NSSDC Provides Network Access to Key Data via NDADS

Jeanne Behnke and Joseph King
National Space Science Data Center
NASA/GSFC/Code 633
Greenbelt, Md 20771
behnke@nssdca.gsfc.nasa.gov
king@nssdca.gsfc.nasa.gov

$S_2 8 - 82$

$434 72$

$p. 7$

Abstract:

The National Space Science Data Center (NSSDC) is making a growing fraction of its most customer-desirable data electronically accessible via both the local and wide area networks. NSSDC is witnessing a great increase in its data dissemination owing to this network accessibility. To provide its customers the best data accessibility, the NSSDC makes data available from a nearline, mass storage system, the NSSDC Data Archive and Dissemination Service (NDADS). The NDADS, the initial version was made available in January 1992, is a customized system of hardware and software that provides users access to the nearline data via ANONYMOUS FTP, an e-mail interface (ARMS), and a C-based software library. In January 1992, the NDADS registered 416 requests for 1,957 files. By December of 1994, NDADS had been populated with 800 gigabytes of electronically accessible data and had registered 1458 requests for 20,887 files.

In this report, we describe the NDADS system, both hardware and software. Later in the report, we discuss some of the lessons that were learned as a result of operating NDADS, particularly in the area of ingest and dissemination.

1. Introduction

The focal point of the NDADS is the mass storage components of two Cygnet jukeboxes, each configured with two SONY 6.5 gigabyte optical disk drives. The two jukeboxes provide the NSSDC 1.2 terabytes of nearline optical disk storage. A VAX cluster computer configuration drives the two jukeboxes, as well as providing network connections to the NASA science community including NSI-DECnet, Internet and US SprintNet. Although the numbers of data sets in the space physics and astrophysics areas are comparable, about 90% of the NDADS data, by byte count, are astrophysics data. These data include a mix of data currently arriving at NSSDC, plus selected data being promoted from NSSDC's offline archives to NDADS. To date, NSSDC has focused on loading space physics and astrophysics data to NDADS. Key space physics data sets presently available from NDADS come from the IMP-8, ISEE-3, DE-1 and 2, Hawkeye, Yohkoh, and Skylab missions. Key NDADS-accessible astrophysics data sets typically include the basic observation data files and accompanying ancillary files (calibration, etc.). The astrophysics missions with data in NDADS are IUE, ROSAT, IRAS, Ginga, VELA5B, HEAO-1 and 2, OAO-3 and the Astronomical Data Center Source Catalogs.

The NSSDC developed the NDADS to support the following requirements:
- (1) the loading of data files to nearline storage and of associated metadata files to an inventory database;
- (2) user access to the (relational) inventory database;
- (3) user access to and retrieval of data;
- (4) data security;

(5) user understanding of the system (through online user guides, etc.);
(6) aggregation of files according to individual project needs;
(7) capability to support additional types of mass storage devices as acquired.

Item 6 on file aggregation is a special concept, whereby related files are grouped into predefined "granules" or "entries." Users are thereby able to request, for example, an astrophysical observation by unique granule/entry ID, and have the system retrieve and stage all the relevant files without the user having to specify each one. This feature makes NDADS more than a typical "file server" system.

The NSSDC must meet several obligations as part of its mission as an archive. One of the primary obligations is that the data must be kept safe and secure. Data integrity is an important requirement as well. Of equal importance is our obligation to disseminate data from the archive. For its own sake, the NSSDC must determine ways to archive the data that are scalable and cost effective. It is important to emphasize that the NDADS is much more than a file server, and hence the reason for the development of the specialized software system, discussed in section 2. Functionally and operationally, NDADS can be divided into two NSSDC activities, ingest and disseminate. In sections 3 and 4, we discuss some of the characteristics and lessons of the ingest and disseminate functions.

2. NDADS Software System

NSSDC developed a specialized software system to manage storing and locating data on NDADS. The NSSDC Storage System (NSS) software was prototyped in mid-1991 and experienced a highly successful two year "experimental" public access period resulting in a second version of the software system completed in 1993. The NSSDC required a system that would support data stored on multiple platforms (UNIX-like) as well as the VAX/VMS™ system platform used in the initial system. The resulting NDADS must also support migrations from the current given hardware and software platforms and mass storage systems. The current NSS software is written for a VAX VMS™ 6.1 platform and uses two commercial-off-the-shelf software packages; the SYBASE relational database management system and CYGNET Jukebox Information Management System (JIMS). It also uses the Software for Optical Archiving and Retrieval (SOAR) for formatting the WORM optical platters, a package that was developed at NASA and available through COSMIC. The modular NSS software is written in C Language to provide us a measure of portability. A client/server approach was used in the development of NSS, allowing a client located on a system outside the NDADS facility to access the NSS server on the NDADS host. The NSSDC also requires a direct applications interface to the NDADS giving the staff better access and control over the system to increase data ingest throughput. The NSS direct applications interface is available through a command line interface and C Call routines.

An important feature of the NDADS is a high level of security and recovery applied to the storing and staging of data from storage devices. The core NSS software processes the data to be stored as part of the transaction management features of the SYBASE. The 'store' transaction is performed in a sequential, 'batch' mode, first storing the pointer to the data on the mass storage system in the database and then actually storing the data on the mass storage device. Since the data is 'stored' as a transaction, any failure that occurs during the store process will trigger the operation to exit and notify the ingest team. Data granules can be tagged as non-proprietary or proprietary, thus restricting access to certain individual user accounts. Proprietary data is that data which has not been granted access to the general public. A complex 'logging' mechanism has been created to track all NSS steps and are used to monitor problems and performance.

The modular design of the NSSDC storage system allows device specific modules ("fetchers") for new storage devices to be integrated into the system quickly and with minimal impact on the rest of the code. Each fetcher module is expected to provide a certain small set of critical services to the "master fetcher", such as mounting a volume, copying a file onto the device, copying a file out of the device, etc. The system is designed to enable the NSSDC to add additional storage devices transparently to the external users without modification of the base software system. Currently, the NSSDC has fetchers for the Cygnet-SONY WORM jukebox, online magnetic disk devices and there are plans to include several other mass storage devices. The NSSDC recently augmented the NDADS with a Digital Linear Tape jukebox connected to an SGI Indigo 2/IRIX workstation (1/95). Figure 1 shows a conceptual design of the NSS system.

3. Ingest Lessons

The NSSDC expects to receive and ingest close to a terabyte of data per year beginning in 1996. To meet ingest requirements, the NSSDC has been studying ways to improve ingest rates. The NDADS ingest process is influenced most by the fact that the nearline system has been WORM disk-based. This fact results in many idiosyncrasies that drive NSSDC processes, for example, the slow transfer rates of the disks, the permanence of the write operation, and the limitation of the number of drives. The ingest process is composed of more steps than was described in the section 2 as part of the NSS software system. Typically, the ingest steps are:

1) assemble the data and determine data staging requirements
2) verify the data (check headers, gross bounds checking,...)
3) archive the data to nearline devices using the NSS software

Ingest is differentiated at the NSSDC by whether the dataset is current and arriving directly (electronically) from a NASA project or if it has been a resident of the NSSDC offline archive.
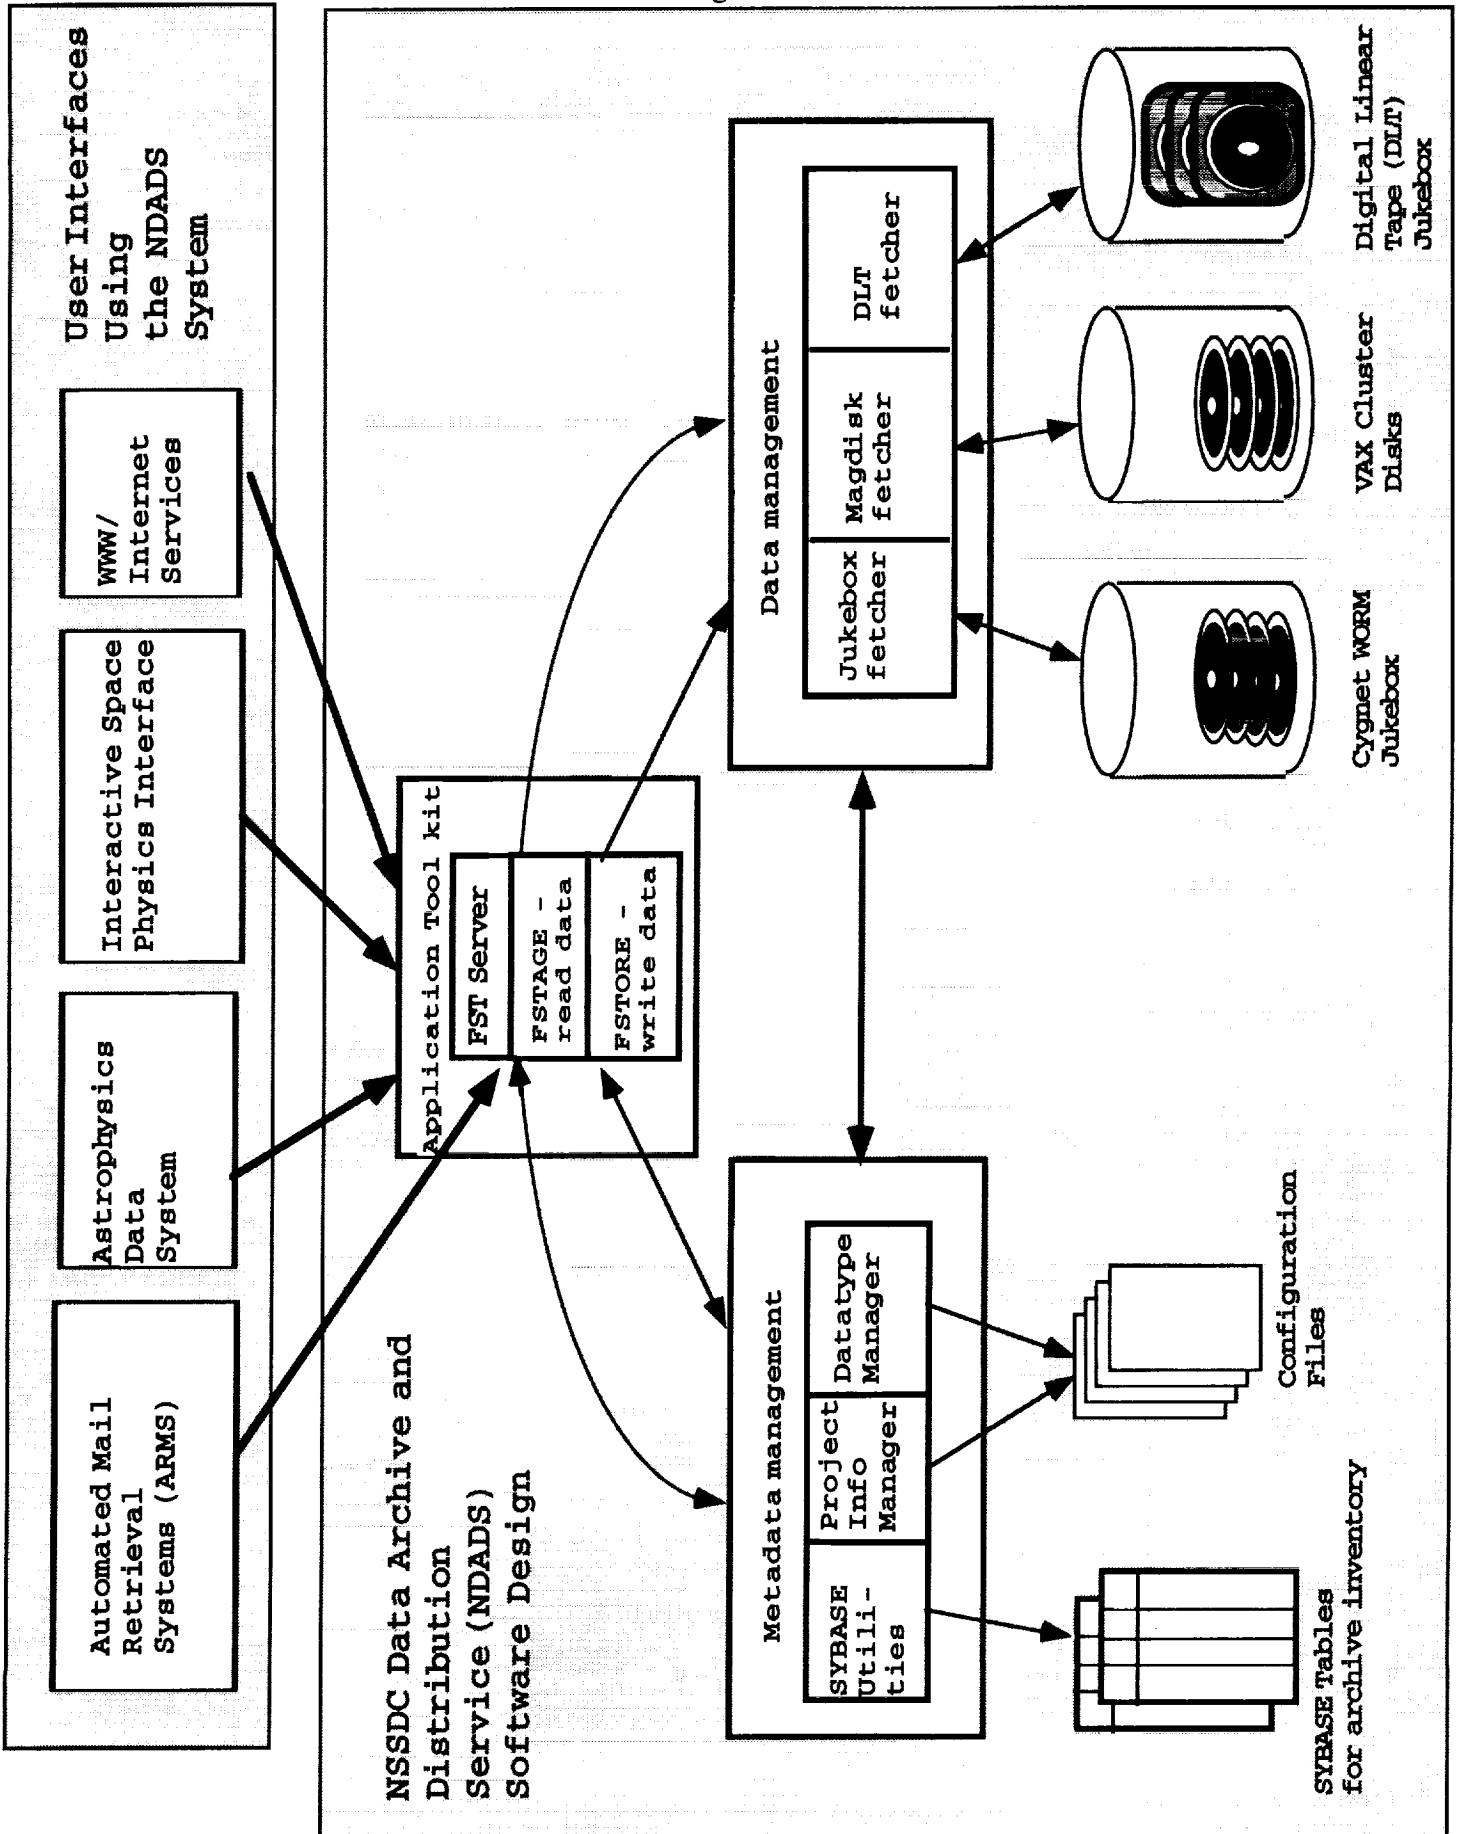
3.1 Offline Data

If the data is already in the NSSDC, it is typically one of 80,000+ 9-track 'legacy' tapes in the archive. In most cases, the data must be converted to files before it is placed in NDADS. Although this step requires customized software, the NSSDC reuses many software modules for the data conversion elements. This step can be time consuming based on the number of errors that are encountered in the dataset conversion process. Data in the archive has several common characteristics:

- it is always an 'old' dataset, often with limited documentation
- a dataset is typically all on the same media and has a finite size
- responsibility for the dataset is completely the NSSDC's
- it is difficult to predict how popular the dataset will be for electronic dissemination
- requires a high degree of human interaction to move the data into the archive

In the case of offline data, the NSSDC uses techniques learned from previous data restoration tasks. It is important to:

1) peer review these legacy datasets before selecting them for placement in the NDADS
2) vigilantly maintain a schedule for transferring the data to NDADS
3) select datasets that have good documentation to support the dataset
4) pre-determine the amount of verification required for storing the datasets

Figure 1

User Interfaces Using the NDADS System

- Automated Mail Retrieval Systems (ARMS)
- Astrophysics Data System
- Interactive Space Physics Interface
- WWW/Internet Services

NSSDC Data Archive and Distribution Service (NDADS) Software Design

Application Tool kit
- FST Server
- FSTAGE - read data
- FSTORE - write data

Data management
- Jukebox fetcher
- Magdisk fetcher
- DLT fetcher

Metadata management
- SYBASE Utilities
- Project Info Manager
- Datatype Manager

Digital Linear Tape (DLT) Jukebox

VAX Cluster Disks

Cygnet WORM Jukebox

Configuration Files

SYBASE Tables for archive inventory

362

5) pre-determine the amount of error correction required before the data is stored

The NSSDC must consider the 'setup' time associated with the above steps as well as the time spent preparing custom programs for reformatting tape data and data verification. We have discovered that a significant portion of manpower resources can be absorbed in these steps.

The NSSDC has reviewed different scenarios involved in ingesting different types of tape-based data to NDADS. Principally, 4mm, 8mm and 9 track offline tapes have been studied to determine the length of time involved in ingest. On the VAX cluster, our evaluation shows that 4mm and 8mm tapes are slower to physically ingest then the 9-track tapes. However, the set-up time for 9-track tapes is almost 4 times longer than that of 8mm and 4mm tape. The shorter set-up time is in partly due to the fact that the data on the 4mm and 8mm's is newer data and in some type of standardized format. The use of standards such as FITS, CDF, and SFDU simplifies the data verification phase as well as accelerates the step of converting to disk files.

3.2 Electronically Delivered Data

The NSSDC has been receiving newer datasets via the network. In these cases, the projects are still actively collecting data and transfer a processed dataset on a regular basis into NSSDC disks. If the dataset is delivered electronically, the NSSDC typically is only required to do basic checks of the data and then copy the data into the nearline system. Several characteristics make these datasets both easier for the NSSDC to work with and more difficult to control, for example:

- the NSSDC can review and affect formats of the data prior to their delivery
- both the NSSDC and the project share responsibility for the data
- easier to predict the popularity of a dataset and its eventual electronic retrieval
- software can be written to completely automate the ingest process, requiring little human intervention
- difficult to predict the quantity of data that will be delivered to the receiving/staging disk, thereby making it difficult to cost effectively determine the size of the disk

As part of the delivery function, the NSSDC contracts with each project a formal arrangement of delivering a list of what was transferred. These transfer lists are commonly referred to as Bills of Lading (or BOLs). In 1992, the NSSDC devised a BOL format that has served as a model for data delivered by other projects. The use of BOLs simplifies the NSSDC's the ability to cross check data delivered electronically by use of routine code. This permits us more accuracy and faster ingest into the nearline system.

The NSSDC does rudimentary verification and validation of the datasets before they are committed to the nearline system. Verification software is written in several programming languages, usually reusing existing code and often supplied by the data provider. The minimum set of tests is applied to newer datasets; i.e. check the filenames and header information, etc... The NSSDC will be working on ways to automate this aspect of data ingestion during this fiscal year. It is becoming increasingly clear that electronically delivered data must be spot checked rather than systematically checked given the large quantities received and the turnover rate from disk to nearline. It is difficult to find the CPU cycles to review all data received electronically.

The NSSDC staff has experimented with several different ways to schedule ingest and it remains are most difficult problem. Problems are routinely encountered in receiving

electronically delivered datasets, either due to system problems for both the project and the NDADS or due to network transfer delays. The data flow problem is compounded by difficulties in scheduling free staging disk space. Electronically delivered data tends to vary in size delivery-to-delivery. To alleviate these problems, it is important to get as much information on delivery plans from the project and to maintain close communications. The NDADS ingest staging space is planned to have available three times the maximum size of a delivery, this allows for potential hardware delays and unforeseen difficulties on NDADS. Along with scheduling of the ingest staging disks, we have in the past tried to manually map out the use of the optical drives to least impact the users who are retrieving data. This way we could insure that all of the drives in a single jukebox were not committed to ingest, thereby prohibiting access to the data for retrieval. This past year, we have developed selected batch queues controlled by the operating system to eliminate the manual intervention. This has improved our ingest throughput without affecting retrieval rates.

Many of the processes used to move the data through ingest pipeline are manually executed and monitored. An ingest team member will manually start one of the steps and monitor to completion. Following successful completion, another job is started and in some cases the jobs are performed in the batch queue. In our evaluation, manual pipeline processing nominally requires at least 4 hours per dataset. By eliminating manual pipeline processing for several electronically delivered datasets, we have increased the ingest throughput without affecting the quality of the load. The steps used for automated ingest of the data are often similar from project to project. The NSSDC is collecting these common steps into a generic ingest software system that can be customized with appropriate configuration files and used on any new dataset to be ingested into NDADS. Because of these measures, the NSSDC shows an increase in ingest rates in 1994, see Table 1.

1994 INGEST RATE IN GB											
JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
9.6	10.2	3.5	18.2	4.4	15.3	21.3	35.2	11.5	17.2	32.4	85

TABLE 1.

4. Disseminate Lessons

The NSSDC is committed to providing its users, both in-house and outside community, four ways to access the NDADS archive:
1) via command line interface
2) via C callable routines
3) via FTP service
4) via an E-mail interface

The first two methods are used principally in-house to directly manipulate the nearline mass-storage systems for better management of ingest and disseminate functions on the NSSDC's behalf. Methods 3 and 4 are provided principally for the outside community. The NSSDC Automated Retrieval Mail System (ARMS) provides an E-mail interface to the NDADS archive. Users send an E-mail request to the account *archives@nssdca.gsfc.nasa.gov*. Within the message, users specify the need for information or data files by adhering to a fixed protocol for the content of the E-mail SUBJ line and, for data requests, by specifying granule ids in the body of the E-mail message. The ARMS Users Manual, detailing the protocols, may be obtained by specifying MANUAL as the subject of the message and leaving the message body blank. The E-mail system is very popular and has supported the distribution of over 260 GBytes of NDADS

364

data. In 1995, we will be working on providing a more fault-tolerant and modular ARMS system to our customer community.

The E-mail system has to its advantage a simplistic interface, but it also requires users to understand the NDADS granule-naming conventions and the granule-file hierarchy. Because of this requirement the NSSDC developed an FTP server to NDADS that makes the full NDADS archive appear to users as a massive FTP-accessible disk farm. The FTP interface allows the NSSDC more versatility in connecting to client/server based user interfaces. One advantage of the FTP service is that NDADS files now have Uniform Resource Locators (URLs). The FTP service incorporates well into World Wide Web pages developed at the NSSDC by space physics and astrophysics disciplines. These Web pages allow retrieval of NDADS data without specifically knowing granule names.

5. Conclusion

The NDADS has been developed to serve the specific needs of the NASA science community. It combines specialized hardware with customized software to significantly enhance the power of the NSSDC scientific database system. The success of this facility can be measured in several ways: the number of requests for data, the turnaround time, capacity, and convenience to the community. Available 24 hours a day every day, NDADS currently satisfies in excess of 1000 requests per month in an average of less than ten minutes. The NDADS service represents three-quarters of all NSSDC data requests. NSSDC believes its NDADS nearline data management environment is evolvable to exploit future changes in both hardware and software. By providing a well-constructed and secure infra-structure, NSSDC will be able to meet the future requirements of managing terabytes of data, cooperatively supporting NASA missions and supporting user interfaces that rapidly change to best meet the needs of scientists and others on the information superhighway.

In the future, the NSSDC expects to need additional storage devices to support the growing archive. The inclusion of the data and storage devices in use at the HEASARC, Compton-Gamma Ray Observatory and other related archives will be of primary importance to the NSSDC as well as intriguing in its possibilities of resource sharing across organizations. Careful planning and consideration will be required to phase-in the future computing requirements of the data center and not disrupt existing capabilities. The NSSDC will also consider improved access to the NSSDC data through Wide Area Information Service (WAIS), World Wide Web (WWW) and related network-based services as well as software application systems used in-house.