SPRHW are determined on a DAAC by DAAC and host by host basis, using the static and dynamic modeling results. These requirements are translated into numbers of drives and controllers based upon the configurations available from the vendor. The current product of choice for external disk arrays is the RAID-5 product from SGI.

The SGI RAID-5 arrays use one redundancy disk for each four data disks. Arrays are built in groups of five disks, with up to four groups (20 disks) in an enclosure. Up to four enclosures can be put in a rack.

An enclosure will support one or two controllers. Each controller can access one or more groups of disks in the enclosure. A group of disks within an enclosure can be accessed by both controllers; however, only one controller may access the group at a time. This form of dual attachment is useful only for implementing failover of the controllers (or hosts) without moving cables. Each controller within an enclosure can also be tied to more than one host; however, only one host may mount the associated file system at a time, and therefore this form of dual attachment is also only useful to implement failover. The controllers in an enclosure can be connected to different hosts, and in fact this configuration is specified for some DAACs.

The throughput to the SGI arrays is limited by either the interface mechanism (at large request sizes) or by overheads in the host and the controller (at small request sizes). The SCSI-2 interface used in the SGI arrays allows a maximum bandwidth of 20 MB per second to a channel. At small request sizes, however, the overheads are expected to limit performance to not more than eight MB per second per controller. In order to avoid contention on the SCSI-2 channels that would significantly reduce throughput for large block requests, and would also modestly impact throughput for small block requests, each array controller is configured on its own SCSI-2 channel.

The number of SCSI-2 channels and array controllers is determined by dividing the throughput rate required for a system (determined from the static and dynamic modeling results) by the eight MB per second rate per channel assumed for small request sizes. This number is then rounded upward. (Although the disk I/O performed by a system will actually be a mixture of large and small block requests, no effort is made to determine an average throughput rate; rather, the slower rate associated with small block requests is used to represent the entire load. This increases the design margin for these components.) The number of enclosures and racks is then matched to this number of enclosures.

SGI offers 2 GB and 4.3 GB disks in its arrays and soon will also offer a 9 GB disk. The 2 GB disks are not specified for use in SPRHW because the SPRHW disk requirements are too large to make use of these drives practical.

Although the number of groups and size of disks within groups can be varied from one controller to another, the number of such variations assigned to a single host has been deliberately limited. This is because it is expected that file systems will be defined so that they are striped across multiple controllers. Such striped file systems are necessary to support very high speed devices such as the HiPPI network. Striped file systems can only be built using partitions of uniform size, which are most easily built when each controller within the stripe is controlling the same configuration of disks.

The disk enclosures will be configured with three power supplies (one redundant supply) and a minimum amount of write cache. (Vendor benchmarks suggest that write cache will not significantly improve performance for the request mix that SPRHW is expected to produce.)

BACKUP AND UPDATE DEVICES. Each science processor will be equipped with an internal CD-ROM, as this is the delivery mechanism employed by SGI for updates to operating system software.

The SPRHW suite will include two tape libraries to provide backup capability for SPRHW and for other subsystems, using the FDDI backbone network. For Ir1 and Release A, two Exabyte EXB-210 and tape libraries were purchased and put in place at GSFC. The EXB-210 is an 8 mm Digital Audio Tape (DAT) library with eleven cartridge slots and two drives. These components will continue to be used for Release B.

MONITORS. In general, the science processors are to function as compute servers, and are not configured with sophisticated graphics processing capabilities. It is expected that the machines will be housed in raised floor environments where space is at a premium. The science processing hosts require a console in order to monitor the system at startup, and to perform some system administration tasks; however, this console need not be a bulky graphics monitor. ECS Operations and Maintenance are investigating ways to rack the system consoles, and to allow systems to share the same console.

NETWORK INTERFACES. An FDDI subnetwork will be implemented at GSFC to support the Planning and Data Processing Segment (PDPS). Each processing component of SPRHW (including the Queuing Server and the Production Planner Stations) will be interfaced to the PDPS FDDI subnetwork; the design of this subnetwork is specified in the Design Specification Overview, CD-305-020-002, and in Section 3.4.1 of this document.

At GSFC the data transfer requirements between DPS and DSS necessitate the implementation of a switched HiPPI network. The HiPPI network will be implemented via a central HiPPI switch with switched 800 Mbps interface ports connected directly to the SPRHW and DSS hosts; this design is also documented in the Design Specification Overview and in Section 3.4.1 of this document. Note that SPRHW-GSFC-4, the Challenge DM used by LIS and COLOR, will not be upgraded with a HiPPI interface; the network I/O rate for this system is modest and can be served by the FDDI network.

ENCLOSURES. The actual enclosures for the science processors and their disk arrays are generally of little interest, except that they impose constraints on expansion capabilities, and they occupy floor space at the DAACs.

The Challenge DM uses a floor-standing half-height cabinet. The Power Challenge XL uses a full height cabinet.

Rack-mounted enclosures are specified for the external disk storage units, in order to save floor space. Units that were cabinet mounted in Ir1 and Release A will be rack-mounted for Release B.

3.4.2.7.1.4 SPRHW Design Discussion

This section provides a discussion of how the specification provided in Section 3.4.2.7.1.3 meets the requirements identified in Section 3.4.2.7.1.1. The sections below provide a general discussion of how the specifications meet the sizing and expandability requirements, and what the general SPRHW failure recovery strategy will be.

SIZING AND EXPANDABILITY. It is important to note that although the science processors are divided into clusters of resources configured to deal with specific processing requirements, and are expected to be tuned and allocated for these purposes, this does not imply that the use of a processing cluster, or a processor within that cluster, is limited to one operating mode or instrument. The DAAC operations staff, using the planning tools provided by ECS, will determine how the DAAC work load is distributed across the pool of SPRHW resources at the DAAC.

Table 3.4.2.7.1-2, *Comparison Of GSFC SPRHW Requirements And Specification*, compares the resources required at GSFC to those provided by the specification for each Epoch of interest. The required values represent 1.2X, 2.2X, and 4.2X levels of processing (at the appropriate Epoch for each instrument), calculated from the static model. The row titled "Specified" shows what the planned configuration provides. Each SCSI-2 disk channel is estimated to provide eight MB per second throughput, and each HiPPI interface is assumed to provide 70 MB per second throughput. Note that question marks are shown for the memory requirements, as these requirements are not accurately known at this time. Requirements for disk space, disk bandwidth, and network bandwidth for AI&T are set to zero, as AI&T requirements for resources other than processing are not easily predictable; however, the specification provides ample resources in these categories. The table shows that the specification meets or exceeds the peak and average requirements for each resource at each Epoch.

Table 3.4.2.7.1-2. Comparison Of GSFC SPRHW Requirements And Specification

			c	g	k	o
			3Q97	3Q98	3Q99	3Q00
CPU	Throughput (MF)	Required	5,853	5,853	10,731	20,479
		Specified	8,875	8,875	11,075	20,737
RAM	Size (MB)	Required	?	?	?	?
		Specified	8,704	8,704	9,728	16,896
	Interleaving (N)	Required	?	?	?	?
		Specified	34	34	42	66
Disk	Space (GB)	Required	21	504	987	1,470
		Specified	916	916	1,340	2,046
	Bandwidth (MB/s)	Required	0.3	121.2	241.9	362.7
		Specified	272.0	272.0	400.0	656.0
Network	Bandwidth (MB/s)	Required	0.3	16.5	32.3	48.3
		Specified	287.0	287.0	357.0	567.0

The ability to expand the SPRHW configurations merits a separate discussion. A DAAC's science processing capabilities can be expanded in two ways: by adding resources to existing hosts and disk arrays, and by adding new hosts and disk arrays. The former approach is referred to as upgrading in place.

For the SMP hosts, the ability to upgrade in place is limited by the number of backplane slots.

The Challenge DM (SPRHW-GSFC-4) could be upgraded in place to a Challenge L; both have five backplane slots, but the L configuration allows them to be used more flexibly. Every system must have at least one processor board, one memory board, and one IO4 board. An L system may therefore have up to twelve processors, configured as three boards of four processors each. The maximum memory configuration would feature three boards, or six GB of memory with six-way interleaving. If only one memory board and one processor board were used, three IO4 boards could be used. This would support six HIO cards, plus six SCSI-2 channels. This combination could be used to support a FDDI interface, a HiPPI interface, and enough SCSI-2 interfaces (a total of 18) to support a 144 MB per second file system. For a system with balanced processing, memory, and I/O, a practical maximum configuration would have two processor boards (eight processors), two memory boards (each with up to two GB of RAM and two leaf controllers), and one IO4 board (with FDDI and SCSI-2 HIO cards). If the processors were upgraded from 150 MHz to 250 MHz, this would provide 250 MF, or 219 percent of the peak requirement at any Epoch, and would provide 512 MB per processor, or 400 percent of the current estimated memory requirement.

The Power Challenge XL (used for the -1, -5, -6, -8, -9, -10, -11, and -12 systems) has 15 backplane slots, although this limit can be expanded to 25 by using a backplane extender. The maximum number of processors supported by Irix 6.2 on the Power Challenge is 36; this would require nine processor boards. Up to eight memory boards can be used, although only eight leaf controllers can be configured on those boards. Up to six IO4 boards can be put in the system. Thus a maximally configured Power Challenge XL would require 23 of the 25 available slots. Such a system would provide 4.95 derated GF, with 16 GB of RAM. The I/O subsystem would support multiple HiPPI connections and SCSI-2 controllers. (In benchmarking experiments at the University of Minnesota, Challenge systems have been configured with as many as four HiPPI interfaces and as many as 40 SCSI-2 controllers -- but not simultaneously.)

At Epoch k, the specification calls for the use of five Power Challenge XLs at GSFC (four for MODIS production and one for AI&T). If each of these systems was upgraded to the maximum configuration described above, these systems would provide an aggregate processing capability of 24.75 GF (derated). This would provide 125 percent growth over the current configuration. The aggregate memory would grow by 789 percent, and the number of disk channels would increase 400 percent.

If growth beyond these scales is required, it will be necessary to add more hosts to the SPRHW configuration. The HiPPI switch supports 16 ports, but if necessary the switch can be cascaded; adding a second switch would thus provide 30 ports for DPS and DSS hosts. If the ratio of DPS

hosts to DSS hosts is assumed to be two to one, this would allow connections for 20 DPS hosts, or a maximum SPRHW compute power of 99 GF (derated). This is roughly 477% of the 20.7 GF specified for GSFC at Epoch o.

If processing power greater than 99 GF (derated) is required, it would be necessary to modify the design slightly by adding a second HiPPI network, and partitioning the HiPPI network traffic. (Cascading the HiPPI network to three switches is not recommended by the vendor, although it may be feasible.)

It should be noted that adding SPRHW hosts potentially has the effect of increasing network traffic, because the SPRHW storage subsystems are locally (host) attached, and can not be shared between hosts. As more hosts are added to process the same load, the potential that data will have to be replicated to multiple host file systems increases, and the aggregate file system size, disk I/O, and network I/O also increase.

FAILURE RECOVERY STRATEGY. The general strategy in the event of the failure of an SPRHW component is to re-distribute the first-time processing load to SPRHW components having sufficient resources to handle the load. Reprocessing and AI&T, although important to ECS over the long haul, can be delayed or reduced until the failed components are replaced. Because SPRHW has redundant capacity to support AI&T and reprocessing, and because scheduling of SPRHW resources is performed dynamically by the queuing and planning resources, most failures will have little impact on the timeliness of production processing.

SPRHW Failure Recovery. The DPS reliability, maintainability, and availability (RMA) requirements of 96% availability and mean down time of less than four hours are met by the SPRHW design. An analysis demonstrating this is provided in *Availability Model/Predictions for the ECS Project*, 515-CD-001-003.

In the event of a failure in the SPRHW suite, the DAAC operators will isolate the fault and determine its severity. Vendor maintenance will be called and repair and/or replacement of the affected part will be initiated.

Depending on the severity and location of the failure, other management actions and decisions may be required. If the failed component only affects resources that were allocated to AI&T, there is no impact to DAAC production operations, and no further action to recover production activity is required.

If the failed component affects resources that were allocated to production, an assessment will be made whether any PGEs need to be re-started or re-run because of possible corruption. If this is the case, ECS production planning capabilities will be used to reschedule the PGEs. If the failed component has affected a limited subset of the production resources, the MSS subsystem will notify the Queuing Server that these resources are no longer available, and the Queuing server will appropriately stop scheduling these resources. If the remaining resources available for production processing are sufficient to keep up with the first-time processing plan, an immediate re-plan may not be necessary; if the first-time processing PGEs are prioritized above reprocessing PGEs (generally the case), the first-time processing will remain on schedule. If it is necessary to shut down SPRHW resources in order to facilitate a repair, a re-plan can be generated, allocating time in the schedule for the repair.

If a failure reduces the working SPRHW resources allocated to production below the level required to support first-time processing, the DAAC operators may choose to re-allocate AI&T resources to production. This may be done in one of two ways: by using the production and queuing software to logically re-allocate and re-plan the use of the system's resources, or by using the physical hardware in the AI&T hosts to replace failed hardware in the production hosts. The first approach, using the ECS system software to schedule production work on the AI&T platforms, is generally preferred, because it is not intrusive to the hardware. However, there may be circumstances when the second approach is simpler and faster. If the AI&T software environment is significantly different than the production environment (suppose AI&T was testing a new version of an operating system), and if the hardware failure in production is simple (a failed power supply), then changing hardware may be easier than changing software.

At GSFC, for Epochs c and g, a system has been configured for AI&T which has sufficient resources to take over the first-time processing responsibilities of any other system in the complement. At Epochs k and o, the AI&T system can provide failover support for the LIS/COLOR system; failover for MODIS first-time processing is provided within the suite of systems allocated to MODIS.

Data backup and recovery are a comparatively minor issue for SPRHW, because SPRHW does not provide long term, secure storage for data. Although SPRHW has large storage arrays, these arrays are used to hold ancillary data files and data granules for short periods of time. If a file system failure occurs within SPRHW, the algorithms, ancillary data files, and data granules can be recovered from the Data Server. Therefore the only data that would be backed up from SPRHW, and restored in the event of a failure, would be operating system and configuration files (i.e., the system disk). In the event that recovery is needed, incremental and full backup tapes would be used to re-build the required file systems.

Network Failure Recovery. The dual-ring implementation of FDDI provides a significant degree of fault tolerance. Most media failures within the FDDI fabric will not result in any loss of service and no reconfiguration would be necessary in these cases. Given the inherent fault tolerance of FDDI, it is not required to have multiple physical communications paths to each host. Hosts within SPRHW will use dual-attached station cards.

Failure recovery for the HiPPI switch will be supported by stocking spares for the Line Replacable Units of the switch (power supplies, interface cards, fan). If an individual interface card fails, a host can be re-configured to a hot spare interface card by moving two cables and sending a single software command to the switch; this process can be done in minutes, without disrupting other hosts. If the control module fails, it will be swapped out and the switch will be reconfigured. In the very rare event of an entire switch chassis failure, the switch would either be replaced or repaired. All of these failure recoveries involve activities on the switch; modifications to the attached hosts are generally not required.

3.4.2.7.2 AQAHW

QA of ECS products will include non-science QA, in-line QA, and SCF-based QA. Non-science QA will generally entail data integrity checks on data products and metadata. In-line QA is a form of science QA in which the content of the product is evaluated using science algorithms. Processing capacity for in-line QA, to the extent that it is specified in the AHWGP inputs and Technical Baseline, is included in the Science Processing HWCI (see Section 7, SPRHW). SCF-based QA is also a form of science QA, and is specified by the product development team. ECS provides support for SCF-based QA to the extent of providing archive and communications capacity for the SCFs to sample the products for QA.

The types of non-science QA to be performed at the DAACs will be specified by the DAAC operations staff. These requirements are as yet largely unspecified. The sections below define the working assumptions being made as to AQAHW requirements, and the specifications which flow from those assumptions.

3.4.2.7.2.1 AQAHW Requirements Analysis

The requirements for AQAHW are based upon the need to have a DAAC-based quality assurance capability to ensure the integrity of the products produced by the DAAC.

DAAC-based non-science QA processing requirements are to be defined through interactions with the DAAC operations personnel. The current design assumption is that DAAC-based non-science QA processing will be performed at the DAACs in parallel with the other forms of QA (SCF-based QA and in-line QA). The design baseline thus includes a local (DAAC) QA workstation to support these DAAC data integrity checks. This local QA workstation is actually similar to a Science User workstation equipped with core Client Subsystem functionality. The QA workstation acts as a client to the Data Processing and Data Server Subsystems. The current operations concept assumes that the QA workstation hosts ECS as well as DAAC supplied processes (as deemed necessary by the DAAC operations personnel), which use a subset of the ECS services to "pull" production data sets using the subscription mechanism. The need for visualization support will be explored as product specific QA processes and requirements are worked jointly with the DAAC operations teams.

3.4.2.7.2.2 AQAHW Technology Assessment

The AQAHW requirements identified to date do not present any significant technical challenges, and therefore no special technology assessment efforts (prototyping, benchmarking, or product evaluations) have been performed to support the specification of AQAHW.

3.4.2.7.2.3 AQAHW Specification

The AQAHW hardware suite at GSFC will consist of a graphics workstation connected to ECS via the PDPS FDDI subnetwork.

The AQAHW graphics workstation was selected to provide a software execution environment equivalent to the AI&T software execution environment in order to facilitate use of the AQA workstation for AI&T when necessary. The AQA workstation is also equipped with a graphics capability sufficient to support sophisticated visualization techniques.

The AQA workstation will be an SGI Indigo2 IMPACT 10000 workstation equipped with 128 MB of RAM. This workstation uses the R10000 chip, which is expected to be available in a 275 MHz implementation in the second half of 1996. The workstation will be equipped with an internal CD-ROM and two 4.3 GB internal disk drives. A SCSI-2 connection will be used to access an external disk array; the external disk array will consist of twenty 4.3 GB disks controlled by a single storage processor and managed in a RAID 5 configuration. This will provide 69 GB of external disk space to the workstation.

The Indigo2 IMPACT 10000 enclosure is a desktop configuration. The workstation will be configured with a 19 inch color graphics monitor.

The AQAHW workstation will reside on the ECS PLS/DPS FDDI network. This network will provide direct access between AQAHW components and the remainder of the PLS/DPS suite. The AQAHW components may also communicate with other ECS components and the external world via the DAAC FDDI switch.

3.4.2.7.2.4 AQAHW Design Discussion

Because the requirements for non-science quality assurance at the DAACs are not well defined, a minimal complement of AQAHW has been specified for each DAAC. This complement could be expanded as necessary in a variety of ways.

If GSFC expresses a need for additional workstations for QA personnel, but does not require visualization support for these staff, X-terminals could be assigned to the AQAHW to provide these seats. These X-terminals would be configured in the same way that the AI&T X-terminals are configured (see Section 3.4.2.7.3.3), and would be attached to the ECS PLS/DPS ethernet subnetwork.

If additional visualization seats were required, additional workstations (configured as above) could be added to the ECS PLS/DPS FDDI subnetwork to support the requirement.

If the data requirements from the Data Server (or from Science Processing) for QA are so large as to burden the ECS PLS/DPS FDDI subnetwork, the AQA workstation may be configured with a HiPPI interface and connected to the DSS/DPS HiPPI network. This would be subject to the availability of a port on the HiPPI switch, and a slot within the workstation. (The workstation graphics boards and interface boards contend for the same set of slots; a workstation configured with the maximum graphics capabilities offered for this product -- Maximum IMPACT graphics -- could not support multiple network interfaces.)

If additional memory is required in the AQA workstation, its memory can be expanded to a maximum of 640 MB.

The Indigo2 IMPACT 10000 workstation is offered with three graphics options; the entry-level option has been selected for use in the AQA configuration. If more sophisticated graphics are required, the graphics in the workstation can be upgraded, subject to the availability of a backplane slot as noted above.

3.4.2.7.3 AITHW

The Algorithm Integration & Test (AI&T) HWCI (AITHW) provides hardware resources to support the integration and test of science software at the DAAC, and system level validation, integration and test. It is important to note that this HWCI provides workstations and tools for software integration and test, but does not provide the compute environment or compute capacity required for science software test. This integration and test compute capacity is included in the Science Processing HWCI (SPRHW).

3.4.2.7.3.1 AITHW Requirements Analysis

The requirements for AITHW are based upon the need to have a software development, configuration management, and test environment at the DAAC to support the integration and test of science software delivered to the DAAC by the instrument teams.

The AI&T activity is expected to be characterized by intense activity at certain program milestones (the initial delivery of science software, for instance) and a much lower level of activity on an ongoing basis. There are no explicit ECS requirements identifying capacity, throughput, or response time requirements for AITHW; therefore, the sizing of the hardware has been based upon the size of the integration and test effort anticipated for each instrument and upon experience with the ECS Ir1 release.

The number of AI&T seats provided at each DAAC has been specified by allocating two seats for each instrument supported at the DAAC. At GSFC, two seats are provided for each of LIS, COLOR, and the three MODIS subsystems (Land, Atmospheres, and Oceans). The number of AI&T stations must be scalable to provide surge capacity during especially busy periods. The AI&T stations must provide the capability to support the AI&T tools, which provide graphical user interfaces. The AI&T stations do not have stringent reliability, maintainability, availability, and backup requirements, as they do not directly support production processing. The AI&T stations must be compatible with and interoperable with the target SPRHW hardware and the AI&T tools server.

The AI&T tools server must provide sufficient capacity to support the use of the development, configuration management, and test tools by the AI&T stations. Because the number of AI&T stations at GSFC will be ten, and the tools server is not also acting as the software build or test environment (which is provided by SPRHW), the AI&T server requirements are minimal. The AI&T tools server does not have stringent reliability, maintainability, availability, and backup requirements, as it does not directly support production processing. The AI&T tool server must be compatible with the AI&T stations and the target SPRHW hardware.

AITHW must provide a network printing capability to support the AI&T task.

3.4.2.7.3.2 AITHW Technology Assessment

The AI&T hardware requirements do not present any significant technical challenges, and therefore no special technology assessment efforts (prototyping, benchmarking, or product evaluations) have been performed to support the specification of AITHW.

Because the CPU resources required for AI&T will be provided by target machines in the SPRHW suite, or by the AI&T tools server, the AI&T station requirements can be met with X-terminals. The use of X-terminals for these stations has several advantages, including cost effectiveness, ease of maintenance, expandability, and compatibility.

3.4.2.7.3.3 AITHW Specification

The AI&T hardware suite at GSFC will consist of ten X-terminals, one server, and one network printer. These components will be connected to ECS via the ECS local area network at the DAAC.

AITHW COMPONENTS. The AITHW components will consist of workstations, a tools server, and a network printer.

AITHW Workstations. The requirements for AITHW workstations will be met at minimum cost by providing NCD HMX-PRO X-terminals. These X-terminals will be configured with 20" color monitors and 16 megabytes of memory to support the use of the tools specified for the AI&T tools server. Users will log on to either the AI&T tools server to use its software, or to a target machine in the SPRHW suite to build and test software in that environment.

The X-terminals will be configured to allow the addition of memory if experience indicates this is necessary; experience with Ir1 resulted in upgrades to the memory of X-terminals purchased for Ir1. Additional X-terminals can be purchased and integrated quickly and easily if expansion of the AI&T seating capacity at a DAAC becomes a critical requirement. Other ECS operations workstations, such as those in the Algorithm Quality Assurance HWCI and the Planning Subsystem HWCI, can also be used to support short term increases in AI&T activity.

Note that GSFC received a Sun workstation (AITHW-GSFC-1) for use as an AI&T workstation in Release A. This workstation will be replaced by an X-terminal, and will be re-used elsewhere in an ECS DAAC. GSFC also received two X-terminals (SPRHW-GSFC-2, and -3) as part of its Ir1 and Release A science processing complement. Although these X-terminals will not be re-designated (that is, with AITHW-GSFC-N designators), they will be included in the count of AI&T workstations.

AITHW Server. GSFC was provided with a Sun 20/50 SPARCstation for Ir1 and Release A. This system is configured with 128 MB of RAM and four gigabytes of internal disk. The processing requirements for the AI&T server are minimal and this equipment is expected to be satisfactory for Release B.

AITHW Network Printer. The AITHW network printing requirement will be satisfied by the printer provided for Ir1, an HP LaserJet 4M+ with 14 megabytes of RAM. This printer has 12 page per minute throughput.

AITHW INTERFACES. The AITHW server will reside on the ECS PLS/DPS FDDI subnetwork. The AITHW workstations and printer will reside on the ECS PLS/DPS ethernet subnetwork. These networks will provide direct access between AITHW components and the SPRHW suite. The AITHW components may also communicate with other ECS components and the external world via the DAAC FDDI and ethernet switches.

3.4.2.8 MSS and CSS Subsystems

The MSS and CSS Subsystem hardware have been sized and configured in a redundant configuration in order to provide for high availability of communications infrastructure and management services. The sizing rationale, therefore, applies to both MSS and CSS servers and will be presented in a single subsection.

The MSS Subsystem consists of a single hardware configuration item (MSS-MHWCI), which provides the servers, workstations, and printers needed for all local system management functions. The MSS-MHWCI provides processing and storage for the following MSS software components:

- Management Software Configuration Item (MCI) - provides system monitoring and control (via HP Openview), the database management system (Sybase), trouble ticketing (Remedy), fault and performance management (Tivoli), physical configuration management (Accugraph), security management, accountability management, billing and accounting system, mode management service, performance trending capability, report generation and distribution, and management data access (custom code / scripts used to import log file data to the relational data base management system)
- Management Logistic Configuration Item (MLCI) - Site and SMC maintenance and operations staffs will rely on configuration management to provide software change control (Clearcase), change request management (DDTS), baseline management (XRP), inventory/logistics/maintenance (ILM) management, training management, policy and procedure management, software distribution management (Tivoli), and software license management.
- Management Agent Configuration Item (MACI) - Agents are processes used to monitor and/or control managed objects distributed across heterogeneous platforms. Current COTS technology for network management uses network protocols such as simple network management protocol (SNMP) to provide a way for the manager, the managed objects, and their agents to communicate. SNMP defines specific messages, referred to as commands, responses, and notifications.

The CSS Subsystem consists of a single hardware CI (CSS-DCHWCI), which provides the server for all CSS functionality. CSS contains a single CI, the Distributed Communications CI, which provides the following services:

- Common Facility Services - includes electronic mail, file access, bulletin board, virtual terminal, and event logger services

- Object Services - includes security, naming, message passing, event, thread, time and life cycle services
- Distributed Object Framework - includes OODCE framework functionality.

3.4.2.8.1 Rationale

The MSS/CSS processing complement for GSFC was designed and sized for both the TRMM and AM-1 missions. The sizing of MSS/CSS subsystem hardware is based on the February 1996 version of the technical baseline. Storage requirements have been rounded upward.

Processing Requirements Processing requirements for the MSS and CSS subsystem are driven by the following types of transactions:

- HP Openview data collection from managed objects and ad hoc queries (server)
- Conversion / import of HP Openview and log file data to MSS Sybase DBMS (server)
- DBMS usage for report generation / ad hoc queries (server)
- Fault & performance management notification (server)
- Trouble ticketing (server)
- Order request tracking (server)
- Billing & accounting (workstation)
- Mode management (server)
- Usage for configuration, baseline, training, license, inventory, change request, software distribution inventory/logistics/maintenance and associated report generation (workstation)
- DCE logical server transactions (directory, security, time).

Server Sizing ECS already has experience with many of the COTS products to be loaded on the MSS server from previous work in Evaluation Packages (EPs) and EDF installations. Based on this experience, a profile of the MSS/CSS server that is operating under nominal load (e.g., HP Openview map is displayed, but no collections are in process) has been developed. To this, processing requirements have been added for specific types of transactions.

In the EDF, an HP 9000/735/125, rated at 160 MIPS, was loaded with HP Openview, DCE client, Sybase server, X-server, and operating system. Tests were run to examine the impact of various types of HP Openview functions on CPU utilization. HP Openview was configured to discover approximately 500 nodes within EDF and then displayed them as a node map. Minimal status polling was performed at 15 minute intervals. A variety of HP Openview on-line reports were generated to show such items as packet throughput and CPU utilization. During the testing, processes resident on the server were monitored. CPU utilization remained extremely low (i.e., less than 3%) except during operator queries and initialization. At system start-up, initialization of the various daemons used by HP Openview generated a load of approximately 50%. After start-up, functions that involved initialization of x-windows screens (e.g., generation of the node map or display of a performance graph) generated loads of 25-40% for a brief (less than 15

seconds) period of time. Multiple SNMP queries on a router increased CPU usage to approximately 20 percent, with the primary driver appearing to be the x-windows server. Simultaneous queries of two routers (to two different x-window screens) consumed a total of 50-60% of the CPU. Based on this benchmark, we assume that a basic configuration of a server, including HP Openview, Sybase, DCE client, and the operating system will require approximately 72 MIPS, and will provide adequate resources for routine HP Openview operations. To this must be added processing capacity to handle DCE server functions, HP Openview monitoring, processing of log files, Sybase report generation / ad hoc query capability, Remedy trouble ticketing, Tivoli monitoring, Tivoli software distribution management, mode management, order request tracking, and mail.

HP Openview and log file-to-Sybase data conversion are primary processing drivers that are expected to vary by DAAC. Table 3.4.2.8.1-2 shows estimated numbers of transactions for HP Openview data collection. HP Openview data collection is driven by the number of managed objects to be monitored and the number and frequency of management information base (MIB) objects to be collected for each. Managed objects for each MIB type were counted based on the Release B hardware plan for GSFC. The number and frequency of data collection for each class of managed objects was provided by MSS developers as specified in the CSMS Database Design and Database Schema Specification, (311-CD-003-003, Appendix B). HP Openview provided an estimate of 100,000 instructions per transaction. Using this information, an average number of instructions per second required for HP Openview data collection was developed. These estimates appear to be reasonably in line with HP-provided performance information, which indicates that an HP 9000/735, a machine rated at 125 MIPS, is capable of performing approximately 1300 collections per second.

An estimate of 100,000 instructions per transaction was assumed for the conversion of each logged event to Sybase, based on the number of source lines of code for the MSS MDA component involved and an estimate of instructions needed to update the Sybase database. Instructions per transaction was multiplied by the number of logged events, including both HP Openview events and events collected from applications via the logging API. HP Openview events (transactions) are described in the previous paragraph. The number of application-generated log entries, see Table 3.4.2.8.1-1 was developed using the following assumptions:

- One log entry is generated for every system transaction, by every process that is included in the transaction thread.
- The number of “pull” transactions is based on the user model and reflects user service requests by DAAC. Pull transactions (e.g., directory, inventory search requests) are assumed to generate a conservative estimate of 5 log entries each from CIDM and data server processes.
- Order request tracking is dependent on the request for data by a user and the request for status of a data product by a user. For every user request for data, an EcRequest is stored in the management DBMS and updated as required by the DSS. The transaction frequency for EcRequest storage is related to the number of granules requested by user per DAAC. Updates made by the DSS to the EcRequest are considered to be a small

percentage of the total granules requested. For every user request for status of a product, the appropriate EcRequest is retrieved from the management DBMS and made available to the client.

- The number of transactions on the “push” side includes external (DAAC-to-DAAC and L0) file transfers (4 log files each), processing-to-archive requests (4 log files each), and PGE execution (2 log files each). Push transactions were based on AHWGP data, which showed that GSFC executes 19,332 PGEs per day and each PGE requires an average of 39 input/output file requests at Release B.
- In addition, major processes generate log entries of approximately 512 K (based on the MSS application MIB) once every 15 minutes. There are estimated to be 15 processes at each DAAC that will generate log entries every 15 minutes.
- Log files and HP Openview data will be kept for 14 days prior to archiving in long-term Data Server Storage.
- For Billing & Accounting, there are expected to be approximately 17,000 total daily user accounts. Each account will be logged and on demand available for information tracking. An approximate number of user accounts per DAAC was estimated from the February 1996 user pull technical baseline.
- To implement the mode management service, multiple modes are assigned to each logged activity and can be simultaneously executed. The overhead required to provide mode management capability is estimated at 30% of the total logged activity.

The MIPs required to import the total number of log files per day are given in Table 3.4.2.8.1-3.

Table 3.4.2.8.1-1. GSFC Log Entry Storage Volume - Release B

Log File	Log Events per Transaction	Transaction Frequency per Hour	Total Logged Events per Hour	Bytes per Transaction	Total Size of Bytes/Hr	14-Day Storage Requirements (MB)
User requests	10	819	8,190	420	3,439,800	1,156
Request tracking	10	410	4,100	420	1,722,000	579
PGE execution	2	403	806	420	338,520	114
External file transfers	4	19	76	420	31,920	11
Processing-to-Archive Requests	4	15,717	62,868	420	26,404,560	8,872
Application MIB poll	15	7	105	512	53,760	18
Billing & accounting activity logging	2	142	284	420	119,280	40
Subtotal Rel B			76,429		32,109,840	10,790
Total Rel B (includes x30% for multiple modes)			99,358		41,742,792	14,027

Table 3.4.2.8.1-2. GSFC HP Openview Collection Processing Requirement

	# MIB Objects	Average Size (Bytes)	GSFC Managed Objects	Collections per hour*	Collections per second	Estimated MIPS
Release B GSFC (Hosts, RDBMS, Router, hubs)	1,953	4	294	241,363	67	7

* Note that the number of collections per hour was derived by multiplying each class of MIB objects (e.g., MIB II objects) by the number of managed objects within that class, and summing the results.

Table 3.4.2.8.1-3. MDA Data Conversion to Sybase Processing Requirement

	Total HP Openview Events / Day	Total Log File Events / Day	MIPS for 8 hour Sybase import
Release B	5,792,712	2,384,592	28

At Release B, ad-hoc queries will be performed and statistical analysis collected from the Sybase database. Ad-hoc reports will be generated that include the following type of information; user accesses, trend analysis, fault occurrences, resource utilization, data production jobs and security events. Benchmarks are being run on a prototype Sybase database to evaluate performance. The prototype database was developed to do real-time benchmarking queries of designated working attributes that are expected to be of reporting interest (i.e. performance).

DCE has been installed in the EDF and used in the Engineering Packages (EPs). Running on an HP 715, rated at 77 MIPS, the DCE server functions used 8% of the CPU, or approximately 6 MIPS. An analysis was performed to determine how much additional load would be placed on the DCE server at Release B.

Load imposed on the DCE server is a function of the number of directory, security and time look-ups from client applications. A client application maintains its own cache containing the most recently accessed directory and security information, and will only access the server when a user is not found in its own cache. Many client applications will only access other clients in the DAAC, and so will never exceed their cache. CIDM and the Data Server APC, however, will be directly accessed by external user clients and so will need to access directory and security information for each user access. At GSFC, the user model reflects a maximum of 819 users accessing per hour. Given that a directory and security lookup typically requires less than 0.5 seconds, it is unlikely that there will be more than 1-2 simultaneous hits on the DCE server. We estimate that 1 additional MIP processor capacity will be sufficient for the level of DCE accesses required.

In the EDF "Mini-DAAC" facility, Tivoli performance was evaluated on an HP 9000 J210/1 rated at 176 MIPS and with 256 MB of RAM. The Tivoli COTS package will be used primarily for performance management, fault management and software distribution (most likely through Tivoli Courier). Performance will be monitored, statistics collected, and faults detected via Tivoli GUI screens. Tests were run to determine Tivoli GUI screen CPU utilization. The benchmark was performed with one user and the following configuration; the Tivoli application, the platform operating system, xwindows, and the performance tool (glance plus). CPU utilization was minimal as expected with no applications running (.5% of the system CPU) and approximately 56 processes active.

Following initialization of the Tivoli application, CPU utilization remained low (<2%), with the Tivoli Management Enterprise (TME) desktop enabled and 61 processes active. CPU loading became more prevalent when an administration GUI was selected from the TME desktop. Peak utilization was recorded at 9% of the system CPU for a period of 10 seconds and 73 processes active. Max peak CPU utilization (approximately 11%) and IO throughput (13.5 MB/s) was recorded when enabling the policy region desktop. In steady state, CPU utilization measured approximately 3% of the system CPU. Opening multiple GUIs did not increase demands on CPU utilization (remained <11%) but linearly required more memory. The Tivoli vendor for this reason recommends 96 MB of RAM dedicated. Total CPU utilization allocation for Tivoli based on benchmark results is estimated at 11% of the system CPU or 20 MIPS. The targeted platform

at each DAAC site will be upgraded from the platform used for benchmark calculations. Recognizing the emphasis by Tivoli for memory and moderate processing needs, additional processing and memory capabilities were added to the MSS management server in Release B to provide adequate resources in support of the Tivoli product.

Remedy was evaluated on the HP 9000/735/125 for CPU utilization. The application required very little CPU allocation (<1%). A more significant load was present when performing browse or ticket assignment functions (approximately 6%). Submittal and processing of a trouble ticket required less than 1% of the CPU capacity.

The server requirements, as dictated by the rationale given above, is synopsisized in Table 3.4.2.8.1-4.

Table 3.4.2.8.1-4. CSS/MSS Server Configuration - Requirements Estimate

Server Load Sources	Estimated R-B MIPS
Basic configuration (includes HP Openview and DCE client)*	72
Additional HP Openview data collection*	7
Sybase Server and Client*	50
Tivoli *	20
Remedy	11
MDA (log conversion to Sybase)	28
MSS Agent*	3
DCE server (including additional processing for peak directory and security transactions)*	7
Word Processor	1
Spreadsheet	1
Other Common Services (Mail, file transfer, etc.)*	5
Total	175
* These items were considered to be potentially active at the same time. MDA database update is assumed to be run in off-peak hours, and not concurrently with Sybase report generation functions.)	

Workstation Sizing There will be two MSS workstations at each DAAC site. Workstation # 1 will primarily contain the MLCI software. This includes software change management (clearcase), change request management (DDTS), baseline control management (XRP), software license manager and inventory/logistics/maintenance (ILM) management. Policy & procedures management and training management will be configured on workstation # 2. Each MSS Workstation will contain the Sybase client, DCE client, Tivoli client, MSS agent, and operator tools.

The DDTs tool was evaluated for performance in the EDF facility on a Sun SPARC 20/50 rated at approximately 130 MIPS and 64 MB of RAM (DAAC targeted platform will be an upgrade version). DDTs is the change request manager and maintains and tracks potential changes (via configuration change requests) to the ECS System. Configuration change requests (CCRs) will be created, logged into the DDTs database and tracked by a CM specialist. Tests were performed to determine CPU utilization for implementation of these tasks. The benchmark was performed with one user, the DDTs application configured with the platform operating system, xwindows and the performance tool (proctool). Following initialization of the DDTs application, CPU utilization as expected was very low, <1%. For each instance that a CCR was either submitted, modified or logged, the CPU utilization remained below 3% and memory utilization less than 8%. Processing utilizations increased significantly when queries were made to the DDTs database for tracking purposes as specified in Table 3.4.2.1-5;

Table 3.4.2.8.1-5. DDTs Benchmark Results

Benchmark Test	# Records	CPU Utilization (% of system CPU)	Memory Utilization (% of system memory)	IO Throughput (KB/s)
CCR submittal/creation	-	1.4 %	6.6 %	4
CCR registration	-	2.3 %	7.1 %	4
EP4 database query	128	9.7 %	7.2 %	11.3
EP6 database query	279	13.5 %	7.3 %	11.7
EP4 + EP6 database query	407	16.5 %	7.5 %	11.8
DDTs (inclusive) database query	1232	30.8 %	7.5 %	12.2

CPU utilizations ranged from 10% of the system CPU to approximately 30%. The number of records in the development environment is expected to be substantially higher than at the DAAC sites. For this reason, a conservative estimate of a maximum of 400 records is used to result in a CPU utilization allocation of approximately 21 MIPS. Memory utilization and IO throughput were moderate and appeared constant for each test performed.

Processing requirements for baseline management COTS (XRP), was estimated from vendor specifications. For a 30 user system, the XRP vendor specifies a processing requirement of 100 MIPS. Each DAAC site is assigned to have 2 XRP users and therefore will require approximately 7 MIPS.

Vendor specifications suggest an allocation of 35 MIPS for the Clearcase Virgin Object Base (VOB) Server. The VOB server is the most compute intensive of the Clearcase server applications due to its required database processing. In the EDF, Clearcase was installed on a SPARCstation 10, equipped with 120 MB RAM, rated at 109 MIPS, and with an ethernet interface. The SPARCstation 10 was initially used for Tool kit development, as well as CM of the Evaluation Packages. With moderate numbers of users, the SPARCstation 10 provided good

performance. At peak use (15-20 simultaneous users viewing items, manipulating the contents of the database, and executing directly out of Clearcase), performance was adversely affected. Usage at the DAAC is not anticipated to require more than 5 simultaneous users, frequency of use is anticipated to be much lower, and applications will not be executed from the Clearcase tool. Additional benchmarks will be run as ECS code and science algorithms become available to help determine the precise Clearcase processing requirements at the DAAC. EDF experience suggests that a workstation configuration in the SPARCstation 20 range should be adequate to support Clearcase, other MLCI COTS, DCE and billing & accounting and tivoli clients.

On MSS Workstation #2, the COTS products expected to exert the larger processing loads are billing and accounting (B&A) and performance trending. Other primary load contributions come from training management, policy and procedures, and the DBMS report generator.

Major B&A processing loads will occur during nightly batch imports to the sybase server which is not resident to MSS Workstation #2. Processing of B&A transactions such as accounts received, purchase orders placed and products delivered are expected to exert a moderate load on the MSS Workstation due to the expected number of user requests for data as provided by user modeling.

Selection for the performance trending statistical analysis package is in progress. Statistical and historical performance data will be analyzed to assure optimum usage of system resources. A determined number of performance attributes will be analyzed by the statistical tool. Performance trending and other resident COTS packages such as training manager, policy and procedure manager and DBMS report generator are expected to require a small to moderate load on the MSS Workstation.

Tables 3.4.2.8.1-6 and 3.4.2.8.1-7 below reflect a best estimate of load to be imposed on each MSS workstation. It assumes that most functions run concurrently. Operator functions can be spread across workstations in such a way as to balance processing loads.

Table 3.4.2.8.1-6. MSS Workstation #1 Configuration - Requirements

Workstation Load Sources	Estimated MIPS
Basic configuration (includes Clearcase and Operating System)*	50
Software License Management*	5
DDTS*	21
XRP*	7
Tivoli Client*	5
Sybase Client*	10
Word Processor	1
Spreadsheet	2
Graphics	1
Inventory/Logistics/Maintenance Management	15
MSS Agent*	2
DCE Client*	5
B&A Client*	5
Other Common Services (Mail, file transfer, etc.)	5
Total	125
* These items are considered to be potentially active at the same time	

Table 3.4.2.8.1-7. MSS Workstation #2 Configuration - Requirements

Workstation Load Sources	Estimated MIPS
Basic configuration (includes Billing & Accounting and Operating System)*	35
Training Management*	5
Performance Trending*	15
DBMS Report Generation*	10
Policy & Procedures*	5
Tivoli Client*	5
Sybase Client*	10
Clearcase Client*	10
Word Processor	1
Spreadsheet	2
Graphics	1
MSS Agent*	2
DCE Client*	5
Other Common Services (Mail, file transfer, etc.)	5
Total	102
* These items are considered to be potentially active at the same time	

Storage Requirements Major datastores for the MSS and CSS subsystems include: HP Openview files, application log files (including request order tracking, billing & accounting and mode management), the Management DBMS, and Clearcase-managed data for software change management.

Other datastores include DCE directory, security data, mail, trouble ticketing, Tivoli, DDTs, baseline control management data (XRP), ILM, billing & accounting client, training management, policy & procedures, and DBMS report generation.

The size of the data storage for HP Openview has been estimated from the determination of the frequency of transmission of the necessary information of all the appropriate attributes of the managed objects during one hour period. It was assumed that fourteen days worth of HP Openview data are stored.

A description of how application log file volume was estimated is in the previous section (Processing Requirements). Log file volume is provided in Table 3.4.2.8.1-1 based on an assumption of fourteen days storage prior to archiving in the data server archive.

The storage requirement for the Management DBMS was based on a worst case assumption that all the records from both the log files and HP Openview are stored in the Management DBMS, with an additional 10% for table overhead and summarization records. It is assumed that one month's worth of data are maintained in the Management DBMS at a time.

Storage requirements for Clearcase are based on the assumption that Clearcase will store two copies of all source code (including ECS application source and algorithms) and two copies of all executables. This will enable recovery of the previous version of any application if required. In addition, Clearcase will store test data and configuration files.

Tivoli sizing estimates are based on the number of performance attributes that will be monitored as specified by MSS developers in the Release B CSMS System Management Subsystem Design Specification, (305-CD-029-002, Section 6.6). These include system, application, process, and disk performance metrics. The monitoring frequency is dependent on the performance attribute. A worst case polling frequency of once per minute for all attributes was used in sizing calculations. The size of a typical fault/performance notification was estimated at 256 bytes.

Approximately 400 trouble tickets per day are estimated to be assigned, or approximately 17 per hour. The size of a trouble ticket is approximately 256 characters. Trouble ticket frequency and size are worst case.

To determine DDTs sizing requirements, the frequency non-conformance reports (NCRs) are generated on a daily basis was identified with the report size. A NCR was evaluated due to its similarity to a configuration change request (CCR). The number of CCRs generated at the DAAC sites are not considered to be more necessary in a developmental environment.

There are expected to be approximately 16 periodic reports that will be produced on a daily, weekly, monthly and annual basis. Reporting areas include data production, fault identification, user accesses, resource utilization, user services activity and trouble ticketing. The size of an average report is estimated to be 15 KB. The aggregate number of reports generated is approximately 1 per hour.

The cumulative datastores of XRP, ILM, billing & accounting client, training management, and policy & procedures was estimated based on vendor provided information and experience in the development facility.

Disk space requirements of the MSS management server COTS applications are listed in Table 3.4.2.8.1-8. These applications will be stored in RAID and available for download to local disk. The RAID device interface is Fast/Wide SCSI which offers application access times comparable with local disk.

Additional RAID storage is allocated for safeguard of HP Openview functions and storage of billing and accounting transaction logs. Other real time functions (i.e. Tivoli and Remedy) will be replicated to the CSS server. A copy of all management data will be stored in RAID on a daily basis and safestored into a DLT tape drive via the management backup server. As required, the management data will then be stored into ECS data server archive.

Table 3.4.2.8.1-8. COTS Product Disk Requirements

COTS Product	Disk Requirement (MB)
HPOV	2,000
Tivoli	100
Trouble Ticket (Remedy)	50
Sybase Server	1,000
Clearcase	2,000
DCE Server	200
Accugraph	50
TOTAL	5,400

The storage requirement for the Sybase DBMS is estimated to be 31 GB, Clearcase 3.9 GB, Tivoli 114 MB, Remedy 2 MB, DDTS 7 MB, DBMS Report Generator 5 MB, and 145 MB for all other datastores combined. The total storage requirement for Release B is estimated to be between 58 and 61 GB (includes additional storage for Sybase swap space). See Table 3.4.2.8.1-9 for storage requirements.

Storage requirements for DCE directory and security stores are based on the number of predicted users as provided by user modeling. The total storage requirement for CSS is estimated to be 496 MB for Release B, as specified in Table 3.4.2.8.1-10.

Table 3.4.2.8.1-9. GSFC MSS Release B Storage Requirements

Datastore	Freq. of Events/Hr	Size in Bytes/ Transaction	Size in Bytes Transmitted/Hr	14-Day Storage Requirements (MB)
HP Openview Datastore	241,363	4	965,452	324
Application log files	99,358	420*	41,742,792	14,027
Sybase DBMS				31,572
Tivoli	1320	256	337,920	114
Remedy	17	256	4,352	2
Clearcase				4,000
DDTS	8	2400	19,200	7
DBMS Report Generator	1	15,000	15,000	5
Other Datastores (ILM, XRP, B&A client, training, policy & procedures)				145
Application Disk Space Requirements				5,400
Total Storage Requirement				55,623

* Application polling generates 512 byte logs. These have been included in the per hour total.

Table 3.4.2.8.1-10. GSFC CSS Release B Storage Requirement

CSS Data Store	# of Users	Size of Record (# Bytes)	14-Day Storage Requirements (MB)
DCE Directory	17,000	1,000	238
DCE Security	17,000	1,000	238
Mail	348	4,000	20
Total Storage Requirement			496

Processor Selection Choice of the MSS/CSS Server platform was based on Release B processing requirements, COTS to be hosted on the platform, and price/performance data provided by EDS. Based on the Release B processing requirements, a medium-range server class platform was chosen. HP is the preferred vendor, since HP Openview and OODCE will be principal COTS products on these platforms, and HP is one of the principal developers of DCE and OODCE.

3.4.2.8.2 Configuration

The following configuration will be provided for the GSFC LSM for Release B, which includes the MSS MHWCI and the CSS DCHWCI.

- MSS Local Management Server and CSS Communications Server: 2 HP 9000 J210/2 processors, 384 MB of RAM and 4 GB of storage.
- RAID Storage: 55 GB total storage
- Workstations:
 - 1 Sun SPARC Ultra 170, 128 MB of RAM and 8 GB of storage (This workstation will house configuration management software)
 - 1 Sun SPARC Ultra 170, 128 MB of RAM and 4 GB of storage
- Management Data Backup Server
 - 1 Sun Ultra 4-slot with 128 MB of RAM and 4 GB of storage
- Printer
 - 1 HP Laser Jet 4M+ Printer, 12 ppm/14 MB

The GSFC DAAC will contain two primary servers for its LSM configuration, cross-strapped to RAID disk to enable warm backup. MSS and CSS applications will run on separate systems but in case of contingency, either system will be capable of running both subsystems.

The HP 9000 J210 is a high performance processor specifically designed for compute intensive and graphic applications.

The configuration at GSFC will include two Sun SPARC Ultra 170 workstations. One of the workstations which will house configuration management software will be configured with higher memory and higher storage (128 MB of RAM and 8 GB of hard drive).

The DAAC configuration at GSFC includes a Digital Linear Tape (DLT) drive which provides better performance transfer rates and higher capacity per cartridge than standard helical scan technology.

3.5 Software/Hardware Mapping

With the exception of the Client subsystem, each subsystem has been designed to incorporate hardware CIs that include the components (processors, servers, archive robotics, etc) on which the software components run. The Management and Communications subsystems include software components that execute on hardware within these subsystems, as well as on hardware across all the subsystems with hardware CIs. Table 3.5-1 provides a mapping of GSFC ECS software components to the applicable hardware components.

Table 3.5-1. GSFC Hardware to Software Component Mapping (1 of 9)

HWCI / units	Subsystem	CSCI	CSC
SPRHW/science processors	Data Processing	AITTL	Binary File Comparison Utility Code Analysis Tools HDF File Comparison Utility Profiling Tools Standards Checkers
		PRONG	Data Pre-Processing PGE Execution Management Resource Management
		SDPTK	All CSCs
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
SPRHW/queuing management server	Data Processing	PRONG	PGE Execution Management Resource Management Product
	Planning	PLANG	All CSCs

Table 3.5-1. GSFC Hardware to Software Component Mapping (2 of 9)

HWCI / units	Subsystem	CSCI	CSC
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
AITHW/AI&T DBMS server	Data Processing	AITTL	Binary File Comparison Utility Code Analysis Tools Data Visualization Tools Documentation Viewing Tools ECS HDF Visualization Tools HDF File Comparison Utility Profiling Tools Report Generation Tools Standards Checkers
		PRONG	COTS COTS Management Data Management
		SDPTK	All CSCs
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
AITHW/AI&T Operations workstations	Client	DESKT	All CSCs

Table 3.5-1. GSFC Hardware to Software Component Mapping (3 of 9)

HWCI / units	Subsystem	CSCI	CSC
		WKBCH	All CSCs, except - User Registration Tools
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Data Processing	AITTL	Data Visualization Tools Documentation Viewing Tools ECS HDF Visualization Tools PGE Processing GUI PGE Registration GUI Product Metadata Display Tool SDP Toolkit-related Tools SSAP Processing GUI Update Data Server GUI
		PRONG	COTS COTS Management Data Management
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
AQAHW/QA workstations	Client	DESKT	All CSCs
		WKBCH	All CSCs, except - User Registration Tools
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Data Processing	PRONG	Quality Assurance Monitor Interface

Table 3.5-1. GSFC Hardware to Software Component Mapping (4 of 9)

HWCI / units	Subsystem	CSCI	CSC
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
PLNHW/PDPS DBMS server	Planning	PLANG	All CSCs
	Data Processing	PRONG	PGE Execution Management Resource Management
	Client	DESKT	All CSCs
		WKBCH	All CSCs, except - User Registration Tools
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
ICLHW/ingest server	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Ingest	INGST	All CSCs
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy

Table 3.5-1. GSFC Hardware to Software Component Mapping (5 of 9)

HWC1 / units	Subsystem	CSCI	CSC
	Client	DESKT	All CSCs
		WKBCH	All CSCs, except - User Preference Tools
ACMHW/administration and operations workstations	Client	DESKT	All CSCs
		WKBCH	All CSCs, except - User Preference Tools
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
ACMHW/APC servers	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Data Server	SDSRV	All CSCs
		STMGT	Service Clients File Peripherals Resource Management
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy

Table 3.5-1. GSFC Hardware to Software Component Mapping (6 of 9)

HWCI / units	Subsystem	CSCI	CSC
DIPHW/distribution servers	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Data Server	DDIST	All CSCs
		STMGT	Service Clients Peripherals Resource Management
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
DRPHW/FSMS servers	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Data Server	STMGT	All CSCs
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
DRPHW/DBMS servers	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs

Table 3.5-1. GSFC Hardware to Software Component Mapping (7 of 9)

HWCI / units	Subsystem	CSCI	CSC
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
	Data Server	SDSRV	DB Wrappers (Illustra DBMS)
DDSHW/document server	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Data Server	DDSRV	All CSCs
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
DMGHW/data specialist workstations	Client	DESKT	All CSCs
		WKBCH	All CSCs
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
DMGHW/administration and operations workstations	Client	DESKT	All CSCs
		WKBCH	All CSCs

Table 3.5-1. GSFC Hardware to Software Component Mapping (8 of 9)

HWCI / units	Subsystem	CSCI	CSC
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
DMGHW/DBMS servers	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	Data Management	DDICT	All CSCs
		DIMGR	All CSCs
		GTWAY	All CSCs
		LIMGR	All CSCs
	Interoperability	ADSRV	All CSCs
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
MSS/MSS workstations	Client	DESKT	All CSCs
		WKBCH	All CSCs

Table 3.5-1. GSFC Hardware to Software Component Mapping (9 of 9)

HWCI / units	Subsystem	CSCI	CSC
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
	ISS	INCI	All CSCs
	Management	MACI	All CSCs
		MCI	All CSCs
		MLCI	All CSCs
		MHCI	All CSCs
MSS/MSS LSM Server	ISS	INCI	All CSCs
	Management	MACI	All CSCs
		MCI	All CSCs
		MLCI	All CSCs
		MHCI	All CSCs
	Communication	DCCI	Electronic Mail Services Event Logger Services File Access Services Life Cycle Services Message Passing Services Thread Services Time Services
CSS/CSS server	Communication	DCCI	All CSCs
	Management	MACI	All CSCs
		MCI	Automatic Actions Fault Management Management Proxy Performance Management Performance Management Proxy
User workstation	Client	WKBCH	All CSCs
		DESKT	All CSCs
	Ingest	INGST	User Network Ingest Interface

Note: "All CSCs" refers to those CSCs for a specific CSCI that is identified in the Component Analysis Table in Section 3.3.2 of this document.

This page intentionally left blank.

4. Future Releases

This document has described the design of ECS subsystems for the GSFC DAAC at Release B. Two other releases are currently planned. The next release, Release C, is scheduled for December 1999. The Release Plan Content Description for the ECS Project describes in detail the capabilities being provided in Releases C and D. In summary, Release C will provide evolutionary enhancements to Release B, Flight Operations for EOS PM-1, and the CORBA and Trader designs.

An updated version of this document will be generated prior to Release Readiness Review (RRR) and will resolve some of the issues highlighted in this version of the document. In particular, areas of resolution include LIS production activities, migration of MSFC V0 data sets, integration of DAO/DAS into the ECS design at the GSFC DAAC, and incorporation of live data streams into ECS (i.e., SeaWiFs).

This page intentionally left blank.

Appendix A

Table A-1 provides detailed Science Processor Hardware (SPRHW) configurations for Release A and Release B.

Table A-1. Detailed SPRHW Configurations for Release A and Release B

Release A	C 3Q97	G 3Q98	K 3Q99	O 3Q00
SPRHW-GSFC-1 Function: GSFC AI&T Cabinet: Power Challenge XL Console: Character CPU: 12 x 75 MHz R8000 RAM: 512 MB/1-way interleaved	SPRHW-GSFC-1 Function: GSFC AI&T Cabinet: Power Challenge XL Console: Character CPU: 16 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-1 Function: GSFC AI&T Cabinet: Power Challenge XL Console: Character CPU: 16 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-1 Function: MODIS (1 of 4) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-1 Function: MODIS (1 of 6) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved
VME-1: FDDI IO4: One HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI SCSI-0 (1,0,1): CD-ROM, 8 mm Tape Stacker SCSI-1 (1,0,2): One 2 GB Internal Disk SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-2 SP1 SCSI-4 (1,2,3): RAID-3 SP2	VME-1: FDDI IO4: Two HIO-1 (1,1): HiPPI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): Unused SCSI-0 (1,0,1): CD-ROM, 8 mm Tape Stacker SCSI-1 (1,0,2): One 2 GB Internal Disk, Two 4.3 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2	VME-1: FDDI IO4: Two HIO-1 (1,1): HiPPI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): Unused SCSI-0 (1,0,1): CD-ROM, 8 mm Tape Stacker SCSI-1 (1,0,2): One 2 GB Internal Disk, Two 4.3 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2	VME-1: FDDI IO4: Two HIO-1 (1,1): HiPPI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): One 2 GB Internal Disk, Two 4.3 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2	VME-1: FDDI IO4: Two HIO-1 (1,1): HiPPI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): One 2 GB Internal Disk, Two 4.3 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2
RAID-1: 20 x 4.3 GB (Rack 1) RAID-2: 20 x 4.3 GB (Rack 1) RAID-3: 10 x 4.3 GB (Rack 1)	RAID-1: 10 x 4.3 GB RAID 5 (Rack 1) RAID-2: 10 x 4.3 GB RAID 5 (Rack 1) RAID-3: 10 x 9 GB RAID 5 (Rack 1) RAID-4: 10 x 9 GB RAID 5 (Rack 1)	RAID-1: 10 x 4.3 GB RAID 5 (Rack 1) RAID-2: 10 x 4.3 GB RAID 5 (Rack 1) RAID-3: 10 x 9 GB RAID 5 (Rack 1) RAID-4: 10 x 9 GB RAID 5 (Rack 1)	RAID-1: 10 x 4.3 GB RAID 5 (Rack 1) RAID-2: 10 x 4.3 GB RAID 5 (Rack 1) RAID-3: 10 x 9 GB RAID 5 (Rack 1) RAID-4: 10 x 9 GB RAID 5 (Rack 1) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)	RAID-1: 10 x 4.3 GB RAID 5 (Rack 1) RAID-2: 10 x 4.3 GB RAID 5 (Rack 1) RAID-3: 10 x 9 GB RAID 5 (Rack 1) RAID-4: 10 x 9 GB RAID 5 (Rack 1) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)
SPRHW-GSFC-2 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-2 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-2 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-2 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-2 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal
SPRHW-GSFC-3 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-3 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-3 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-3 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal	SPRHW-GSFC-3 Function: GSFC AI&T Cabinet: Desktop Console: Graphics CPU: X Terminal
SPRHW-GSFC-4 Function: LIS Cabinet: Challenge DM Console: Character CPU: 1 x 150 MHz R4400 RAM: 256 MB/1-way interleaved	SPRHW-GSFC-4 Function: LIS and COLOR Cabinet: Challenge DM Console: Character CPU: 4 x 150 MHz R4400 RAM: 512 MB/2-way interleaved	SPRHW-GSFC-4 Function: LIS and COLOR Cabinet: Challenge DM Console: Character CPU: 4 x 150 MHz R4400 RAM: 512 MB/2-way interleaved	SPRHW-GSFC-4 Function: LIS and COLOR Cabinet: Challenge DM Console: Character CPU: 4 x 150 MHz R4400 RAM: 512 MB/2-way interleaved	SPRHW-GSFC-4 Function: LIS and COLOR Cabinet: Challenge L Console: Character CPU: 6 x 150 MHz R4400 RAM: 512 MB/2-way interleaved

Table A-1. Detailed SPRHW Configurations for Release A and Release B

Release A	C 3Q97	G 3Q98	K 3Q99	O 3Q00
IO4: One VME-1: FDDI HIO-1 (1,1): SCSI HIO-2 (1,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): One 2.0 GB Internal Disk SCSI-2 (1,1,1): RAID-1 SP1 SCSI-3 (1,1,2): Unused SCSI-4 (1,1,3): Unused RAID-1: 10 x 4.3 GB RAID 5 (Rack 1)	IO4: One VME-1: FDDI HIO-1 (1,1): SCSI HIO-2 (1,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): One 2.0 GB Internal Disk, One 4.3 GB Internal Disk SCSI-2 (1,1,1): RAID-1 SP1 SCSI-3 (1,1,2): RAID-1 SP2 SCSI-4 (1,1,3): Unused RAID-1: 20 x 4.3 GB RAID 5 (Rack 1)	IO4: One VME-1: FDDI HIO-1 (1,1): SCSI HIO-2 (1,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): One 2.0 GB Internal Disk, One 4.3 GB Internal Disk SCSI-2 (1,1,1): RAID-1 SP1 SCSI-3 (1,1,2): RAID-1 SP2 SCSI-4 (1,1,3): Unused RAID-1: 20 x 4.3 GB RAID 5 (Rack 1)	IO4: One VME-1: FDDI HIO-1 (1,1): SCSI HIO-2 (1,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): One 2.0 GB Internal Disk, One 4.3 GB Internal Disk SCSI-2 (1,1,1): RAID-1 SP1 SCSI-3 (1,1,2): RAID-1 SP2 SCSI-4 (1,1,3): Unused RAID-1: 20 x 4.3 GB RAID 5 (Rack 1)	IO4: One VME-1: FDDI HIO-1 (1,1): SCSI HIO-2 (1,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): One 2.0 GB Internal Disk, One 4.3 GB Internal Disk SCSI-2 (1,1,1): RAID-1 SP1 SCSI-3 (1,1,2): RAID-1 SP2 SCSI-4 (1,1,3): Unused RAID-1: 20 x 4.3 GB RAID 5 (Rack 1)
	SPRHW-GSFC-5 Function: GSFC AI&T Cabinet: Power Challenge XL Console: Character CPU: 16 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-5 Function: MODIS (1 of 3) Cabinet: Power Challenge XL Console: Character CPU: 16 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-5 Function: MODIS (2 of 4) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-5 Function: MODIS (2 of 6) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved
	IO4: Two HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): HIPPI SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2	IO4: Two HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): HIPPI SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2	IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HIPPI HIO-6 (3,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2	IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HIPPI HIO-6 (3,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2
	RAID-1: 10 x 9 GB RAID 5 (Rack 2) RAID-2: 10 x 9 GB RAID 5 (Rack 2) RAID-3: 10 x 4.3 GB RAID 5 (Rack 2) RAID-4: 10 x 4.3 GB RAID 5 (Rack 2)	RAID-1: 10 x 9 GB RAID 5 (Rack 2) RAID-2: 10 x 9 GB RAID 5 (Rack 2) RAID-3: 10 x 4.3 GB RAID 5 (Rack 2) RAID-4: 10 x 4.3 GB RAID 5 (Rack 2)	RAID-1: 10 x 9 GB RAID 5 (Rack 2) RAID-2: 10 x 9 GB RAID 5 (Rack 2) RAID-3: 10 x 4.3 GB RAID 5 (Rack 2) RAID-4: 10 x 4.3 GB RAID 5 (Rack 2) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)	RAID-1: 10 x 9 GB RAID 5 (Rack 2) RAID-2: 10 x 9 GB RAID 5 (Rack 2) RAID-3: 10 x 4.3 GB RAID 5 (Rack 2) RAID-4: 10 x 4.3 GB RAID 5 (Rack 2) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)
	SPRHW-GSFC-6 Function: GSFC AI&T Cabinet: Power Challenge XL Console: Character CPU: 16 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-6 Function: MODIS (2 of 3) Cabinet: Power Challenge XL Console: Character CPU: 16 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-6 Function: MODIS (3 of 4) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved	SPRHW-GSFC-6 Function: MODIS (3 of 6) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved

Table A-1. Detailed SPRHW Configurations for Release A and Release B

Release A	C 3Q97	G 3Q98	K 3Q99	O 3Q00
	<p>IO4: Two HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): HiPPI</p> <p>SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2</p>	<p>IO4: Two HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): HiPPI</p> <p>SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2</p>	<p>IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused</p> <p>SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2</p>	<p>IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused</p> <p>SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2</p>
	<p>RAID-1: 10 x 9 GB RAID 5 (Rack 3) RAID-2: 10 x 9 GB RAID 5 (Rack 3) RAID-3: 10 x 4.3 GB RAID 5 (Rack 3) RAID-4: 10 x 4.3 GB RAID 5 (Rack 3)</p>	<p>RAID-1: 10 x 9 GB RAID 5 (Rack 3) RAID-2: 10 x 9 GB RAID 5 (Rack 3) RAID-3: 10 x 4.3 GB RAID 5 (Rack 3) RAID-4: 10 x 4.3 GB RAID 5 (Rack 3)</p>	<p>RAID-1: 10 x 9 GB RAID 5 (Rack 3) RAID-2: 10 x 9 GB RAID 5 (Rack 3) RAID-3: 10 x 4.3 GB RAID 5 (Rack 3) RAID-4: 10 x 4.3 GB RAID 5 (Rack 3) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)</p>	<p>RAID-1: 10 x 9 GB RAID 5 (Rack 3) RAID-2: 10 x 9 GB RAID 5 (Rack 3) RAID-3: 10 x 4.3 GB RAID 5 (Rack 3) RAID-4: 10 x 4.3 GB RAID 5 (Rack 3) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)</p>
	<p>SPRHW-GSFC-7 Function: GSFC Queuing Server</p>	<p>SPRHW-GSFC-7 Function: GSFC Queuing Server</p>	<p>SPRHW-GSFC-7 Function: GSFC Queuing Server</p>	<p>SPRHW-GSFC-7 Function: GSFC Queuing Server</p>
			<p>SPRHW-GSFC-8 Function: MODIS (4 of 4) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved</p>	<p>SPRHW-GSFC-8 Function: MODIS (4 of 6) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved</p>
			<p>IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused</p> <p>SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2</p>	<p>IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused</p> <p>SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2</p>

Table A-1. Detailed SPRHW Configurations for Release A and Release B

Release A	C 3Q97	G 3Q98	K 3Q99	O 3Q00
			SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2 RAID-1: 10 x 9 GB RAID 5 (Rack 4) RAID-2: 10 x 9 GB RAID 5 (Rack 4) RAID-3: 10 x 4.3 GB RAID 5 (Rack 4) RAID-4: 10 x 4.3 GB RAID 5 (Rack 4) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)	SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2 RAID-1: 10 x 9 GB RAID 5 (Rack 4) RAID-2: 10 x 9 GB RAID 5 (Rack 4) RAID-3: 10 x 4.3 GB RAID 5 (Rack 4) RAID-4: 10 x 4.3 GB RAID 5 (Rack 4) RAID-5: 10 x 4.3 GB RAID 5 (Rack 5)
			SPRHW-GSFC-9 Function: GSFC AI&T Cabinet: Power Challenge XL Console: Character CPU: 8 x 275 MHz R10000 RAM: 1 GB/8-way interleaved IO4: Two HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): HiPPI SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 RAID-1: 10 x 9 GB RAID 5 (Rack 6) RAID-2: 10 x 9 GB RAID 5 (Rack 6) RAID-3: 10 x 4.3 GB RAID 5 (Rack 6) RAID-4: 10 x 4.3 GB RAID 5 (Rack 6)	SPRHW-GSFC-9 Function: GSFC AI&T (1 of 2) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2 RAID-1: 10 x 9 GB RAID 5 (Rack 6) RAID-2: 10 x 9 GB RAID 5 (Rack 6) RAID-3: 10 x 4.3 GB RAID 5 (Rack 6) RAID-4: 10 x 4.3 GB RAID 5 (Rack 6) RAID-5: 10 x 4.3 GB RAID 5 (Rack 7)
				SPRHW-GSFC-10 Function: GSFC AI&T (2 of 2) Cabinet: Power Challenge XL Console: Character CPU: 24 x 275 MHz R10000 RAM: 2 GB/8-way interleaved IO4: Three HIO-1 (1,1): FDDI

Table A-1. Detailed SPRHW Configurations for Release A and Release B

Release A	C 3Q97	G 3Q98	K 3Q99	O 3Q00
				<p>HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2</p> <hr/> <p>RAID-1: 10 x 9 GB RAID 5 (Rack 8) RAID-2: 10 x 9 GB RAID 5 (Rack 8) RAID-3: 10 x 4.3 GB RAID 5 (Rack 8) RAID-4: 10 x 4.3 GB RAID 5 (Rack 8) RAID-5: 10 x 4.3 GB RAID 5 (Rack 7)</p>
				<p>SPRHW-GSFC-11 Function: MODIS (5 of 6) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved</p> <hr/> <p>IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2</p>

Table A-1. Detailed SPRHW Configurations for Release A and Release B

Release A	C 3Q97	G 3Q98	K 3Q99	O 3Q00
				RAID-1: 10 x 9 GB RAID 5 (Rack 9) RAID-2: 10 x 9 GB RAID 5 (Rack 9) RAID-3: 10 x 4.3 GB RAID 5 (Rack 9) RAID-4: 10 x 4.3 GB RAID 5 (Rack 9) RAID-5: 10 x 4.3 GB RAID 5 (Rack 7)
				SPRHW-GSFC-12 Function: MODIS (6 of 6) Cabinet: Power Challenge XL Console: Character CPU: 18 x 275 MHz R10000 RAM: 2 GB/8-way interleaved <hr/> IO4: Three HIO-1 (1,1): FDDI HIO-2 (1,2): SCSI HIO-3 (2,1): SCSI HIO-4 (2,2): SCSI HIO-5 (3,1): HiPPI HIO-6 (3,2): Unused SCSI-0 (1,0,1): CD-ROM SCSI-1 (1,0,2): Two 9 GB Internal Disks SCSI-2 (1,2,1): RAID-1 SP1 SCSI-3 (1,2,2): RAID-1 SP2 SCSI-4 (1,2,3): RAID-2 SP1 SCSI-5 (2,0,1): RAID-2 SP2 SCSI-6 (2,0,2): RAID-3 SP1 SCSI-7 (2,1,1): RAID-3 SP2 SCSI-8 (2,1,2): RAID-4 SP1 SCSI-9 (2,1,3): RAID-4 SP2 SCSI-10 (2,2,1): RAID-5 SP1 SCSI-11 (2,2,2): RAID-5 SP2 <hr/> RAID-1: 10 x 9 GB RAID 5 (Rack 10) RAID-2: 10 x 9 GB RAID 5 (Rack 10) RAID-3: 10 x 4.3 GB RAID 5 (Rack 10) RAID-4: 10 x 4.3 GB RAID 5 (Rack 10) RAID-5: 10 x 4.3 GB RAID 5 (Rack 7)

This page intentionally left blank.

Abbreviations and Acronyms

ACMHW	Access Control and Management HWCI
ACRIM	Active Cavity Radiometer Irradiance Monitor
ADC	Affiliated Data Center
ADEOS	Advanced Earth Observing System
ADS	Archive data sets
ADSHW	Advertising Service HWCI
ADSRV	Advertising Service CSCI
AHWGP	Ad hoc working group on production
AITHW	Algorithm Integration & Test HWCI
AITTL	Algorithm Integration and Test Tools (CSCI)
AM	Ante meridian
ANSI	American National Standards Institute
APC	Access/Process Coordinators
API	Application Programming Interface
APID	Application Process Identifier
AQA	Algorithm QA
AQAHW	Algorithm QA HWCI
ASAP	As soon as possible
ASCII	American Standard Code for Information Interchange
ASF	Alaska SAR Facility (DAAC)
ATM	Asynchronous Transfer Mode
BDS	Bulk Data Service
BONeS	Block-oriented Network Simulator
CCSDS	Consultative Committee for Space Data Systems
CD ROM	Compact disk read only memory
CDR	Critical Design Review
CDRL	Contract Data Requirements List
CERES	Clouds and Earth's Radiant Energy System

CI	Configuration Item
CIDM	Client, Interoperability, Data Management
CIESIN	Consortium for International Earth Science Information Network
CLS	Client Subsystem
COLOR	Ocean Color (EOS Color Mission)
CORBA	Common Object Request Broker Architecture
COTS	Commercial off-the-shelf
CPU	Central processing unit
CSC	Computer Software Component
CSCI	Computer Software Configuration Item
CM	Configuration Management
CSDT	Computer Science Data Types
CSMS	Communications and Systems Management Segment
CSS	Communication Subsystem (CSMS)
D3	Type of StorageTek digital helical scan tape
DAA	DAN Acknowledge
DAAC	Distributed Active Archive Center
DADS	Data Archive and Distribution System
DAN	Data Availability Notice
DAO	Data Assimilation Office
DAR	Data Acquisition Request
DAS	Data Assimilation System
DAS	Data Availability Schedule
DAT	Digital Audio Tape
DBA	Database administrator
DBMS	Database Management System
DCCI	Distributed Computing Configuration Item
DCE	Distributed Computing Environment
DCN	Document Change Notice
DDA	Data Delivery Acknowledgement
DDICT	Data Dictionary CSCI

DDIST	Data Distribution CSCI
DDN	Data Delivery Notice
DDSHW	Document Data Server Hardware CI
DDSRV	Document Data Server CSCI
DDTS	Distributed Defect Tracking System
DESKT	Desktop CI
DEV	Developed code
DID	Data Item Description
DIM	Distributed Information Manager
DIMGR	Distributed Information Management CSCI
DIPHW	Distribution & Ingest Peripheral Management HWCI
DMGHW	Data Management HWCI
DMS	Data Management System
DMS	Data Management Subsystem
DOF	Distributed Object Framework
DP	Data Processing
DPR	December Progress Review
DPREP	Science Data Pre-Processing CSCI
DPS	Data Processing Subsystem
DR	Data Repository
DRPHW	Data Repository HWCI
DS	Data Server
DSM	Distribution Storage Management
DSS	Data Server Subsystem
DT	Data Type
EBnet	EOSDIS BAckbone Network
ECS	EOSDIS Core System
EDC	EROS Data Center (DAAC)
EGS	EOS Ground System
EMASS	E-Systems Modular Automated Storage System
EDOS	EOS Data and Operations System

EOS	Earth Observing System
EOSDIS	Earth Observing System Data and Information System
EP	Evaluation Package
EP	Early Prototype
ERS	European Remote-Sensing Satellite
ESDIS	Earth Science Data and Information System
ESDT	Earth Science Data Types
EXB	Exabyte
F&PRS	Functional and Performance Requirements Specification
FC	Fiber Channel
FDDI	Fiber distributed data interface
FDF	Flight Dynamics Facility
FLOPS	Floating Point Operations per Second
FOO	Flight of Opportunity
FOS	Flight Operations Segment
FSMS	File and Storage Management System
Ftp	File transfer protocol
GB	Gigabyte
GDAO	GSFC Data Assimilation Office
GFLOPS	Giga (billions) Floating Point Operations per Second
GOES	Geostationary Operational Environmental Satellite
GRIB	Gridded Binary
GSFC	Goddard Space Flight Center
GTWAY	Version 0 Interoperability Gateway CSCI
GUI	Graphic user interface
GV	Ground Validation
HDF	Hierarchical Data Format
HiPPI	High Performance Parallel Interface
HMI	Human machine interface
HTML	Hypertext Markup Language
HWCI	Hardware Configuration Item

I&T	Integration and Test
I/O	Input/Output
IAS	Image Assessment System
ICD	Interface Control Document
ICLHW	Ingest Client HWCI
IDL	Interface Definition Language
IDR	Incremental Design Review
IEEE	Institute of Electrical and Electronics Engineers
IERS	International Earth Rotation Service
IGS	International Ground Station
ILM	Inventory/Logistics/Management
IP	International Partner
IR-1	Interim Release 1
IRD	Interface Requirements Document
IS	Ingest Subsystem
JERS-1	Japanese Earth Remote Sensing Satellite
JPL	Jet Propulsion Laboratories
LAN	Local Area Network
LANDSAT	Land Remote Sensing Satellite
LaRC	Langley Research Center
LIM	Local Information Manager
LIMGR	Local Information Management CSCI
LIS	Lightning Imaging Sensor
L0	Level 0
LSM	Local System Manager
MACI	Management Agent CSCI
MB	Megabyte
Mbps	Megabits per second
MBps	Megabytes per second
MCI	Management CSCI
MD	Maryland

MFLOP	Millions of Floating Point Operations per Second
MOC	Mission Operations Center
MODIS	Moderate-Resolution Imaging Spectrometer
MPP	Massively Parallel Processor
MRF	Medium Range Forecast
MSFC	Marshall Space Flight Center
MSS	Management Subsystem (CSMS)
MTBF	Mean time between failures
MTTR	Mean time to restore
NESDIS	National Environmental Satellite Data and Information Service
NMC	National Meteorological Center
NOAA	National Oceanic and Atmospheric Administration
NSIDC	National Snow and Ice Data Center (DAAC)
O/A	Orbit/Attitude
ODC	Other Data Center
ODL	Object Description Language
OODCE	Object oriented DCE
ORNL	Oak Ridge National Laboratory (DAAC)
OSM	Open Storage Manager
OTS	Off-the-shelf
PAM	Permanent Archive Manager
PCI	Periphewral Component Interface
PDPS	Planning and Data Processing System
PDR	Preliminary Design Review
PDS	Production Data Set
PDS	Production Data Specialist
PGE	Product Generation Executive
PGS	Product Generation System
PLNHW	Planning HWCI
POSIX	Portable Operating System for UNIX
PR	Precipitation Radar

PRONG	Processing CSCI
Q	Quarter
Q/A	Quality Assurance
QA	Quality Assurance
QAC	Quality and Accounting Capsule
RAID	Redundant Array of Inexpensive Disks
RAM	Random Access Memory
REL	Release
RID	Review Item Discrepancy
RMA	Reliability, Maintainability, Availability
RTF	Rich Text Format
S/C	Spacecraft
SAA	Satellite Active Archives (NOAA)
SCF	Science Computing Facility
SCSI II	Small Computer System Interface
SDF	Software Development File
SDP	Science Data Processing
SDPF	Sensor Data Processing Facility
SDPS	Science Data Processing Segment
SDPTK	SDP Toolkit CSCI
SDSRV	Science Data Server CSCI
SDST	Science Data Support Team
SFDU	Standard Format Data Unit
SMC	System Management Center
SMP	Symmetric Multi-Processor
SNMP	Simple Network Management Protocol
SPRHW	Science Processing HWCI
SSM/I	Special Sensor Microwave/Imager
STMGT	Storage Management CSCI
STK	StorageTek
TBD	To be determined

TBR	To be resolved
TDRSS	Tracking and Data Relay Satellite System
TMI	TRMM Microwave Imager
TONS	TDRSS Onboard Navigation System
TRMM	Tropical Rainfall Measuring Mission
TSDIS	TRMM Science Data and Information System
UR	Universal Reference
USNO	United States Naval Observatory
V0	Version 0
VC	Virtual Channel
VCDU-ID	Virtual Channel ID
VOB	Version Object Base
WAIS	Wide Area Information Servers
WAN	Wide Area Network
WKBCH	Workbench CI
WKSHC	Working Storage HWCI
W/S	Workstation
WORM	Write Once Read Many
WS	Working Storage
WWW	World Wide Web