

5. Release B Design Component Overview

This section provides an overview of the Release B components. Section 5.1 and 5.2 provide the decomposition of SDPS and CSMS into their Computer Software Configuration Items (CSCI) and Hardware Configuration Items (HWCI). Section 5.3 provides an overview of the changes and enhancements in the Release B design over that provided in Release A. Section 5.4 and Section 5.5 provide a brief summary of the role each CSCI and HWCI plays within the ECS architecture. Section 5.6 discusses the network topologies that interconnect the ECS components. Note that detailed mappings of the ECS L4 requirements to specific CSCIs and HWCIs are documented in the detailed 305 design volumes, 305-CD-021-002 through 305-CD-038-002.

5.1 SDPS Components

The SDPS consists of seven subsystems. Each subsystem consists of one or more Computer Software (CS) and Hardware (HW) Configuration Items (CIs), composed of a logical grouping of software or hardware components. These components consist of Commercial-Off-the-Shelf (COTS) hardware, and custom-developed and OTS software. OTS software may include COTS and/or reuse components. Reuse includes ECS component reuse, reuse of heritage code from other programs, freeware or shareware. Many of the software components are developed by combining OTS and custom-developed software, sometimes referred to as “wrappers” or “glue code”, to integrate and encapsulate the OTS software. Collectively, the CIs provide the functionality identified in the Release B SDPS/CSMS System Requirements Specification, Document number 304-CD-005-002. In addition, some functionality requires the integration of components from several subsystems, including some outside of SDPS. The SDPS subsystem and their CIs are listed below. The Release B CSCIs are described in Section 5.4, and the Release B HWCIs are described in Section 5.5.

- Client Subsystem (CLS)
 - Desktop CSCI (DESKT)
 - Workbench CSCI (WKBCH)
- Interoperability Subsystem (IOS)
 - Advertising Service CSCI (ADSRV)
 - Advertising Service HWCI (ADSHW)
- Data Management Subsystem (DMS)
 - Local Information Manager CSCI (LIMGR)
 - Distributed Information Manager CSCI (DIMGR)
 - Data Dictionary CSCI (DDICT)
 - Version 0 Interoperability Gateway CSCI (GTWAY)
 - Data Management HWCI (DMGHW)
- Data Server Subsystem (DSS)

- Science Data Server CSCI (SDSRV)
- Document Data Server CSCI (DDSRV)
- Storage Management Software CSCI (STMGT)
- Data Distribution Service CSCI (DDIST)
- Access and Control Management HWCI (ACMHW)
- Working Storage HWCI (WKSHW)
- Data Repository HWCI (DPRHW)
- Distribution and Ingest Peripheral Management HWCI (DIPHW)
- Ingest Subsystem (INS)
 - Ingest Services CSCI (INGST)
 - Ingest Client HWCI (ICLHW)
- Planning Subsystem (PLS)
 - Production Planning CSCI (PLANG)
 - Planning HWCI (PLNHW)
- Data Processing Subsystem (DPS)
 - Processing CSCI (PRONG)
 - Science Data Processing (SDP) Toolkit CSCI (SDPTK)
 - Algorithm Integration and Test CSCI (AITTL)
 - Science Processing HWCI (SPRHW)
 - Algorithm Integration and Test HWCI (AITHW)
 - Algorithm Quality Assurance (QA) HWCI (AQAHW)

The breakdown of SDPS CSCIs into planned Computer Software Components (CSCs) is documented in the specific 305 design volumes, 305-CD-021-002 through 305-CD-027-002. In addition, the planned deployment of the CSCIs and the specific CSCs at each of the individual DAACs is documented in the DAAC specific 305 volumes, 305-CD-030-002 through 305-CD-037-002.

5.2 CSMS Components

The CSMS consists of three subsystems; the Management Subsystem (MSS), the Communications Subsystem (CSS), and the Internetworking Subsystem (ISS). Each subsystem consists of one or more Configuration Items (CIs), composed of a logical grouping of software or hardware components. These components consist of Commercial-Off-the-Shelf (COTS) hardware, and custom-developed and COTS software. Many of the software components are developed by combining COTS via custom-developed software, sometimes referred to as "glue code".

The Systems Management Subsystem (MSS) common management services and management application services map to the Management CI (MCI), the Management Logistics CI (MLCI), the Management Agent CI (MACI), and the MSS Management Hardware CI (MHCI). The MHCI

includes site specific configurations of workstations and servers for management of DAACs, EOC/ ICC, for WAN management, and for system-wide coordination and monitoring purposes.

The Communications Subsystem (CSS), comprised of three service superclasses, is mapped to the Distributed Computing CI (DCCI), and the Distributed Communications Hardware CI (DCHCI). All or parts of the DCCI are installed at every ECS machine to enable distributed communications. Machines are configured as clients and/or servers to meet specific implementation requirements. The DCHCI includes communication servers for security, directory, mail, and bulletin board services (software servers may share a physical server) as required to support specific site implementations.

The Internetworking Subsystem (ISS) includes three superclasses which are mapped to the Network CI (NWCI), and the Internetworking Hardware CI (INHCI). Part or all of the NWCI is installed on every ECS machine (end system), and on communication routers (intermediate systems) based on specific site implementation requirements. The INHCI includes routers, plant cabling (e.g., copper cables and optical fiber), and modem access devices required to support specific site implementations.

Collectively, the CIs provide the functionality identified in the Release B SDPS/CSMS System Requirements Specification, 304-CD-005-001. The CSMS CIs are listed below. The Release B CSCIs are described in Section 5.4 and the Release B HWCI are described in Section 5.5.

- Communications Subsystem (CSS)
 - Distributed Computing Software CI (DCCI)
 - Distributed Communications Hardware CI (DCHCI)
- Systems Management Subsystem (MSS)
 - Management Software CI (MCI)
 - Management Agents CI (MACI)
 - Management Logistics CI (MLCI)
 - Management Hardware CI (MHCI)
- Internetworking Subsystem (ISS)
 - Internetworking CI (INCI)
 - Internetworking Hardware CI (INHCI)

The breakdown of CSMS CSCIs into planned Computer Software Components (CSCs) is documented in the specific 305 design volumes, 305-CD-028-002 and 305-CD-029-002. In addition, the planned deployment of the CSCIs and the specific CSCs at each of the individual DAACs and the SMC is documented in the DAAC and SMC specific 305 volumes, 305-CD-030-002 through 305-CD-038-002.

5.3 Summary of Changes to and Enhancements in Release B Design

The Release B design will continue to evolve from the Release A design due to various factors, including additional requirements for Release B, new COTS selections, completion of trade studies and prototypes, changes to the technical baseline, and refinement of the object models. A summary

of some of the major Release B additions or enhancements to the Release A design are summarized below in Table 5.3-1.

Table 5.3-1. Release B Changes and Enhancements to Release A Capabilities (1 of 2)

Release B Enhancement	Release A Capability	Subsystems affected
Two way interoperability with NOAA; increased access with GCMD/GCDIS.	One way interoperability with NOAA.	CLS, DMS
More robust multi-DAAC planning and scheduling; support for inter-DAAC resource conflict resolution via use of common planning data; support for Targets Of Opportunity (TOOs)	Basic inter-DAAC planning and scheduling	PLS
Continued full TRMM support, plus support for Landsat7, COLOR, AM-1, ADEOS II, SAGE III, RADAR ALT and ACRIM. Support for ERS, JERS and RADARSAT at ASF. Support for DAO at GSFC.	Complete data handling/processing of TRMM, CERES, and LIS instrument data. Interface testing for ASTER GDS to EDC DAAC; LPS to EDC DAAC; MODIS SCF to GSFC and EDC DAACs; and AM-1 MOPITT, MISR, and CERES SCFs to LaRC DAAC.	All
Greatly increased (by at least an order of magnitude) maximal network rates, data processing and required data product storage, esp. for AM-1 mission support.	Moderate data rates, data processing and data storage requirements.	All
Replacement with V1 ECS client plus continued V0 interoperability, V0 2-way interoperability	Use of Release A Version 0 client capability, one-way interoperability	CLS, DMS
Enhanced Release A WAN to support additional data traffic requirement for Release B.	Use of Release A ESN WAN for inter-DAAC communications.	ISS
Enhanced local information management capabilities and implementation of distributed information management. Enhanced management reporting.	Basic local information and resource management.	DMS
More complex data searches, including multiple data set coincident searches.	Simple data searches	CLS, DMS, DSS
Enhanced processing on-demand in addition to simple storage and retrieval from archive.	No on-demand processing.	DPS, DSS
Robotic control of file servers; enhanced attached storage capability; client APIs for scientists to gain access to data storage and retrieval services.	Basic data storage and retrieval capabilities.	DSS
Enhanced metadata capabilities - expandable metadata attributes and geographic metadata search - including WRS parameters.	Basic metadata generation and search.	DSS, DMS

**Table 5.3-1. Release B Changes and Enhancements
to Release A Capabilities (2 of 2)**

Release B Enhancement	Release A Capability	Subsystems affected
System & network management to include additional security features and DAAC interfaces (esp. the ASTER GDS and ASF support)	System & Network Management for Release A network configuration	CSS, ISS, MSS
Additional DCE services	Distributed Object Framework	CSS
SMC integrated with LSM. SMC capabilities for the security management and performance management are automated.	SMC activated. Many of the SMC capabilities are performed manually.	MSS
Completion (by end of Rel. B operations) of data migration from V0. Increased capability to translate data to HDF.	Initiation of data migration from V0 and data translation to other formats.	DSS
Mode management	No comparable Release A capability	MSS
DAO data assimilation processing	Integration of Release A standalone environment	DPS, PLS
DARs	No DARs	CLS, DPS, DSS
Ability to add Extended Data Providers	None	DMS

Descriptions of major additions or enhancements in Release B for specific subsystems are highlighted below.

CLS Subsystem

An implementation plan for reusing the Version 0 Client as the Release A Client was produced. In Release B, the Version 0 Client is replaced, although the existing V0 client will still be supported via the V0 interoperability gateway. A new group of client tools is provided in place of the V0 client. These client tools provide similar functions to the Version 0 Client while communicating in the ECS query language (ESQL) and protocol and providing more flexibility in how the system is accessed.

IOS Subsystem

The advertising service CSCI is used in total from Release A. The Advertising Service server will be modified to support the Earth Science Query Language and subscription notifications for inserts, updates, and deletions of advertisements. The search and access of advertisements is fully implemented at Release A. At Release B, the search interface will be enhanced to provide linkages to the Data Dictionary Service to obtain definitions of terms for those advertisements that also have schema entries in the Data Dictionary Service.

DMS Subsystem

There are three new CSCIs implemented as part of Release B, the Data Dictionary Service CSCI (DDICT), the Distributed Information Manager CSCI (DIMGR), and the Local Information Manager CSCI. In addition, the Release A V0 Interoperability Gateway (GTWAY) will be enhanced to provide two-way interoperability at Release B. In Release A, the GTWAY provides interoperability between V0 and ECS. At Release B, the GTWAY will be enhanced to provide interoperability between the new ECS clients and the V0 systems. The LIMGR, DIMGR, and GTWAY will also provide reusable components from which other gateways and interoperability components may be built. For example, these will be reused in a special ASTER gateway component.

Data Server Subsystem

The design activities since Release A CDR have not significantly altered the design presented, but instead have focused on the natural evolution of the system to meet the Release B requirements. The principal requirements to be considered in the Release B timeframe are the increased volumes and diversity of earth science data, and the services to be provided on that data.

The FSMS product selected for Release B is the same as for Release A, e.g., the Archival Management and Storage System (AMASS) product marketed by EMASS Inc. The file system design allows all UNIX File System (UFS) access methods to be employed (e.g., ftp, rcp, uucp, nfs, RPC, native, etc.) while removing some of the limitations of the UFS. Primary among these is reliance on UNIX Index Node (inode) structures. AMASS maintains all inode information in database files rather than in associated disk structures. This minimizes or eliminates many of the file search problems inherent in searching large numbers of files in multiple directories. In addition, AMASS organizes files as groups of blocks which can be individually retrieved. This differs from UFS resident systems which require staging the entire file. AMASS utilizes a disk-based I/O buffer for communications rate matching between disk and tape resources.

EMASS Automated Tape Library has been selected for the Data Repository component of the Data Server. The driving selection factor was the library's ability to accommodate multiple media form factors. This crucial ability enables recording technology migration with minimum migration cost associated.

Document management will continue to be implemented with a public class interface as well as through HTTP access. The additional document formats in Release B (Word, WordPerfect, Interleaf) will be supported internally, and the public interfaces to these new document formats will remain relatively unchanged from Release A. Text search functionality and the document repository component will continue to be provided by components that are mostly COTS.

The Data Server in Release B will migrate to the Illustra Object-Relational DBMS product for its earth science data catalog. Illustra provides the server and data type extensibility that the Data Server requires for the diversity of data types that will be managed in Release B. Although the detailed design of the Data Server is driven by the need to seek extensible, flexible and scalable components, the design of the Data Server software interfaces remain a set of stable, consistent classes which provide an interface to a dynamic set of data types and data type services.

Ingest Subsystem

The Ingest design will be modified to support an increase in Release B data rates but the overall Ingest architecture is unchanged from Release A. The Ingest Subsystem design provides an interface with a gateway translating external messages (tcp/ip) into OODCE requests for ECS servers.

Planning Subsystem

The basic capabilities of the Planning Subsystem have not changed from Release A. Its capabilities include the AutoSys/AutoXpert Job Scheduling COTS selection. The interfaces to the Planning and Data Processing subsystems are appropriately encapsulated to give later flexibility augmenting, modifying or replacing the underlying COTS as ECS matures.

The Job Scheduling COTS is a component of the Data Processing subsystem, which accounts for the reallocation of some of the Planning capabilities.

The Planning design within the Production Management capabilities performs two main activities:

1. coordinates the production by providing a Data Processing Request to the Data Processing subsystem when all the data required for the task are present at the Data Server, and
2. provides a display of the active production, and it's status according to the plan.

These two capabilities are performed by separate CSCs. The subscription manager CSC provides the first of these capabilities. The graphic capabilities of AutoXpert in the Data Processing subsystem provides the second.

The selection of AutoSys and AutoXpert affect the sequence of events “activating” a plan. This now involves rolling a portion of the “long term” plans generated in the Planning subsystem into AutoSys.

Data Processing Subsystem

The selected COTS products for the Data Processing Subsystem are Platinum Technology's AutoSys and AutoXpert. They are integrated into PDPS to provide the basis for the monitoring and management of ECS' science data production facility.

Design decisions continue to be driven by a desire to minimize custom code development, tempered by the need to provide proper encapsulation of the COTS to insure later flexibility of adding or modifying the underlying COTS product as ECS matures and evolves.

The following provides a top-level view of the current Planning and Data Processing Subsystem Architecture.

1. PDPS shares a common database, i.e., one instance of a Sybase RDBMS. This allows PDPS to eliminate the large amount of common persistent data structures which existed in the PDPS preliminary design. For detailed information on the PDPS Database, refer to the Release B Planning Subsystem Design Specification, 305-CD- 026-002.
2. The AutoXpert product provides mechanisms for monitoring and managing the active plan. Active plan management is performed within the Processing CSCI and the management of subscription notifications is performed in the Planning CSCI.

3. AutoSys receives all Data Processing Requests at the beginning of the day. The Data Processing Requests which do not have all data dependencies fulfilled are kept in a "HELD" state until the dependencies are fulfilled, at which time the Planning CSCI releases the job.
4. The Data Pre-Processing CSC within the Processing CSCI performs the Science Data Pre-Processing functions.

MSS Subsystem

The following are the major changes to the MSS Release A baseline:

- 1) Mode Management Service

The Mode Management Service (MMS) provides the ability to initiate/terminate a mode of execution. A mode is defined as a unique system activity, such as operations (production), testing, training, etc., where process distinction and data integrity must be maintained within the activity.

- 2) Report Generation

Extends the management- application-specific reporting offered by Release A COTS packages to include reports derived from multiple application areas.

- 3) Billing and Accounting

The Enterprise Monitoring and Coordination (EMC) Billing and Accounting Application Service (BAAS) provides the mechanisms for ECS to price user data transactions, invoice users for system usage, and meet ECS' needs to track and to provide financial data.

- 4) Ground Event Planning

Ground event planning has been moved from MSS to PLS.

CSS Subsystem

Major Release B capabilities include the use of DCE version 1.1, DCE multi-cell architecture, Remote File Access, Subscription Service, Process Framework, and Server Request Framework. These capabilities are detailed in the following sections.

DCE Version and Cell Architecture Changes. The DCE encapsulation prototype and associated benchmarking of OODCE services has provided insight into the performance of OODCE, which is a key contributor to the development of performance allocations to satisfy Level 3 requirements. DCE 1.1 has not yet been delivered by all vendors. Currently, some vendors are planning to ship DCE 1.1 in Summer, 1996.

The DCE cell configuration provides one cell per DAAC/site configuration, with an additional "Isolation cell" to isolate configurations from external networks. The "Trade-off Studies Analytical Data" (ECS DID 211/SE3) provides the basis for placement of an OODCE server at each ECS site. There will be a separate cell for EOC and SMC. Studies indicate that DCE 1.1 cells can reasonably support up to approximately 300k users. Release B is not expecting more than 25k users. The ECS architecture also includes maintaining (Directory & Security) replicas at each DAAC, the SMC and the EOC to improve performance. The multi-cell architecture to provide scalability will be implemented in Release B.

Remote File Access Enhancements. The addition of Distributed File Service (DFS) constitutes a major change in the Remote File Access facilities to be provided in Release B. DFS was not used in Release A because it implemented a single cell configuration. Furthermore, it was also felt that the number of users as well of the amount of data flow do not warrant DFS. Release B on the other hand will be implementing a multicell configuration and the number of users as well as the data flows are expected to be significantly higher. These factors have motivated the inclusion of DFS as a part of the Remote File Access facility. DFS has a rich set of functionality and provides complete security to the file level. DFS integrates well into the existing infrastructure. Performance problems of DFS that were a cause for concern in Release A are expected to be mitigated in future releases of DCE (1.1) by including significant enhancements to improve the performance. Another important change related to Remote File Access facility is the inclusion of a batch mode feature in the FTP service.

Subscription Services Capabilities. The role of the subscription server is to support the detection of previously identified events and to perform specified actions on behalf of clients who have previously registered to those events. Examples of events include science granule insertion, metadata update, new advertisement, and new schema export to DDICT. The Subscription Server Process interacts with active clients that have previously registered events with it. Some clients will not be active when an event occurs. When an event occurs, and the client registered to that event is not active, the Subscription Server sends an email message to the client.

Server Request Framework. The Server Request Framework (SRF) provides a framework for constructing ECS Servers and Client/Server APIs. It provides a single common implementation of reliable asynchronous request processing services. ECS Client applications use SRF to handle callbacks associated with asynchronous requests and notifications associated with subscriptions. The Server uses SRF to track service requests, generate callback events, accept callbacks (e.g., event notifications) and send them to the corresponding client object. This infrastructure uses a factory model for creating server objects (and client counterparts) dynamically.

Process Framework. The Process Framework (PF) provides an extensible mechanism for ECS Client and Server applications to transparently include several infrastructure features, such as interfaces to DCE-based Distributed Object Services, Life Cycle Services, and Mode Management. The primary objective of the PF is to ensure design and implementation consistency for all ECS Client and Server applications.

ISS Subsystem

The ESN WAN network (carrying DAAC-DAAC processing flows) and the Ecom network (carrying L0 data to the DAACs from EDOS) have been consolidated into a single network called the EOSDIS Backbone Network (EBnet). EBnet will be designed and implemented by GSFC Code 540. The topologies presented in this design document reflect current understanding between ECS and EBnet. Note that the EBnet consolidation impacts the detailed method by which ECS interfaces to some external systems; it does not impact internal ECS DAAC and SMC network architectures.

5.4 Computer Software Configuration Item (CSCI) Description

5.4.1 Client Subsystem

The Release B Client Subsystem contains two CSCIs, Desktop CSCI and Workbench CSCI. These CSCIs are summarized below. For more information on these CSCIs refer to the Release B SDPS Client Subsystem Design Specification (305-CD-021-002).

Desktop CSCI (DESKT)

The Desktop CSCI provides the generic underpinning for the SDPS user interfaces. All science user interfaces, and all developmental operator interfaces will be based on this CSCI. From a design perspective, it provides a set of generic object classes from which the science user interface objects will inherit their behavior, and a number of services for associating programs with icons and user actions.

Workbench CSCI (WKBCH)

The Workbench CSCI provides the user interfaces to SDPS services, as well as a number of basic tools for viewing and/or manipulating SDPS data objects (e.g., guide documents, browse images, production history and quality assurance data).

The WKBCH includes a collection of tools that provide data search and access functions. The Graphical User Interfaces (GUI) are applications built with the Motif widget set and operate under the X-Window system, allowing the user to display multiple windows simultaneously, and support a mouse for easy user interaction. The search tool also allows search areas to be specified from a global map, and provides an interactive data browse facility and coverage map of data products. The WKBCH tools will communicate with the Release B server components (such as the Science Data Server), using the Distributed Object Framework developed by the Communications Subsystem. Table 5.4.1-1 summarizes the tools available in the workbench with a brief description of their functions.

Table 5.4.1-1. Workbench Components (1 of 2)

Component Name	Description of Functions
Earth Science Search Tool	Used to submit searches to the various components in SDPS and other data providers. This includes searching Earth Science data, guide documents, advertisements, and almost any other searchable repository in ECS.
Product Request Tool	Used to request/order Earth Science data. Uses references to the data obtained from the Earth Science Search Tool or other means such as referrals from colleagues.
HyperText Viewer	Used to access services provided using World Wide Web servers. These services include the Data Dictionary, Advertising Service, User Registration, and Comment/Survey.
Visualization Tool	Used to view browse data and product data or any data formatted in ECS-HDF. This is the EOSView tool provided in Release A with enhancements for Release B. User-supplied tools can be used for more sophisticated analysis of the data.
E-mailer Tool	Used to send desktop objects via electronic mail.

Table 5.4.1-1. Workbench Components (2 of 2)

Component Name	Description of Functions
Logger/ Reviewer Tool	Used to review client sessions or transactions that have been submitted.
News Reader Tool	Used to read and submit articles to the ECS bulletin board. A user supplied news reader can be supplied as well.
HyperText Authoring Tool	Used to create HTML documents that will be incorporated in ECS.

Some of these components are off-the-shelf and some are custom developed. For more information on these tools refer to the Release B SDPS Client Subsystem Design Specification (305-CD-021-002).

No special hardware CI has been defined as host for the Workbench and Desktop CSCI. The two CSCIs will be available to scientists for installations on their workstations, but they also will be deployed on workstations within the DAAC in support of normal operations (e.g., those supporting algorithm integration and test).

5.4.2 Interoperability Subsystem

The Release B Interoperability Subsystem contains one CSCI, the Advertising Service CSCI (ADSRV), and one HWCI, the Advertising Service HWCI (ADSHW). This CSCI is summarized below. For more information on this CSCI refer to the Release B SDPS Interoperability Subsystem Design Specification (305-CD-022-002). The Release B Interoperability Subsystem reuses the Release A Advertising Service CSCI with the changes as noted in the design specification. The Advertising Service HWCI is physically located on the Data Management HWCI. The requirements of the ADSHW has been taken into account when sizing the physical workstations and servers that make up the Data Management HWCI.

Advertising Service CSCI (ADSRV)

The Advertising CSCI manages a database of information describing the services and data offered by EOSDIS service providers. The user interfaces to the CSCI are part of the Workbench CSCI. However, the CSCI supports access through Internet protocols (HTTP), and thus will be accessible to users who do not have an SDPS Client Subsystem installed, as well.

The Advertising Service provides the interfaces needed to support Client defined interactive browsing and searching of advertisements. Although there will be a single format for submitting advertisements to the service, advertisements should be accessible via several different interfaces to support database searching, text searching, and hypertext access and retrieval according to several different viewing styles (e.g., plain ASCII text, interactive form, or HTML document).

A data server or other provider will advertise its data collections and services with the Advertising Service. The advertisement will include a listing of all products (and other Earth Science Data Types) available in the collection and a set of product attributes. Advertisements include directory level metadata, therefore, the attributes reflected in the advertising service include the ECS Core Metadata Directory-Level attributes. The workbench will send user queries which access only directory level metadata directly to the advertising service (rather than sending it as a distributed

query to the various sites which provided the advertising information). A user who wishes to find out what data sets are available on the network can search (i.e., formulate a query) or browse (i.e., navigate through hypertext pages of advertisements) the advertising information. Both types of 'directory searching' are available on the user's desktop; the user can choose whichever approach is most convenient in the current work context.

5.4.3 Data Management Subsystem

The Release A Data Management Subsystem contains one CSCI, the Version 0 Interoperability Gateway CSCI. The Release B Data Management Subsystem consists of three new CSCIs, the Data Dictionary CSCI, the Distributed Information Manager CSCI, and the Local Information Manager CSCI. The Version 0 Interoperability Gateway CSCI provides additional functionality in Release B. For more information on these CSCIs refer to the Release B SDPS Data Management Subsystem Design Specification (305-CD-023-002).

Data Dictionary Service CSCI (DDICT)

This CSCI stores and provides access to descriptions of data products, their attributes, and the valid values of those attributes. Users query the DDICT to get these descriptions to enhance their knowledge of the system and what it provides. Terms can have different meanings based on the context they are used in. For example, the "Sea Surface Temperature" as a geophysical parameter can have a different meaning (units, location, etc.) dependent on which data product the parameter is used in. The DDICT provides these details to the client. The client is responsible for differentiating the different meanings to the user and making it clear which definition is being used and when.

The DDICT data is also used by the other Data Management CSCIs to decompose queries and pass them on to other components. The attribute information is mapped to data products which are in turn mapped to components. Thus, the Distributed Information Manager and Local Information Manager can determine from the contents of a query which other components should be accessed to satisfy the request. The DDICT is the manager of this schema information.

Distributed Information Manager CSCI (DIMGR)

The DIMGR provides distributed search and access services. The DIMGR accepts queries and data access requests for execution. It acts as a search agent on behalf of users by identifying the sources of the data and transforming the search and access operations into requests which are acceptable to other data sources, such as Local Information Manager or Science Data Server. Users interface with the DIMGR to determine the status of the search or to obtain the search results. The details of the underlying requests to the other agents are hidden from the user, except when the user requests these details. The DIMGR uses the information in the DDICT database in order to decompose the queries and determine the optimal request structure to satisfy the request. LIMGRs and Science Data Servers make themselves available to DIMGRs by exporting information to the DDICT which updates the DDICT database. The DIMGR scope can be reconfigured to encompass the new information that has been established in the DDICT.

Local Information Manager CSCI (LIMGR)

The LIMGR operates in a similar manner to the DIMGR except that it works in a local site environment rather than a wide area network environment. The LIMGR accepts search and data access requests and issues these to other search agents (such as Science Data Servers). Science Data Servers make themselves available to LIMGRs by exporting information to the DDICT, which updates the DDICT database. The LIMGR scope can be reconfigured to encompass the new information that has been established in the DDICT. The LIMGR can provide the sole interface to a site or the site can opt to make the Science Data Servers directly accessible as well. The LIMGR can be used by external agencies to provide access to their data while hiding the details of the site from the rest of the ECS system. More information on external data provider options is available in External Data Provider Options (442-TP-001-001).

Version 0 Interoperability Gateway CSCI (GTWAY)

The GTWAY provides a bi-directional gateway between ECS and Version 0. It enables V0 IMS users to query ECS databases, and users of the ECS Client Subsystem to query Version 0 databases. In Release A the Version 0 to ECS direction is being implemented. During Release B the reverse direction will be developed.

The V0 Gateway provides interoperability with V0 for directory queries, inventory queries, browse requests and product orders. Version 0 queries originating from the Version 0 System Client are sent to a Version 0 gateway which operates at each DAAC. The gateway translates an incoming V0 ODL request into ECS query format and submits it to the local ECS data server. The results are returned through the Gateway, which then reformats it into V0 ODL structures and returns it to the Version 0 System Client. The structure of the V0 ODL messages is documented in "Messages and Development Data Dictionary for v4.5 of IMS Client" (IMSV0-PD-SD-002 v1.0.11 950515). The Gateway uses information stored in the DDICT database to resolve mappings between ECS and V0 valid values. The Advertising Service (ADSRV) CSCI and Document Data Server (DDSRV) CSCI make use of this data as well to resolve ECS to V0 mappings.

For the ECS to V0 direction, the V0 Gateway accepts queries specified in the ECS query language and protocol and translates the queries to the V0 ODL messages. The messages are sent to the V0 IMS Servers. Inventory search, product request, browse requests, and guide searches are resolved by the V0 Gateway in the ECS to V0 direction. The responses are returned to the Gateway and are formulated as ECS query results and returned to the client. The V0 Gateway will have all the capabilities of the Distributed Information Managers and Local Information Managers, such as session management, to the level that they can be supported by the V0 protocol.

5.4.4 Data Server Subsystem

The Release B Data Server Subsystem contains four CSCIs: Science Data Server CSCI, Document Data Server CSCI, Storage Management Software CSCI and Data Distribution Service CSCI. These CSCIs are summarized below. For more information on these CSCIs refer to the Release B SDPS Data Server Subsystem Design Specification (305-CD-024-002).

Science Data Server CSCI (SDSRV)

The Science Data Server CSCI (SDSRV) manages the access to ECS science and related data, and the services on the data. Science metadata describing the data products cataloged in the ECS are

stored in databases managed by the SDSRV. The SDSRV also provides the library interfaces that are to be used by other ECS components requiring access to ECS data and their services. A subscription service is provided by the SDSRV, to allow users to register requests that are related to prescribed changes in the state of data, such as the arrival of new data, its insertion into the Data Server, and updates to its catalog description.

Document Data Server CSCI (DDSRV)

The Document Data Server CSCI (DDSRV) manages and provides access to the document holdings. It accepts, stores, indexes, and delivers documents in several different formats (e.g., PostScript and HTML), and will support popular Internet document access protocols (WAIS, http). The potential exists for merging the DDSRV with the SDSRV, with documents just becoming another type of data managed by an SDSRV. The current approach is to continue to manage two types of data and services (data and document management) separately.

Storage Management CSCI (STMGT)

The Storage Management CSCI manages and provides access to archive data. It also provides a stable interface to the other software within the data server subsystem to insulate them from future changes in storage technology of which ECS will want to take advantage.

The facilities to adapt the physical storage of data in the data server to policy, while minimizing impact to availability, is provided by the Storage Management CSCI (STMGT CSCI). This CSCI provides an isolation layer between the search and access views of the archived data in the clients domain, and the physical storage mechanisms of the data internal to the archive. Through the use of unique data identifiers, the STMGT CSCI externalizes its data holdings to the SDSRV CSCI, while hiding the actual physical storage of its data. This allows the STMGT CSCI to optimize its archive storage and data migration strategies, while maintaining a consistent reference to the data for its clients.

Data Distribution CSCI (DDIST)

Data Distribution CSCI (DDIST) is responsible for providing the distribution services to the data server. DDIST orchestrates the delivery of data to its end destination (e.g., user, DAAC). DDIST receives tasking, in the form of distribution requests, from the Science Data Server and Document Data Server CSCIs and coordinates the activities of the Storage Management CSCI in transferring the data to the media specified by the requester. DDIST also supports operator management of distribution by allowing operators to view, cancel, suspend/resume, and change the priorities of requests. Distribution of this data can be via either electronic or physical media. Electronic distribution may be requested via either push or pull. With push, DDIST uses network resources managed by Storage Management to transfer the data to a remote destination specified by the requester. For pull, the data is placed in an area managed by Storage Management, from which the request originator can retrieve the data. Physical media distribution can be via 4 mm or 8mm tape, 6250 bpi 9-track tape, FAX 3480/3490 or CD-ROM. DDIST uses resources managed by Storage Management to transfer the data to the physical media.

5.4.5 Ingest Subsystem

The Ingest Subsystem contains one CSCI, the Ingest CSCI. This CSCI is summarized below. For more information on this CSCI refer to the Release B SDPS Ingest Subsystem Design Specification (305-CD-025-002).

Ingest CSCI (INGST)

The Ingest CSCI is responsible for the receipt of data arriving at a site and the initial physical placement of data into the site's storage hierarchy. These data may be delivered through a wide variety of interfaces (network file transfer, hard media, etc.), with a wide variety of management approaches to these interfaces. This interface heterogeneity and the need to support extensibility and new data/interfaces as algorithms and provider functionality change, leads to a design in which the ingest functionality is isolated from other subsystems within the segment design.

Each instance of the Ingest CSCI has similar functionality. However, in each instance the CSCI has to deal with the characteristics of the specific interface it is managing. The Ingest CSCI implements a table-driven design to identify and invoke appropriate processing for a given interface.

5.4.6 Planning Subsystem

The Release B Planning Subsystem contains one CSCI, the Production Planning CSCI. This CSCI is summarized below. For more information on this CSCI refer to the Release B SDPS Planning Subsystem Design Specification (305-CD-026-002).

Production Planning CSCI (PLANG)

The Production Planning CSCI (PLANG) provides the ability to create, modify, and implement a production plan for a site, and manage all planning related data. A production plan is generated from basic planning information (such as data dependencies, dependencies among different production steps, descriptions of production resources and their availability schedules) with the help of a production strategy which defines, for example, the priorities for various types of processing. The plan defines a schedule for the pending production requests (i.e., instructions which specify which products should be produced for what periods of time). When completed, the plan consists of series of Data Processing Requests (DPR) which implement the production requests.

Multiple candidate plans can be created, but only one plan can be active at any one time. The CSCI submits the Data Processing Requests in the plan to the Data Processing Subsystem as data becomes available. Execution status is recorded against the plan to assess progress.

The following functionality is new in Release B:

- On-demand processing (production initiated by a user request, as opposed to routine processing);
- Inter-DAAC planning (comparing plans from multiple DAACs and identifying any data dependencies);
- Limited automatic replanning (the ability to have the software automatically notify the operator when a replan should be considered under certain user-configurable conditions);

- Data availability schedules (schedules of when a site predicts it will make data sets available, for use by other DAACs in generating local plans);
- Enhanced planning workbench using strategies (operator-defined production strategies which allow for specification of priorities and resources to help control processing);
- Support for large reprocessing jobs (breaking up very large jobs into smaller ones and reordering priorities if deemed necessary);
- DAAC-wide resource planning;
- Support for more complex production rules required by the Release B instrument teams.

5.4.7 Data Processing Subsystem

The Release B Data Processing Subsystem contains three CSCIs: Processing CSCI, SDP Toolkit CSCI and Algorithm Integration & Test CSCI. These CSCIs are summarized below. For more information on these CSCIs refer to the Release B SDPS Data Processing Subsystem Design Specification (305-CD-027-002).

Processing CSCI (PRONG)

The Processing CSCI (PRONG) is responsible for the initiation, managing, and monitoring of the generation of ECS Data Products. An ECS Data Product is generated through the execution of Product Generation Executives (PGEs) which are provided by the instrument teams. The Processing CSCI supports the execution of a PGE by performing the following activities:

- Supports Operations staff interfaces to monitor the Processing environment,
- Interfaces with the Data Server to predictively stage data required by a PGE for execution,
- Support for predictive staging of input data,
- Performs preprocessing of L0 data received from the Ingest Subsystem
- Allocates hardware resources, i.e., central processing units (CPU), memory, and disk space, required by the PGE for execution,
- Interfaces with the Data Server to destage the data generated by the execution of the PGE.

Requests for processing are transmitted to PRONG from the Planning CSCI (PLANG) in the form of Data Processing Requests (DPR) which describe the details of the processing requirements as defined in the production plan for that product.

SDP Toolkit CSCI (SDPTK)

The SDP Toolkit CSCI provides a set of software libraries which are used to integrate Science Software into the EOSDIS environment. By promoting the POSIX standard, these libraries allow the Science Data Processing environment to support the generation of data products in a heterogeneous computer hardware environment.

This CSCI is part of the incremental track design. Its design is not part of this design specification, but an overview is included to provide subsystem design context. The following documents provide guidance on the roles and responsibilities of the SDP Toolkit to support the execution of science software:

333-CD-003-002 SDP Toolkit Users Guide for the ECS Project
193-801-SD4-001 PGS Toolkit Requirements Specification for the ECS Project, FINAL,
10/93 [AKA GSFC 423-16-02]

Algorithm Integration & Test CSCI (AITTL)

The Algorithm I&T CSCI is a set of tools which are used to integrate and test new science software, new versions of science software and user methods into the Science Data Processing operational environment. The CSCI provides the software capabilities needed to transition the science processing algorithms and user methods which have been developed externally within the SCF or at a user site into the operational environment of the DAAC and to validate the results of these algorithms/methods within the operational environment. The CSCI consists for the most part, of OTS tools, including software development environments, test and integration tools (e.g., debuggers), software analysis tools, and the like. The CSCI also includes the user interfaces needed by I&T staff. New in Release B is a graphic user interface (GUI) to display a list of Science Software Archive Packages (SSAPs) and to retrieve, add, delete, or edit SSAPs or SSAP files. This CSCI also encapsulates the production rules needed to generate Data Processing Requests (DPRs) from a production request.

5.4.8 Communications Subsystem

The Release B Communications Subsystem contains one CSCI, the Distributed Computing Software CSCI. This CSCI is summarized below. For more information on this CSCI refer to the Release B CSMS Communications Subsystem Design Specification (305-CD-028-002).

Distributed Computing Software CSCI (DCCI)

DCCI is a collection of “middleware” providing additional services in each ECS release. Interim Release 1 provides ftp, virtual terminal and DCE core services: Directory, Security, Time, and Remote Procedure Calls (RPCs). Release A provides mail, bulletin board, event logger, Message Passing and object oriented DCE services along with some enhancements. Release B provides Distributed File Service(DFS), Subscription Service, Process Framework(PF), and Server Request Framework(SRF). Both IR-1 and Release A use a single DCE cell where all the users, platforms, and services are maintained. Release B will use multi-cell DCE configuration.

DCCI is distributed across all ECS components. On client and server platforms, DCCI provides SDPS and FOS applications with access to legacy services such as mail, bulletin board, file transfer and host access as well as object-oriented infrastructure services upon which to execute client-server operations. Client platforms outside the ECS installation are provided with a subset of DCCI services which are integrated within the ECS Toolkit software. In addition to installation on SDPS and FOS platforms, DCCI services are also installed on CSS and MSS servers and workstations distributed throughout the ECS.

A brief summary of the services provided by DCCI is described here.

Directory Naming Service

The Directory Naming Service provides a reliable mechanism by which distributed applications can associate information with names. Its primary purpose is to allow clients to locate servers. Its

capabilities, however, are general-purpose, and it can be used in any application that needs to make names and their attributes available throughout a network.

CSS will provide implementation of both the DNS and the X.500 by supporting BIND and OSF Global Directory Service and OSF Cell Directory Service (CDS). It also provides application programmers the ability to store, retrieve, list information in the locally supported namespaces. The DNS and X.500 namespaces are used to connect the locally supported CDS namespaces. The functionality provided here will be implemented on top of XDS/XOM interfaces. As such, application programmers can use the above mentioned services (store, retrieve, list) in CDS as well as OSF GDS.

Security Service

The security service provides secure transfer of data on local and wide area networks. It provides mechanisms to verify the identity of users, and to determine whether users are permitted to invoke certain operations (authentication and authorization). Transmission of data is protected through the use of checksums and encryption of data. Authentication is provided by trusted third party (secret key) authentication. Authorization is based on Access Control Lists. The protocol used for authentication is Kerberos. All of these features are implemented within the ECS domain by employing OSF/DCE Security Services.

Multicast

Multicasting is a mechanism through which a single copy of data is transferred from a single point to several places. Multicasting allows a sending application to specify a multicast address and send one copy of the data to that address. This data is then distributed through the Multicast backbone to all the applications listening at that address. This reduces the network traffic and improves the performance.

Message Passing Service

The Message Passing Service allows for the exchange of information between applications running on different platforms. Clients send data to servers, which process the data and return the result back to the client. This interaction can be classified into three categories: synchronous, asynchronous, and deferred synchronous.

CSS will provide two implementations of Message Passing. The first model will provide for asynchronous and synchronous message passing—byte streams only—with store and forward, recovery and persistence. It will also include the concept of groups where a list of receivers belong to a group. A message sent to the group will be delivered to all the addresses registered in that local group. The second model will provide for asynchronous and deferred synchronous communication without recovery.

Both implementations are designed to take advantage of DCE Pthreads. Message Passing Service is generally intended to handle low volumes of data per message, unlike kftp which is used for bulk data transfer.

Thread

A thread is a light weight process without the actual process overhead. Threads provide an efficient and portable way to provide asynchronous and concurrent processing, which is a requirement of

network software. Threads can maintain thread specific data and can also share data with other threads in an application. This service provides functionality to create, maintain (scheduling, locking, etc.) threads.

Time

The Time Service keeps system (host) clocks in the ECS network approximately synchronized by adjusting the time kept by the operating system at every host. This service changes the clock tick increments (rather than the actual clock) so that host clocks will be synchronized with some reference time provided by an external time provider. CSS will also provide a way to simulate time by applying a supplied delta time to the actual time. Within ECS, OSF Distributed Time Service (DTS) will be used to synchronize the system clocks. DTS makes use of the NASA Time Provider (NASA-36) as the external time provider.

LifeCycle

Managing a system involves managing individual applications. An operator may want to start a new application, shutdown/suspend a running application due to anomalies. An application may not be active all the time to accept requests. In order to effectively use the CPU and memory it is desired to control the applications as well as some objects residing in the application by starting them on demand.

LifeCycle services can be broadly classified into two categories: Application and Object level. LifeCycle services for applications involve Startup, Shutdown, Suspend and Resume functionality on applications. This functionality lets the M&O manage server applications. MSS provides the application related LifeCycle functionality. CSS provides the internal APIs that are needed for the MSS to control the applications. LifeCycle services for objects provide the application programmer with the functionality to create and delete server objects residing in different address spaces.

Process Framework

The ECS contains several infrastructure features which facilitate the implementation of client-server applications. The Process Framework provides an extensible mechanism for ECS Client and Server applications to transparently include these infrastructure features. Therefore, its importance grows with future releases of ECS. Furthermore, the framework is used solely by ECS custom developed applications and as such is not meant for COTS applications. The primary objective of the PF is to ensure design and implementation consistency for all ECS Client and Server applications. This is achieved by encapsulating the implementation details of ECS infrastructure services and removing the need for programmers to rewrite common initialization code.

In general, the following capabilities are needed in the ECS client and server applications and have been accommodated in an appropriate fashion:

- (1) Ability to initialize the process application and infrastructure in a consistent way and provide some basic process information
- (2) Interface to Mode Management
- (3) Interface to Error-Event logging

- (4) Ability to set DCE Directory/Naming Service options
- (5) Ability to set DCE Security management parameters
- (6) Support for Life Cycle services
- (7) Interface to Asynchronous Message Passing
- (8) Interface to Server Request Framework, and
- (9) Interface to the Batch FTP service

Subscription Server

The role of the subscription server is to support the detection of previously identified events and to perform specified actions on behalf of clients who have previously registered to those events. Examples of events include science granule insertion, metadata update, new advertisement, and new schema export to DDICT. The Subscription Server Process interacts with active clients that have previously registered events with it. Some clients will not be active when an event occurs. When an event occurs, and the client registered to that event is not active, the Subscription Server sends an email message to the client. The Subscription Server will be reviewed in more detail in Section 6.

Server Request Framework

The Server Request Framework provides a framework for constructing ECS Servers and Client/Server APIs. It provides a single common implementation of asynchronous request processing services. ECS Client applications use SRF to handle callbacks associated with asynchronous requests and notifications associated with subscriptions. The Server uses SRF to track service requests, accept callbacks (e.g., event notifications) and sends them to the corresponding client object. This infrastructure provides factories for creating server objects (and client counterparts) dynamically. The Server Request Framework will be reviewed in more detail in Section 6.

Distributed Object Framework (DOF)

In an object oriented processing architecture, objects may be distributed in multiple address spaces, spanning heterogeneous platforms. The basic contract between an object and its users is the interface that the object provides and users can use. Objects can be spread across the network for reasons of efficiency, availability of data, etc. From the perspective of the requester of a service, invocation should be the same no matter where the object physically resides.

The distributed object framework will be implemented using OODCE. The set of core DCE services are naming, security, threads, time, rpc and DFS. In order to aid the application programmer, another layer(object-based) of abstractions is provided with OODCE. Four generic classes: DCEObj, DCEInterface, DCEInterfaceMgr, and ESO will be available for application programmers to implement client-server applications. The DOF will be reviewed in more detail in Section 6.

Electronic Mail (E-Mail)

E-mail is a standard component of Internet systems. It is useful for asynchronous, relatively slow notification of many different types. Also, E-mail is persistent, and will continue to try to deliver

even if there are temporary network outages. The CsEmMailRelA class provides object-oriented application program interface (API) to create and send e-mail messages under program control.

Remote File Access

The file access service provides functionality for file transfers and management. Remote File Access (RFA) refers to the ability to mount remote files and access them just like local files. RFA contains the FTP service and DFS.

FTP is an internet standard application for file transfers. It allows a user to retrieve or send files from/to a remote server. The files transferred can be either ASCII or binary files. FTP also provides an insecure password protection scheme for authentication. KFTP builds on the standard FTP but adds a layer for strong Kerberos authentication. The CsFtFTPRelA object provides an object for managing FTP sessions between clients and servers to allow programmers to transfer files between machines.

DFS (Distributed File System) is used for the secure data sharing and transparent file access. It is used for the worldwide file sharing service encapsulation. It provides location independent file access with high performance, availability and security. With the remote file access (RFA) capabilities offered by DFS, users can access remote file as if they are on the local file system. Under the operating system shell prompt, DFS services can be invoked using local commands. The whole process is transparent - upon authentication at login with the file servers up, no difference is noticed between a DFS file and files on local disk.

Bulletin Board

Bulletin Board is another internet standard application, however unlike e-mail, Bulletin Board messages are directed to all readers of a named group. It uses the Network News Transfer protocol (NNTP) for sending and receiving messages. The CsBBMailRelA object provides object-oriented application program interface to send e-mail.

Virtual Terminal

Virtual Terminal access refers to remote logon to a machine. This software is provided on all platforms. The server telnetd will listen to incoming telnet clients and will allow remote logons. There is also a secure version of telnet and telnetd using Kerberos authentication which CSS will provide where available. The Telnet service is allowed only within ECS due to security considerations.

X is a Graphic User Interface conforming to the X/Open standard. While X is not a specific CSS Release B service this description is listed here for informational purposes. It consists of a client and a server where the client displays the actual interface. Developing applications in X is cumbersome and complex. OSF Motif is another standard, layered on top of X which provides a high level application programming interface to make the application development easier. Applications developed with Motif will work with an X server. The X client/server connection presents some significant security risks; therefore ECS will not support applications where the X client and the X server reside on different platforms outside a DAAC. Users can download data from ECS and can use the X application to view the data on their local machines. Alternatively, dedicated or secured circuit access from a user client is used to connect to an ECS X application.

Event Log

Event log provides the programmers the capability to record events in to files. Events are broadly classified into two categories: management events and application events. Each event is recorded with all the relevant information for identifying and for later processing. Management events need to be recorded in a history file and on some occasions reported to the Network Node Manager. Application events are only recorded into a programmer specified file. Event log provides a uniform way for the application programmers to generate and report (record) events.

5.4.9 Systems Management Subsystem

The Release B Systems Management Subsystem contains three CSCIs: Management Software CSCI, Management Agents CSCI and Management Logistics CSCI. These CSCIs are summarized below. For more information on these CSCIs refer to the Release B CSMS Systems Management Subsystem Design Specification (305-CD-029-002).

Management Software CSCI (MCI)

The Management Software CSCI provides a variety of management services to support the Systems Management Architecture presented in Section 6. A brief summary of the services provided by MCI is described here.

Mode Management

Mode Management, from a system perspective, consists of procedural activities and infrastructure control. The procedural aspects of mode management address mode planning, resource allocation, and system configuration activities. Infrastructure control ensures the software subsystems will recognize the different modes of execution and that data integrity and process distinction will be maintained within each mode. Some examples of procedural activities include: allocating the required hardware and software resources, configuring the directory namespace, configuring data storage entities, loading mode specific test drivers and configuration files, etc. The infrastructure control will be maintained through the use of a mode identifier. This identifier will automatically enable managed applications to communicate and perform data I/O within a specific mode. The MSS portion of mode management, the Mode Management Service (MMS), provides mode initiation, monitoring, and controlling capabilities. These capabilities, including the Release B MMS design approach, is presented in Section 6.7

Billing and Accounting

ECS operations are supported by integrated and automated billing and accounting functions. The Enterprise Monitoring and Coordination (EMC) Billing and Accounting Application Service (BAAS) provides the mechanisms for ECS to price user data transactions, invoice users for system usage, and meet ECS' needs to track and to provide financial data. One of the BAAS' primary function is to provide bill-back capabilities. The DPS and DSS will provide the BAAS with accounting and resource data for science user orders which have been fulfilled so that the data may be priced. For purposes of estimating the price of a new product request, pricing algorithms maintained in pricing tables in BAAS will be made available to the DSS. The invoicing functionality allows ECS to inform accounts of their activity during a particular billing cycle and of the charges associated with such activity.

Report Generation

The Report Generation Service is new with Release B. It extends the management- application-specific reporting offered by Release A COTS packages to include reports derived from multiple application areas. Access to reports by end users is simplified through a web-based interface. The need for custom ad-hoc reports and queries is implemented through the RDBMS COTS reporting/query tool and management application COTS tools.

The primary user of the Report Generation Service is the ECS/DAAC manager or member of the M&O staff who is responsible for analyzing trends in workload, capacity utilization, system performance, security, reliability, and user satisfaction. The Report Generation Service provides for the generation of a range of standard management reports. A standard report, also referred to as a canned report, is one for which a template specifying format and content has been previously defined and saved. These standard reports are maintained by M&O database specialists. Standard reports can be run automatically on a periodic basis (e.g., daily, monthly, quarterly) based on setup parameters associated with the report. ECS management and M&O personnel can access these reports for viewing from their desktop through the Web-based user interface. Optionally, they can apply time and domain scope to the standard report templates to generate ad-hoc reports. Data underlying reports can be saved in text format for import into analysis tools such as a spreadsheet.

In support of the data specialist on the M&O staff, the Report Generation Service provides a workbench for use in maintaining existing report templates and constructing new ones. The workbench also supports generation of ad-hoc queries. This workbench is implemented with the report-writer/ query tool COTS associated with the Management RDBMS. The interface to the workbench is through the COTS RDBMS client interface to the MSS server.

Report Generation Services are available to both SMC and LSM M&O personnel. In general, the default scope of reports at the SMC include all of ECS whereas the scope at an LSM is the local management domain.

Fault Management

Fault Management addresses the detection, diagnosis, isolation and resolution of faults associated with the managed objects within ECS. The managed objects comprise networks, hosts and applications. A fault is an unacceptable change in the state of a managed object. Fault Management provides for the detection of changes in state of managed objects in order to be able to distinguish the unacceptable changes that constitute faults from acceptable changes. Fault Management, therefore, provides the capabilities for real-time configuration management to include the startup, shutdown and discovery of ECS applications. Further, since the service maintains the status of resources, it provides the capability to provide the status of these resources, such as processors and associated disks, upon requests from subsystems such as the Planning Subsystem. The Process Framework provides the context for the usage of the mechanisms by developers. For more information on the Process Framework, refer to the Rel B CSMS Communications Design Specification, 305-CD-028-002.

Performance Management

The Performance Management Application Service provides the capability to continuously gather statistical and historical data on the operational states of applications, operating system resources and network components, to analyze the data collected by comparing with established criteria, adjust measurement criteria or initiate other corrective actions as necessary in order to ensure an optimal utilization of resources. The service allows for the benchmarking and trends analysis of network component performance, in addition to collecting performance data on scientific algorithms. The Performance Management Application Service has two instances: one at each of the DAACs and one at the SMC. The site Performance Management Application Service collects and processes performance data local to the site.

Site performance management data is periodically summarized and sent to the SMC for analysis by the SMC Performance Management Application. The SMC Performance Management Application Service, which has capabilities similar to those of the site Performance Application Services, operates on performance data collected system-wide by the various site Performance Management Application Services in order to evaluate system-level performance and system-wide trends. In addition, the SMC Performance Management Application Service is also capable of connecting directly to each of the DAACs as required to monitor the performance of site elements.

Security Management

The mechanisms used to provide security in ECS comprise three distinct parts: network security, distributed communications security, and host-based security. The network security is based on router address filtering. The distributed communications security addresses communications between software entities such as clients and servers employing mechanisms such as Kerberos/DCE for real-time authentication exchange. The host-based security addresses the compliance to established directives (e.g. password usage guidelines) and intrusion detection (e.g. viruses and break-ins).

The Security Management Application Service provides for the management of these three mechanisms that are used to protect and control access to ECS resources. It implements security rules and authentication procedures, maintenance of authorization facilities, maintenance of security logs, intrusion detection procedures. Network security management involves the management of routing tables used for address-based filtering (network authorization). Distributed communications security involves the management of the authentication database (the DCE registry database), the authorization database (DCE Access Control List Managers). Host-based security management addresses the protection of these mechanisms, in addition to the management of compliance to established security policy, and intrusion detection.

Accountability Management

The Accountability Management Service provides the capabilities of User Registration and the generation of reports from audit trails.

ECS provides for two generic classes of users: guest users and registered users. Guest users are users that have not formally registered to become registered users. Registered users are those guest users that have submitted requests for a registered user account, and have had accounts created for them, based on an approval process. Registered users are allowed access to services and products beyond those available to guest users.

User registration provides the operators the capability to create accounts against requests submitted by guest users wishing to become authorized ECS users. The registration service provides the capabilities for the creation, modification and maintenance of accounts with user profiles. The user profile information contains user identification, user class, field(s) of research, investigating group affiliation (if any), and shipping address, electronic mail address (if any). The Accountability Management Service makes the user profile available to the various subsystems, such as the Data Server subsystem for information such as the user's electronic mail address and the shipping address, used for the distribution of data products ordered.

The Audit Trail capability provides the means to verify the integrity of the system. This comprises the generation of a user audit trail and a security audit trail with data collected from a variety of sources.

The Accountability function provides for end-to-end tracking of user orders and requests. It also allows ECS to gather and track information on science user data orders, and to cost these orders based on different costed resources (e.g., disk utilization, CPU, media, connect time) or standard products ordered using pricing algorithms associated with each one.

Physical Configuration Management

The Physical Configuration Management Service (PCMS) provides the capability to track, manage, and control all the physical elements in the network. It integrates graphics with data to create a complete electronic model of the physical infrastructure of the network. It provides tools to locate physical proximity of down nodes, place newly discovered nodes, and manage circuit changes. It supports a variety of network administration applications including inventory, billing, and troubleshooting. It has mechanisms to track everything from maintenance data, network protocol data to software registration. In addition, it provides integration support to several Trouble Ticket applications

Request Tracking

The Request Tracking Key Mechanism is intended for use by the developers of ECS applications to report request status changes back to a central database to be displayed to an operator. The mechanism is also used to report resource cost that was collected during the life of the request. The cost data is reported back and used for cost accounting. The types of requests which are tracked are as follows: Product Orders, Ingest Requests, User Requests, and Operator Requests. The mechanism also accounts for spawning of requests. When an ECS applications creates one or more sub-requests for a request, a parent-child type of relationship is established in the mechanism so that the spawned requests can be tracked independently of each other and the operator will be able to get status information for the entire tree of requests.

Trouble Ticketing

The Trouble Ticketing Service (TTS) provides the DAACs a consistent means of reporting, classifying and tracking problem occurrence and resolution.

TTS provides several methods for a user to report a trouble ticket. The primary method is through the MsTtUserInterface class. This set of HTML documents allow a graphical form for on which the user can disseminate as much information as possible to characterize the nature of the problem. Additionally, TTS provides a textual electronic mail template for generation of trouble tickets from any e-mail package. Finally, if information regarding a problem is received any other manner, a support staff member may enter a trouble ticket directly into TTS.

To the support staff responsible for handling the trouble tickets, provides a common means of statusing, prioritizing, and categorizing reported problems and the resolutions. However, within this common environment, TTS does offer the flexibility for the support staff at an individual DAAC to customize several aspects of the trouble ticketing process. The two primary functions which provide this flexibility are the definition of escalation rules and active links.

Escalation rules are simply time activated events which execute on trouble tickets which meet a set of specified criteria. Actions which can be taken include notification (of either a user or support staff member), writing to a logfile, setting a field value on the trouble ticket, or even running a custom process. Qualifications can be expressed on any data which tracks for trouble tickets. Additionally, TTS put some pre-defined active links and escalation rules into effect at installation time so as to provide a degree of consistency across the DAACs.

TTS also provides a graphical interface to search the trouble ticket database and produce a variety of reports on the trouble ticket data.

Finally, in addition to the detailed trouble ticket information at each DAAC, TTS summarizes data at the SMC. This data allows reports to be produced which can be used to track more "global" trends in reported problems. An example of this could be a report indicating the number of trouble tickets entered per day across all DAACs.

Management Data Access

The Management Data Access (MDA) Service is responsible for centralizing, processing and providing access to the information which is logged into the ECS Management Data log file on each managed host. This log data includes performance, security audit trail, fault, and ECS application processing information.

MDA will centralize the log file data at each DAAC. It is responsible for transferring these log files to the MSS server from each managed host on a scheduled basis. This schedule is configurable, allowing time intervals, absolute time specification, and or size threshold parameters to be set for each ECS managed host's log file. Any of these parameters may trigger a transfer of the file to the MSS server. The parameters may be updated through MDA's graphical user interface.

For the purposes of shorter term analysis of event data, MDA allows browse access to detailed log file data, through its user interface. Provided a host, time period, and optional selection filter information, MDA will retrieve the requested data and display it using its log file browser. Once displayed, options are given for sorting, additional filtering, and saving this data. It should be noted

that while this access is typically used to access DAAC (or SMC/EOC) local log files, it also provides the capability to browse the log file data located at other sites.

For longer term analysis and reporting of management event data, MDA will process and accumulate metric data and load it to the Management RDBMS. In addition to this metric data, some specific event detail information will be loaded to the RDBMS.

Common Management

As documented in 193-00632, DCE Migration Study for the ECS Project, and 193-00156, DME Migration Study for the ECS Project, the Open Software Foundation's (OSF) Distributed Management Environment (DME) is the selected distributed management architecture for ECS project. DME is an open architecture that is capable of evolving with new technologies and offers an integrated Distributed Enterprise System and Network Management Architecture for ECS.

Even though a full DME compliant implementation will not exist for Release B, most industry enterprise management players are adopting these technologies and migrating their existing products toward the DME architecture.

To mitigate risk, a DME precursor product (HP OpenView) is the ECS Management Framework for Release B. This selection provides an ECS migration path to management applications under the full DME architecture.

Ground Events Planning

Ground Events Planning Service consists of the Production Planning Workbench that is part of the Production Planning applications provided by the Planning and Data Processing subsystem (PDPS). This service provides an interface to submit operations ground events such as maintenance, testing, training, etc., and to develop resource utilization plans and schedules based on approved system configurations and priorities. PDPS will provide this service in Release B.

Management Agents CSCI (MACI)

The MSS provides ECS M&O Staff with the capability to manage the ECS enterprise, i.e., to perform network and system management services on all ECS resources, including all SDPS, FOS, and CSMS components.

The enterprise management system is based on the manager-agent model. It consists of management applications, a managed object model, and a management protocol. The management applications reside on managing system(s). They provide the interfaces for the human enterprise manager to perform management tasks. The managed object model consists of managed objects which are defined to represent the resources being managed. The underlying resources can be physical devices, system software or applications. The management agent is the implementation which substantiates the managed objects. It normally resides on each remote host performing monitoring and control functions for the management applications which are on the managing system(s). The management applications communicate with agents through the management protocol.

The MSS is composed of a variety of management applications providing services such as fault, performance, security, and accountability management for ECS networks, hosts, as well as SDPS and FOS applications. The management applications reside on MSS Server. The management

information of remote objects need to be conveyed to the management applications through the Management Agent Service which primarily resides on remote hosts.

The MSS Management Agent Service provides the following functions:

- Enables the management applications to retrieve and to set managed object values.
- Performs local polling on remote hosts to monitor the state of managed resources.
- Handles event logging and notifications.
- Provides instrumentation API to application developers to enable the manageability of ECS applications.
- Defines the managed object model to represent the management characteristics of ECS applications.

SNMP has been chosen as the management protocol since it is the defacto and Internet standard protocol for network management in TCP/IP environment. The MSS management applications pass SNMP requests to the agent to retrieve management information. For setting management information, it uses DCE RPC to send requests to remote agents for security reasons.

MSS management applications need to monitor the state of managed resources. It can be done by polling of remote resources. But, remote polling has certain impact to the network traffic. Therefore, the agent can perform local polling for the management applications to avoid the costly remote polling.

Event handling will be provided by Management Agent Service to satisfy the need to dispatch events for orderly and prompt resolution to fix problems. All events will be logged locally on each host. Performance data will also be logged.

A set of Instrumentation API will be provided to ECS application developers to use for the manageability of ECS applications. ECS applications can be categorized into two general types, distributed-object-based or non-distributed-object-based applications. Application developers can determine the performance metrics along with their threshold to monitor. Fault types can also be monitored and counted. For managing non-SNMP resources such as COTS, proxy agents will have to be used, and supplied by the resource provider. The front-end of the proxy agent uses the instrumentation API provided by MSS. The back-end of the proxy and counted. For managing non-SNMP resources such as COTS, proxy agents will have to be used, and supplied by the resource provider. The frond counted. For managing non-SNMP resources such as COTS, proxy agents will have to be used, and supplied by the resource provider. The frontend is the interface unique to each resource.

MSS requires that on each managed host, standard SNMP MIB II, Host Resource MIB, and the MIBs of network devices are supported by vendor agents. In addition, a managed object model is defined by MSS for ECS applications in SNMP MIB format. The Management Agent Service implements this application MIB. The information contained in the MIB is composed of different types of attributes: configuration, performance, fault, dynamic, static, and traps.

Management Logistics CSCI (MLCI)

The Management Logistics Configuration Item (MLCI) implements the Configuration Management Services [Baseline Manager, Software Change Manager, Change Request Manager Services, Software License Manager, Software Distribution Manager Services, Inventory Manager, Maintenance Manager, Logistics Manager Service, Training Manager and Policies and Procedures Manager]. It provides tools with which ECS staffs at the DAACs, EOC, and SMC track deployed ECS baselines and control changes to the hardware and software that comprise them; distributes software and controls commercial off-the-shelf (COTS) products' licenses; tracks the inventory of all key equipment end items down to the component or other appropriate level; tracks and records use of logistics supplies for the MSS; tracks maintenance of SMC/LMS accountable equipment; tracks training and certification of SMC/LMS personnel; and maintains and distributes information on system (ECS) prescribed policies and procedures.

CMS maintains electronic stores of baseline data, software, and system change requests that enter the operational environment, making them and a variety of reports available for system maintenance and operations activities. It accepts ECS and algorithm software and non-real time configuration management data from formatted files or via operator interface. M&O staffs, sustaining engineers, and AIT teams rely on CMS data stores to make, track and audit configuration changes and to help enforce ECS CM rules. They also use CMS to produce formatted files containing change requests, site baseline records, software, documentation, and reports that can be made available for distribution system-wide via CSS services such as e-mail, ftp, and the ECS bulletin board.

For this release, the MLCI design includes three service managers from Release A and five new service managers:

- Baseline Manager
- Software Change Manager
- Change Request Manager
- Software Distribution Manager
- Software License Manager
- Inventory/Logistics/Maintenance Manager
- Training Manager
- Policies and Procedures Manager.

5.4.10 Internetworking Subsystem

The Release B Internetworking Subsystem contains one CSCI, the Internetworking CSCI. This CSCI is summarized below. The Internetworking Subsystem detailed design is not presented in its own subdocument. Instead it is presented in the DAAC-specific subdocuments of DID 305. For more information on this CSCI refer to the DAAC and SMC specific subdocuments listed below.

305-CD-030-002	Release B GSFC DAAC Implementation
305-CD-031-002	Release B LaRC DAAC Implementation
305-CD-033-002	Release B EDC DAAC Implementation

305-CD-034-002	Release B ASF DAAC Implementation
305-CD-035-002	Release B NSIDC DAAC Implementation
305-CD-036-002	Release B JPL DAAC Implementation
305-CD-037-002	Release B ORNL DAAC Implementation
305-CD-038-002	System Monitoring and Coordination Center Implementation

Internetworking CSCI (INCI)

INCI provides internetworking services based on protocols and standards corresponding to the lower four layers of the OSI reference model as described below.

Transport Protocols

ECS provides IP-based connection-oriented and connectionless transport services. The connection-oriented service is implemented using TCP, while UDP is used for connectionless transport. Higher layer applications use one or the other based on such requirements as performance and reliability.

Transmission Control Protocol (TCP), specified in RFC 793, is a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols which support multi-network applications. It provides for reliable inter-process communication between pairs of processes in host computers attached to networks within and outside ECS. Because TCP assumes it may obtain potentially unreliable datagram service from the lower level protocols, it involves additional overhead due to the implementation of retransmission and acknowledgment processes.

The User Datagram Protocol (UDP), specified in RFC 768, provides a procedure for application programs to send messages to other programs with minimal overhead. The protocol is transaction oriented and delivery of data is not guaranteed, since there is no acknowledgment process or retransmission mechanism. Therefore, applications requiring ordered and reliable delivery of data would use TCP.

Network Layer Protocols

The network layer provides the functional and procedural means to transparently exchange network data units between transport entities over network connections, both for connection-mode and connectionless-mode communications. It relieves the transport layer from concern of all routing and relay operations associated with network connections.

The Internet protocol (IP), specified in RFC 791, is the ECS-supported network protocol, based on its dominance in industry usage and wide community support. As part of IP support, ICMP and ARP will also be supported. As the IETF-specified new generation of IP becomes available for deployment, it will be supported by ECS networks.

Physical/Datalink Protocols

Physical and datalink protocols describe the procedural and functional means of accessing a particular network topology. For the Release A DAAC and SMC networks, the datalink/physical protocols to be implemented are FDDI and Ethernet. (FDDI is a 100Mbps token-passing network topology, and Ethernet is a 10 Mbps bus topology.)

5.5 Release B Hardware Architecture

5.5.1 Hardware Architecture Overview

SDPS and CSMS subsystems have been subdivided into 14 hardware configuration items (HWCI). This section provides a brief description of the hardware configuration items and the roles that each plays within overall ECS context. Sizing and selection rationale for the HWCI can be found in the respective subsystem specific subdocument. Specific configurations and candidate hardware selections can be found in the DAAC-specific and SMC subdocuments.

Figure 5.5-1, ECS DAAC Release B Hardware Architecture, further details the relationships between the HWCI and also illustrates their association with external systems and organizations in the Release B timeframe.

Section 5.5.1 describes each HWCI and provides a brief summary of the role each plays within the overall hardware architecture. Section 5.5.2 overviews the performance analysis process used to size the hardware components for Release B. Section 5.5.3 provides the process and rationale for selection of components classes of hardware for Release B. Refer to the DAAC and SMC specific subdocuments for the COTS hardware components for each site configuration (305-CD-030-002 through 305-CD-038-002).

5.5.1 Hardware Component Descriptions

5.5.1.1 Interoperability Subsystem

The Release B Interoperability Subsystem shares hardware resources with the Data Management hardware CI (DMGHW). For more information on this HWCI refer to the Release B SDPS Interoperability Subsystem Design Specification (305-CD-022-002).

5.5.1.2 Data Management Subsystem

The Release B Data Management Subsystem contains one HWCI, the Data Management Server HWCI. This HWCI is summarized below. For more information on this HWCI refer to the Release B SDPS Data Management Subsystem Design Specification (305-CD-023-002).

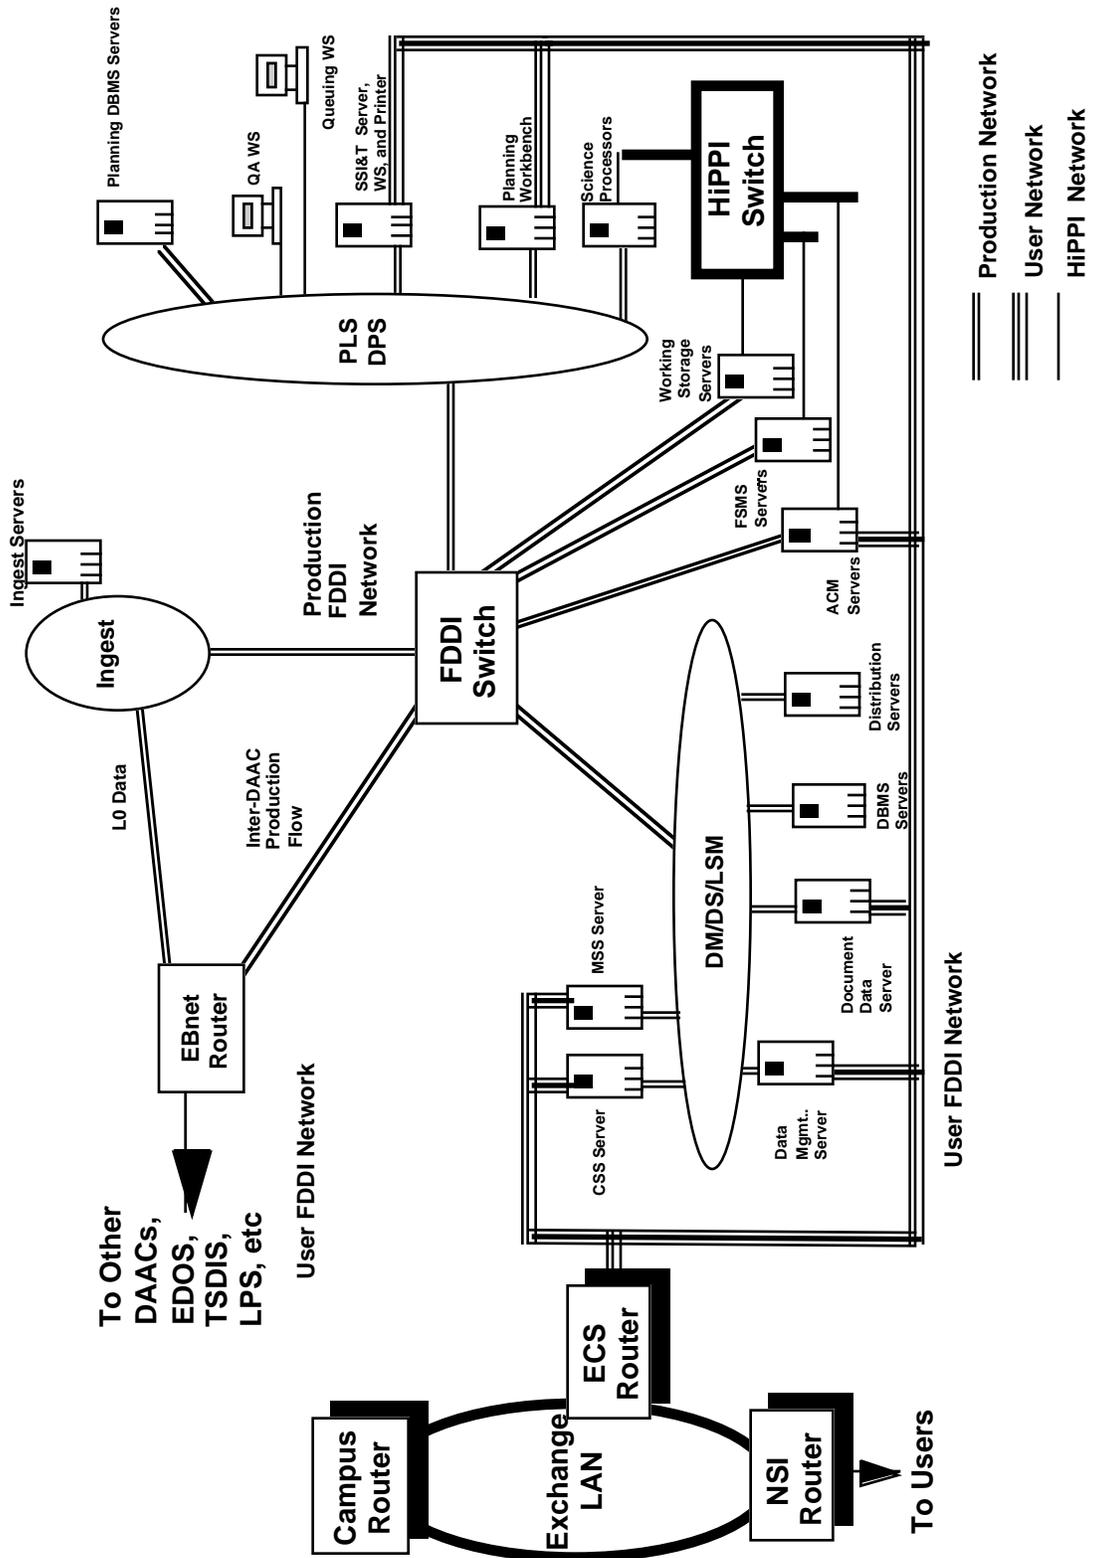


Figure 5.5.1-1. ECS DAAC Release B Hardware Architecture

Data Management Server HWCI (DMGHW)

This includes the hardware associated with the LIMGR, DIMGR, GTWAY, and DDICT CSCI. The DMGHW includes the servers and operator positions required by the four Data Management CIs including the following:

- The physical server, disk, channel, (etc.) hardware needed to process the service requests and administrative functions associated with these CSCIs, and to store the administrative and temporary data required for their operation.
- The workstations, X-Terminals, (etc.) needed to support the operator interfaces to these CSCI at each site. This includes hardware for: DBMS administration, data specialists, user support, phone/mail support, etc.

The HWCI does not include the operations position hardware associated with the data servers at each site. These hardware requirements are covered by a separate HWCI.

5.5.1.3 Data Server Subsystem

The Release B Data Server Subsystem contains five CSCIs: Access Control and Management HWCI, Working Storage HWCI, Document Data Server HWCI, Data Repository HWCI, and the Distribution and Ingest Peripherals HWCI. These HWCI are summarized below. For more information on these HWCI refer to the Release B SDPS Data Server Subsystem Design Specification (305-CD-024-002).

Access Control & Management HWCI (ACMHW)

The Access Control & Management HWCI 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 supports many of the infrastructure requirements of the Data Server. This hardware configuration item controls logical data server access, maintains client sessions, and directs service requests to other appropriate Data Server subsystem configuration items. The Access Control and Management hardware is broken down into two components; Administration Stations (AS) and Access/Process Coordinators (APCs). The number, type, and configuration of the APCs and Administration stations vary according to site needs and number of data servers supported. DAAC-specific configurations can be found in the respective DAAC-specific volume of this document set.

Working Storage HWCI (WKSHW)

The Working Storage (WS) hardware configuration item of the data server supplies storage used for temporary file and buffer storage within the Data Server architecture. In Release B WS may also be used at some sites to support the higher levels of a hierarchical storage scheme that utilize other data repositories as lower levels in the storage schema. Any data that resides in WS and is not designated as temporary data will be copied to a permanent data repository (see DRPHW - Data Repository HWCI) and maintained there. WS will hold all interim products, that are not permanently archived, until they can be deleted.

WS provides the disk, and at most DAAC sites, tape staging capacity for data acquires and inserts. Because of its role at the higher levels of the archiving hierarchy, WS may hold production related data that is to be accessed in the near future to increase performance.

Document Data Server HWCI (DDSHW)

The Document Data Server (DDSRV) provides storage and retrieval services on earth science-related documents and their metadata. The DDSRV HWCI will host the applications which provide the document access and retrieval services, which include HTTP access. Full text and keyword searching is provided, as well as the support for HyperText presentation of document metadata.

Data Repository HWCI (DRPHW)

This HWCI provides the permanent storage devices associated with the Data Server Subsystem (and some forms of Ingest Data Servers like the L0 Ingest Client). This includes archive robotics, drives, database repositories (with embedded database software), and file servers. Disk resources used for staging data after retrieval until they are processed or distributed, or after ingest until they are archived, are provided by WS HWCI.

Distribution and Ingest Peripherals HWCI (DIPHW)

This HWCI provides the pool of peripherals needed for hard media data distribution and data ingest (the HWCI is shared by the Data Server and Ingest Subsystems). This makes it possible to share peripherals which are capable of both input and output across the two subsystems, thus providing for more cost effective utilization (as requirements permit). The HWCI includes disk, tape and other media ingest and/or preparation devices (e.g., 8mm tape, CD-ROM, printers) as needed to fulfill requirements of the site. The HWCI also covers the workstations needed by ingest and distribution operators.

5.5.1.4 Ingest Subsystem

The Release B Ingest Subsystem contains one HWCI, the Ingest Client HWCI. This HWCI is summarized below. For more information on this CSCI refer to the Release B SDPS Ingest Subsystem Design Specification (305-CD-025-002).

Note that the Ingest Subsystem includes instantiations of the Data Server Subsystem HWCI needed for archiving and staging Level 0 Data. The Ingest subsystem also shares Data Server Subsystem input / output peripherals contained in the Distribution & Ingest Peripherals HWCI.

Ingest Client HWCI (ICLHW)

This HWCI includes any servers and/or workstations required for Ingest management, control, monitoring and/or processing. It includes any X-Terminals and/or workstations associated with ingest technician operator positions.

5.5.1.5 Planning Subsystem

The Release B Planning Subsystem contains one HWCI, the Planning HWCI. This HWCI is summarized below. For more information on the planning subsystem CSCIs refer to the Release B SDPS Planning Subsystem Design Specification (305-CD-026-002).

Planning HWCI (PLNHW)

This HWCI provides workstations (including user interface hardware), and servers as needed, to support production planning, the maintenance of planning data, and the interaction with and reaction to the processing environment during execution, e.g., to accept and process notifications of PGE completion and submit new DPRs.

5.5.1.6 Data Processing Subsystem

The Release B Data Processing Subsystem contains three HWCI: Science Processing HWCI, Algorithm QA HWCI and Algorithm Integration and Test HWCI. These HWCI are summarized below. For more information on these HWCI refer to the Release B SDPS Data Processing Subsystem Design Specification (305-CD-027-002).

Science Processing HWCI (SPRHW)

This HWCI provides all processing pools/strings associated with the following forms of processing: standard, reprocessing, on-demand, and testing. This includes processing platforms and working storage required during processing. The precise architecture of the disk storage resources used to stage data for processing or after processing for archiving is still under investigation and depends on technology decisions (see WKSHW).

The HWCI also includes workstations and servers for managing the production queues and dispatching processing requests.

Algorithm QA HWCI (AQAHW)

This HWCI provides the workstations, X-Terminals, and other devices needed for algorithm quality assurance (QA). For example, the HWCI supports the manual QA of algorithm results within the DAAC. The HWCI will execute the Client Subsystem, as well as additional user interface software needed to give the QA staff access to QA-related information.

Algorithm Integration and Test HWCI (AITHW)

This HWCI provides the servers, workstations, X-Terminals, and other devices needed by the algorithm I&T staff. Hardware needed to run tests in simulated production mode is part of the SPRHW. The HWCI will execute, for example, software development tools, test and integration tools, and the Client Subsystem.

5.5.1.7 Communications Subsystem

The Release B Communications Subsystem contains one HWCI, the Distributed Computing HWCI. This HWCI is summarized below. For more information on this HWCI refer to the Release B CSMS Communications Subsystem Design Specification (305-CD-028-002).

Distributed Computing HWCI (DCHCI)

The Distributed Communications Hardware CI (DCHCI) logically includes an enterprise communications server, a local communications server, and a bulletin board server. To provide for warm standby, the CSS servers and MSS servers at all DAAC sites and the SMC are cross-strapped and are configured to include the CSS Distributed Computing Software CI (including both OODCE client and server); the MSS Management Software CI; and the MSS Agent Software CI.

The complete configuration of the CSS and MSS HWCI, based on the combined requirements of the subsystems and site-specific requirements, are presented in the site-specific subdocuments. Additional detail on the analysis of MSS HWCI sizing and performance is contained in the MSS subdocument.

5.5.1.8 Systems Management Subsystem

The Release B Communications Subsystem contains one HWCI, the Management Hardware HWCI. This HWCI is summarized below. For more information on this HWCI refer to the Release B CSMS Systems Management Subsystem Design Specification (305-CD-029-002).

Management Hardware HWCI (MHCI)

This HWCI provides the servers and workstations needed to host the enterprise monitoring, local management and configuration management software, CM data, and backup copies of all ECS "infrastructure" software.

5.4.1.9 Internetworking Subsystem

The Release B Internetworking Subsystem contains one HWCI, the Internetworking Hardware HWCI. This HWCI is summarized below. Section 5.5, LAN Architecture, provides additional overview material pertaining to the ISS Hardware. The Internetworking Subsystem detailed design is not presented in its own subdocument. Instead the Internetworking Subsystem is presented in the DAAC specific subdocuments of DID 305. For more information on this HWCI refer to the DAAC and SMC specific subdocuments listed below.

305-CD-030-002	Release B GSFC DAAC Implementation
305-CD-031-002	Release B LaRC DAAC Implementation
305-CD-033-002	Release B EDC DAAC Implementation
305-CD-034-002	Release B ASF DAAC Implementation
305-CD-035-002	Release B NSIDC DAAC Implementation
305-CD-036-002	Release B JPL DAAC Implementation
305-CD-037-002	Release B ORNL DAAC Implementation
305-CD-038-002	System Monitoring and Coordination Center Implementation

Internetworking Hardware HWCI (INCI)

This HWCI provides the networking hardware for the intra-DAAC, DAAC to V0, DAAC to EBnet, SMC, and EOC connectivity, including: FDDI switches, concentrators and cabling; Ethernet routers, hubs and cabling; HiPPI switches and cabling; and network test equipment.

5.5.2 Performance Analysis Approach

Our performance analysis for Release B builds on the Release A CDR data analysis using the ECS Technical Baseline, dated February 1996.

For each of the hardware CIs, an analysis was performed to determine the performance and storage requirements for the various components. This analysis varied by subsystem, but used the following techniques:

- Dynamic modeling
- Static modeling and analysis
- Benchmarking

Issues that were considered in the analysis of performance and storage included failover / back-up strategies driven by RMA requirements and operational requirements (e.g., flexibility of configuration in order to enable test and operations, number of operator staff).

The methodology used to size hardware for each CI is described in the CI's portion of this document. A description of the performance analysis used to size hardware for specific DAACs, analysis results, and rationale for the choice of Release B hardware components, is contained in the DAAC-specific portions of this document.

Since Release B IDR, significant effort has gone into refining our dynamic modeling of the system. Dynamic modeling simulations have been run to verify and refine the approach briefed at IDR with respect to reprocessing; also to examine the effects of various changes in assumptions regarding pull and push parameters (e.g., 4X distribution vs. 2X distribution; increase / decrease in user subscriptions; decimation ratio; prioritization of processing requests). These have been briefed at ECS modeling workshops held in January and February 1996; results have been folded into the hardware analysis performed for CDR.

In addition, an end-to-end DAAC system model has been developed to ensure that all subsystems and flows are appropriately represented in a single system-level analysis. This model includes system management and communications infrastructure flows, as well as subsystem specific data and message flows required to perform activities at the ECS DAACs.

Since IDR, ECS has initiated and performed additional benchmarks on several components (both COTS and custom) to better understand performance constraints. Benchmarking results are cited in subsystem hardware rationales in the DAAC design documents.

The results of the modeling efforts will not only be used to enhance our hardware sizing, but as the input data and DAAC topologies stabilize, will be used to validate the design, by modeling system behavior under various operational conditions.

5.5.3 Release B Hardware Component Classes

The selection of COTS hardware components for Release B is the result of a continuing effort that began in the Release A timeframe to identify and validate candidate products, coordination of vendor demonstrations, hands-on use through the EP process and loaner equipment, and evaluation of vendor stability (past performance, financial status) and support. Other influences included upgradability/expandability, product training and documentation, and lifecycle cost. The results, as illustrated in Figure 5.5-2, Release B Architecture Component Classes, show the component classes and vendors selected for the overall DAAC hardware suite. The rationale for these selections is presented in Table 5.5-1, HWCI Class Selection Rationale.

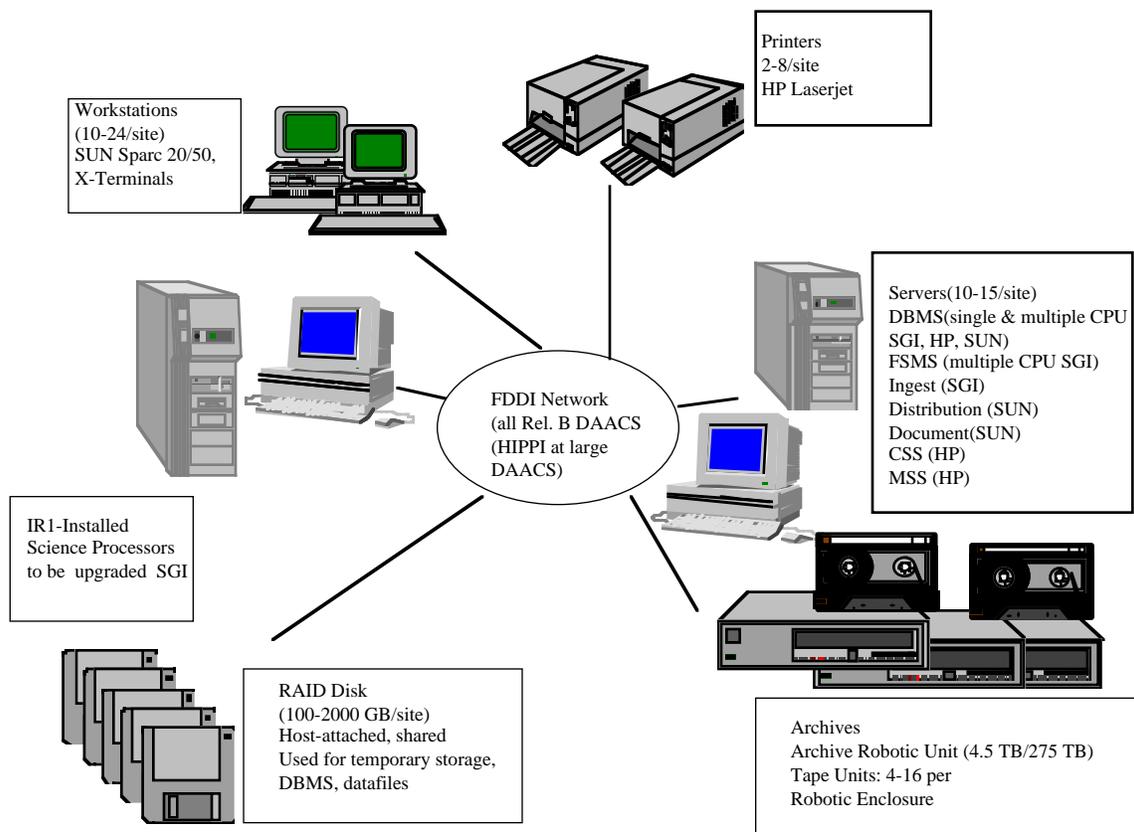


Figure 5.5-2. Release B Architecture Component Classes

Table 5.5-1. HWCI Class Selection Rationale (1 of 2)

Hardware CI	Platform Family	Rationale
Data Management Server HWCI (DMGHW)	HP K200 - Servers HP 715 - Ops Workstation Sun Sparc 20 - Workstations	Software compatibility; Suitability/ Scalability; Price/Performance Software compatibility; Price/ Performance; Price/Performance
Access Control & Management HWCI (ACMHW) <i>(Data Server Subsystem)</i>	SGI Challenge L and XL - Servers SGI RAID - Storage Archive Robotics at some sites Sun Sparc 20 - Workstations	Software compatibility with regards to total Data Server solution; Price/Performance
Distribution & Ingest Peripherals HWCI (DIPHW) <i>(Data Server Subsystem)</i>	Sun Sparc 20 - Servers HP LaserJet 4M - Printers	Software compatibility; Price/Performance
Data Repository HWCI (DRPHW) <i>(Data Server Subsystem)</i>	SGI Challenge XL - Servers Archive Robotics -Archives Sun Sparc 20 - Doc. Server	Software compatibility (archive solution primary processor is SGI); Robotics selected via Request For Proposal (RFP); Document server selected based on software compatibility.

Table 5.5-1. HWCI Class Selection Rationale (2 of 2)

Hardware CI	Platform Family	Rationale
Working Storage HWCI (WKSHW) <i>(Data Server Subsystem)</i>	SGI RAID - Storage Archive Robotics at some sites	Software and hardware compatibility with data server/ingest overall solution.
Ingest Client HWCI (ICLHW) <i>(Ingest Subsystem)</i>	SGI Challenge L - Server	Software and hardware compatibility with data server overall solution.
Algorithm Integration & Test HWCI (AITHW) <i>(Data Processing Subsystem)</i>	Sun Sparc 20 - Workstations HP LaserJet 4M - Printers	Price/performance; Price/performance.
Algorithm QA HWCI (AQAHW) <i>(Data Processing Subsystem)</i>	Sun Sparc 20	Price/performance.
Science Processing HWCI (SPRHW) <i>(Data Processing Subsystem)</i>	SGI Power Challenge XL - SP SGI RAID Sun Sparc 20 - Workstations	Suitability and scalability; Software and hardware compatibility; Price/performance.
Planning HWCI (PLNHW) <i>(Planning Subsystem)</i>	Sun Sparc 20 - Servers and Workstations	Software compatibility and price/performance.
Distributed Computing HWCI (DCHCI) <i>(Communications Subsystem)</i>	HP J210/1 - Servers Sun Sparc 20 - BB Server HP RAID	Software compatibility and price/performance.
Management Hardware HWCI (MHCI) <i>(Management Subsystem)</i>	HP J210/1 Server Sun Sparc 20 Workstations HP RAID	Software compatibility and price/performance.
Inter networking Hardware HWCI (INCI) <i>(Inter networking Subsystem)</i>	CISCO Routers Synoptics FDDI Concentrators, Cabletron Hubs	Result of RFP; Result of RFP; Result of RFP;

Notes:

- The following elements were taken into consideration in determining the vendor platforms:
 - Processing and storage capacities were determined for Release A and extrapolated out to Release B (this identified the scalability required for a particular platform);
 - COTS SW was assigned to each platform;
 - A survey across the DAACs with regards to platforms currently in place and in use by the science community and the DAAC operations staff was taken to determine which platforms could easily be integrated into the solution (suitability);
 - A common denominator based on equipment class (i.e. workstation, DBMS server, Science Processor), suitability and COTS SW availability was identified (this was done in order to obtain the best possible pricing (a much better price is achieved when many of the same configurations are ordered) on the most suitable and scalable platforms that were identified.):
workstation: SUN Sparc 20/50, server: HP J210/1; HP K200; SGI Challenge series; SUN Sparc 20/71, science processors: SGI Power Challenge XL; SGI Indy
 - Mass storage (RAID) could be supplied by platform vendors when required in the design;
 - Price/performance was the determination for the print selection
- Justification for each vendor was performed early on in the project, as the end result of the Request for Proposal (RFP) process.

5.6 Release B LAN Architecture Overview

5.6.1 Release B DAAC LAN Architecture

This section provides an overview of the DAAC network architecture during Release B. DAAC-specific topologies are presented in section 3.4.1 of the DAAC-specific volumes for each Release B DAAC.

The generic architecture for the Release B DAAC LANs is illustrated in Figure 5.5.1-1. The topology consists of a User Network (generally FDDI, but also switched Ethernet at some sites), a Production Network (FDDI at all sites), and a HiPPI Network (for processing to data server flows at some sites). 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' need to transfer large volumes of data. Each of the networks is discussed in more detail below.

The Production Network consists of multiple FDDI rings supporting the DAAC subsystems and connections to external production systems (such as EDOS 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.) At high-bandwidth sites, some Data Server hosts are contained on a dedicated FDDI ring in order to provide adequate bandwidth for DAAC-to-DAAC processing flow requirements. 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.

The User Network is an FDDI-based LAN (except at small sites, which will be switched Ethernet) connecting the users (via NSI, local campuses, general Internet, etc.) to the 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 the Data Manager hosts and to a subset of the 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 NSI, the local DAAC campuses, and other Internet providers through an ECS router which will provide the necessary routing and filtering controls.

The individual FDDI rings for both the User and Production Networks will be implemented with FDDI concentrators to provide ease of wiring and central points of management. All DAAC hosts will have FDDI interfaces and will be attached directly to the FDDI rings (although at some small sites, such as ORNL, the User Network will be implemented via switched Ethernet instead of

FDDI). 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. The interfaces of these machines that are on the User Network will have single-attached interface cards (or Ethernet cards where applicable). 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. Printers, which (generally) will be the only Ethernet devices in the DAAC, 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 documentation (305-CD-028-002) for more detail on BDS.

5.6.2 DAAC Addressing and Routing Architecture

The Planning and Data Processing Subsystems, the LSM, and Data Server/Data Management subsystems (collectively known as the Production Network) are connected to the FDDI switch on virtual LAN ports. They are assigned a class "C" address. The Ingest subsystem and the EBnet router are connected to the FDDI switch on routed ports. They are assigned class "C" subnet addresses. The User Network and SMC subsystem are connected to the ECS Router on routed ports. They are assigned class "C" subnet addresses. The Data Server and Processing subsystems are connected to the HiPPI switch. They are assigned private addresses as specified in RFC 1597.

5.6.3 SMC Network Architecture

The SMC network architecture, as illustrated in Figure 5.6.3-1, consists of two FDDI LANs connected to the GSFC DAAC router. The Enterprise Communications Server (ECS) and the Enterprise Management Server (EMS) connect directly to one of the FDDI rings, and the Management Workstations and printers are attached to Ethernet networks bridged to the FDDI ring via an Ethernet-to-FDDI hub. Since the Bulletin Board Server (BBS) is accessible by the general public, it is attached to a separate FDDI ring to facilitate increased security and to segregate BBS traffic from the rest of the SMC. (Section 5.6.4 below discusses network security in more detail.)

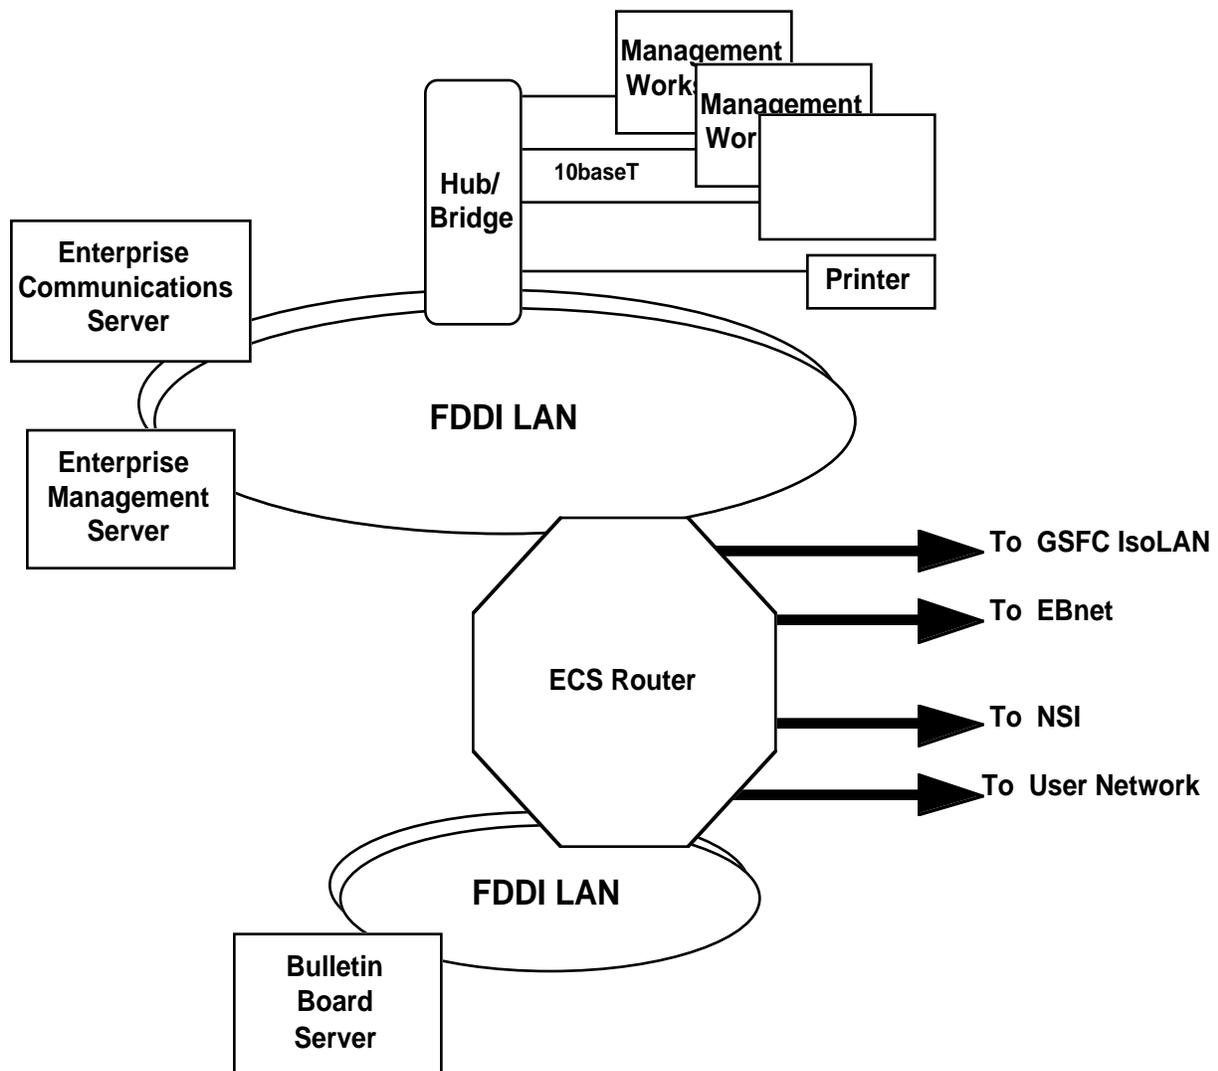


Figure 5.6.3-1. SMC Network Architecture

Because the SMC has been assigned requirements dictating very high availability, the FDDI LANs will be implemented via physically wired rings as opposed to concentrators. Physical rings eliminate concentrator hardware from the network and create a less complex topology, thereby increasing availability. The use of physical rings is feasible in this case due to the very small number of hosts on the FDDI network (two hosts on one ring and one host on the other). Of course, the workstations and printers will be attached to the Ethernet-to-FDDI hub, which will in turn be part of the physically wired FDDI ring.

5.6.4 Network-based Security Architecture

The Release B network architecture will provide basic levels of security to isolate and protect particular hosts and subsystems within the DAACs and SMC. Note that this section describes only network-based security; ECS has implemented other security measures, such as DCE-based authentication and authorization, Kerberized telnet and FTP, and DCE access control lists (ACLs), which are described elsewhere in this document.

The most basic security device employed at the network level is the creation of a separate User Network. This network allows the Production Network to be isolated so that no user access (e.g., no access from NSI, the campuses, or Internet) is required on the Production Network. Since most security threats will come from the Internet at large and not from within EOSDIS, this simplifies the security architecture since the only hosts that will be reachable from users are specifically designed to interact with users, and will therefore have all proper security mechanisms in place.

The Instrument Teams (IT) will need access to the Planning Workbench and a Science Processor in order to perform remote Science Software Integration and Test (SSI&T). The ITs will gain access via the User Network using the X-11 protocol. Access will be controlled by the ECS router which will only pass X-11 traffic from known IT hosts. At the Instrument team sites, access to their workstations will be controlled by a router which will only pass X-11 traffic from known ECS hosts.

At each ECS router connecting to external EOSDIS networks (such as EBnet) and external user networks (such as NSI), security filters will be in place to control access to the DAAC. These network and transport-layer filters control what traffic passes through the switch, and they are able to control access to individual hosts as well as to whole subsystems. Figure 5.6.4-1 shows a graphical representation of the DAAC network security architecture.

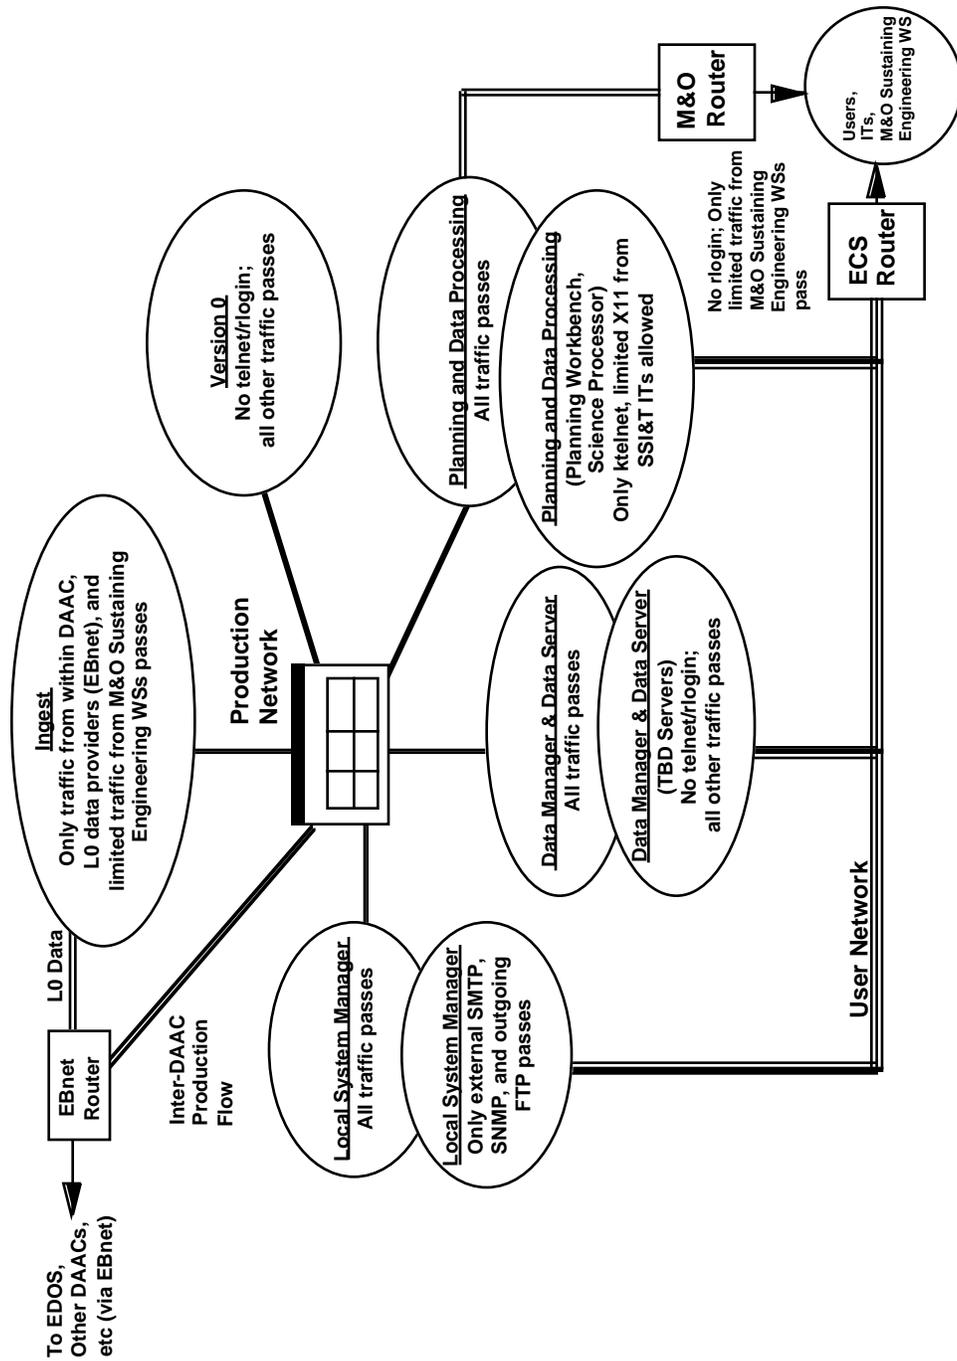


Figure 5.6.4-1. Network-based Security Architecture for Release B DAACS

The highest level of network security is associated with the Ingest subsystem, because of its role in receiving critical L0 data. No direct user access to Ingest will be permitted, and the only permissible traffic from outside the DAAC itself will be from the L0 data providers, such as EDOS and SDPF (through EBnet). All other traffic to Ingest will be filtered by the FDDI switch. For the LSM subsystem, no filtering will be performed on traffic originating from within ECS (e.g., other DAACs, EOC, SMC), but the only permissible external traffic is SNMP (for network management interaction with external network providers such as NSI) and SMTP (for mail messages from external systems). The PDPS subsystem is also secure, in that no traffic from outside ECS will be allowed. For the Data Manager and Data Server subsystems, only telnet and rlogin traffic originating from outside the DAAC will be prohibited; all other traffic will be allowed. For the interface to V0, no telnet or rlogin traffic will be allowed; all other traffic will be allowed.

The M&O Sustaining Engineering workstations which will be located within offices at each DAAC will require access to the Production Network and Ingest. Access will be controlled by a router which will pass limited traffic from M&O Sustaining Engineering workstations into the Production Network. This router will either be provided by the M&O organization or another port on a DAACs ECS router. The exact implementation will be site dependent.

No filtering is performed on traffic originating from within the DAAC.

The network security architecture for the SMC is illustrated in Figure 5.6.4-2. The design is similar to that of the DAACs, since the primary security mechanism is preventing interactive traffic from networks outside the SMC. The Enterprise Management Server and Enterprise Communications Server will allow all traffic originating from within ECS (e.g., other DAACs and the EOC, and limited traffic from the M&O Sustaining Engineering workstations), but will only allow SNMP and SMTP from external systems. Only telnet and rlogin will be blocked for the Bulletin Board Server, since its purpose is to provide information to the public. Only rlogin will be blocked from the M&O Sustaining Engineering workstations.

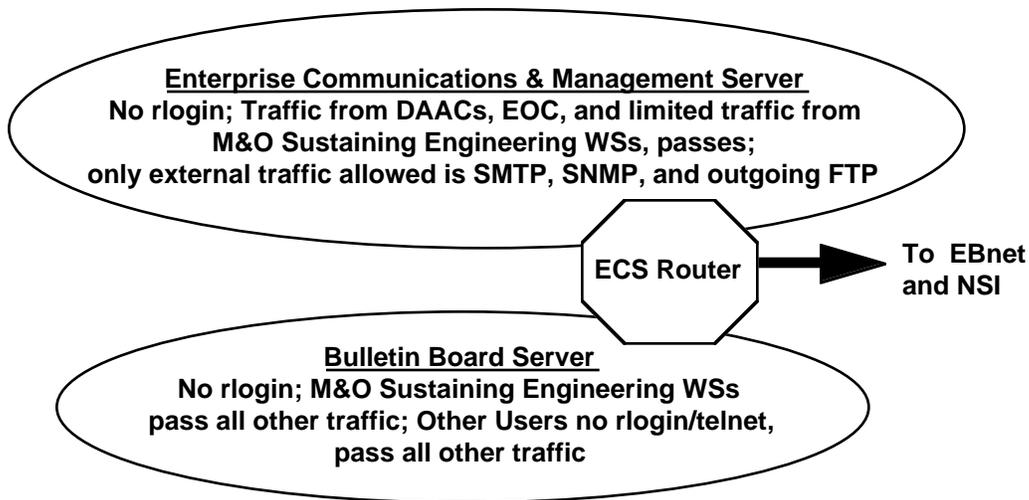


Figure 5.6.4.2 Network Security Architecture for the SMC

5.6.5 H/W and Network COTS Choices for Release B

5.6.5.1 Release B Network COTS Hardware

The HiPPI switches and ECS routers have not yet been selected for Release B. The FDDI switches, FDDI-to-Ethernet hubs, and FDDI concentrators are identical to the Release A hardware. The Release B DAAC and SMC LANs will contain five types of COTS hardware: FDDI concentrators, FDDI-to-Ethernet hubs, FDDI switches, FDDI/Ethernet switch/routers (ECS Router), and HiPPI switches. As described above, the FDDI rings within the DAACs will be implemented via FDDI concentrators, and the FDDI switch will be used to connect multiple Production Network FDDI rings together (refer to Figure 5.5.1-1). The FDDI-to-Ethernet hubs will be used to connect printers and X-terminals in the DAACs, workstations in the SMC. The FDDI/Ethernet switch/routers will be used to provide access to external networks (NSI and Campus nets) via the User Network, and the HiPPI switches will connect the Data Server and Processing hosts with a high-speed fabric to be used for transferring large volumes of data between the two subsystems.

5.6.6 Backup and Recovery Overview

5.6.6.1 Network Backup and Recovery

The DAAC and SMC networks are designed to provide exceptional availability coupled with quick and straightforward failure-recovery procedures. The required level of redundancy is determined by F&PRS (Level 3) functional RMA requirements. Specific RMA analysis for these functional requirements can be found in the Availability Models/Predictions Document (515-CD-002-002).

As mentioned briefly in the previous sections, the DAAC LAN FDDI rings will be implemented with FDDI concentrators. Each production ring will have at least two concentrators, and each high-end server or processor will be dual-homed to separate concentrators (e.g., each server/processor will be connected to two different units). This allows complete and uninterrupted Production Network connectivity to exist in the event of a concentrator failure or in the event of a damaged or severed FDDI cable. The switch-over to a backup concentrator for a dual-homed host is basically instantaneous and results in no data loss; in fact, the application processes will be unaware of the event and will continue as normal. Interfaces on the User Network will be via single-attached FDDI interfaces (and in some instances via Ethernet interfaces).

Since workstations are used primarily for administrative tasks and are thus not part of the critical processing flow, they are single-attached to single FDDI concentrators. In the event of a concentrator failure, several recovery procedures could be performed. The most simple, perhaps, is to simply move the FDDI cable of the effected workstation from the failed concentrator to the backup. This is possible because every FDDI ring will have at least two concentrators, each generally having spare capacity. If no extra capacity is available on the remaining concentrator, then another workstation can be used or the failed unit can be replaced. The failure recovery for the User Network is similar, since it will be implemented via single-attached FDDI interfaces connected to an FDDI concentrator.

Replacing the FDDI concentrator is simple and straightforward. The units are capable of performing properly at power-up without prior configuration (i.e., they are "plug and play"). (Note, however, that configuration would be required to configure the unit's SNMP management capabilities, but such configuration is not required for proper operation.) Thus, replacing a failed

unit is simply a matter of transferring cables to the replacement unit. This is also the case in replacing a failed FDDI-to-Ethernet hub.

The FDDI switch will be highly redundant. It will have n+1 redundant, load-sharing power supplies and cooling fans to allow the unit to operate if a power supply or fan fails. The FDDI switch will also have redundant control processors and hot-swappable interface cards. In the event of an interface card failure, the card will be replaced (while the rest of the switch operates as normal) and normal operation will continue without further reconfiguration (because the configuration data is contained in the control processor). (The FDDI/Ethernet switch/router will have a similar level of redundancy.)

If a control process fails, the redundant control processor will take over. Note that because L0 data is routed directly to the Ingest hosts without passing through the FDDI switch, the critical receipt of L0 data is not affected by switch failures.

The HiPPI switch will be a redundant device with high MTBF. Individual interface cards will generally be hot-swappable, so in the event of failure they can be changed without disrupting other hosts. If the control module fails, it would need to be swapped out and the switch reconfigured. In the very rare event of an entire switch chassis failure, the switch would either need to be replaced or repaired. All these failure recoveries involve activities on the switch; modifications to the attached hosts are generally not required.

5.6.7 Summary of Changes to LAN Architecture Since IDR

The following list summarizes the changes since IDR:

- There will not be an ECS DAAC installation at MSFC; therefore, hardware procured for MSFC will be reused at other sites.
- An IP-based protocol, BDS, will be used at those DAACs that require HiPPI, vs. "raw" HiPPI interfaces.
- The Data Management Server will be implemented with a single, redundantly-configured HP K-400, vs. two K-200s per site (resulting from price/performance evaluation).
- At DAACs with more complex scheduling requirements, planning and queueing functions will be implemented with Sun 20/71 workstations and a Sun UltraServer with two CPUs, to ensure adequate throughput for Autosys and Planning Workbench activities.
- The Data Server (DRPHWCI) will use NTP and 3490 linear tape technology in conjunction with D3 helical scan technology (resulting from price/performance analysis of current technical baseline requirements).

This page intentionally left blank.