6.5 THE NOAA NATIONAL OPERATIONAL MODEL ARCHIVE AND DISTRIBUTION SYSTEM (NOMADS): GROWING PAINS AND A LOOK TO THE FUTURE

Glenn K. Rutledge *, Danny Brinegar, Andrea Fey
NOAA National Climatic Data Center, Asheville, North Carolina

Jordan Alpert
NOAA National Centers for Environmental Prediction, Camp Springs, Maryland

Dan Swank
STG Inc. Asheville, North Carolina

Michael Seablom
NASA Goddard Spaceflight Center, Greenbelt, Maryland

1. INTRODUCTION

The National Oceanic and Atmospheric Administration (NOAA) National Operational Model Archive and Distribution System (NOMADS) project recently completed its third year of operation. NOMADS (Rutledge, 2005) servers now provide up to 4 terabytes (TB) of model data from 1-4 million individual downloads every month. NOMADS is a Web-services based pilot project that leverages community and in-house software development efforts across multiple institutions to provide and promote distributed format neutral access and model inter-comparisons to Numerical Weather Prediction (NWP) model input and output, climate models, reanalysis, reference quality in-situ datasets and now under development satellite datasets.

2. BACKGROUND

To address a growing need for retrospective Global Climate Model (GCM) and NWP input and output data, the National Climatic Data Center (NCDC) along with the National Center's for Environmental Prediction (NCEP) and the Geophysical Fluid Dynamics Laboratory (GFDL) initiated the highly collaborative NOAA National Operational Model Archive and Distribution System (NOMADS). NOMADS (see http://nomads.ncdc.noaa.gov and Figure 1) provides real-time and retrospective access to model and observational data by a wide variety of users using the Internet. NOMADS is an inter-operable network architecture with fully integrated data access and manipulation tools using a distributed, Web-services based format independent methodology. NOMADS allows temporal, spatial, and variable sub-setting to address the ever increasing spatial resolution, volume and varied formats of models necessary for archive and access at NCDC. To permit format neutral Web-services based data access, NOMADS uses the Open-source Project for a Network Data Access Protocol (OPeNDAP) transport protocol (formerly called DODS), (Gallagher, 1995; Davis l999). OPeNDAP provides "format-neutral" access to data in many different formats and locations.

The scientific modeling community is a vast intellectual resource. This community is extremely interested in obtaining GFDL, NCAR, and NCEP products for historical cases and for operational and research purposes. The community is also interested in collaborating to improve weather and climate forecasts, but until now virtually no resources existed to support this collaboration.

3. BENEFITS

Under the NOMADS framework improvements to models will lead to longer lead times and improved accuracy for forecasts over all time scales. This results in fewer lives lost and less property damage. NOMADS will allow researchers to more easily access and understand data thereby increasing scientific productivity. Using increasingly detailed and complex models, NOMADS provides a capability for scientists to assess the potential impact of long term climate change on the Nation since various models will be more readily available for inter-comparison.

The NOMADS solution reduces software costs NOAA-wide by integrating local efforts to archive and gain access to historical model data and for the first time creates a unified archive of climate and weather data. The NOMADS project leverages many man-years of software development within NOAA and elsewhere.

---

* *Corresponding author address:* Glenn K. Rutledge, NOAA/NCDC, 151 Patton Ave., Asheville, NC 28801; e-mail: Glenn.Rutledge@noaa.gov.

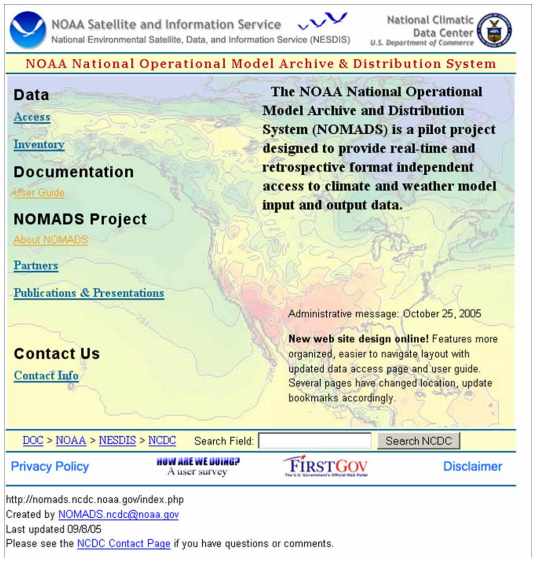

Figure 1. The New NOMADS Web Page

Heavy dependence is made upon technology transfer and re-use of previously developed successful technologies such as the OPeNDAP, the Center for Ocean-Land-Atmosphere Studies (COLA) developed GrADS Data Server (GDS, Doty, et al., 2001), and the Pacific Marine Environmental Laboratory (PMEL) developed Live Access Server (LAS, Hankin, 2001). NOMADS allows an option for the integration of large datasets by performing some of the data manipulation, sub-setting, and compression on the servicing host rather then transmitting large amounts of data across relatively slow speed lines. NOMADS can easily be expanded to include other datasets such as satellite and radar data.

4. NOMADS WEB ENHANCEMENTS

The expansion of the NOMADS project and new dataset availability resulted in a disjoint collection of Web pages with confusing navigation. The NOMADS Web site contained hundreds of pages of documentation and instructions developed on an as-needed basis. Many of these pages were not cross-referenced and contained redundant information. The NCDC NOMADS team recently consolidated this information into a more intuitive, user-friendly site which includes a new user guide. Future plans are to expand the user guide and to compartmentalize dataset documentation using "PHP" in order to offer information to the user in multiple formats based upon individual need.

4.1 *NOMADS Live Access Server*

The efforts to merge the NOMADS Live Access Server (LAS) to the existing GDS were met with several challenges, particularly surrounding dataset aggregation and incorporation of in-situ data. The NOMADS LAS (http://nomads.ncdc.noaa.gov:8085/las/servlets/dataset) requires OPeNDAP data to be aggregated. NOMADS has successfully implemented the LAS with several model archive datasets including the Global Forecast System (GFS), the North American Meso-scale (NAM, formally Eta), the North American Regional Reanalysis (NARR), and the National Weather Service (NWS) Service Records Retention System (SRRS), the legally mandated NWS required archive suite of data. NOMADS then investigated a real-time access system for these data. However, updating LAS for the high volume real-time NWP is computationally expensive. Currently it takes up to nine hours to update the LAS for each set of model output fields that are ingested into the system. NOMADS is investigating whether we can improve the performance of these update scripts or whether we need to break the aggregate sets into large archives and smaller, frequently-updated real-time datasets. Additionally, a NOMADS goal is to provide NCDC reference quality in-situ datasets. One such dataset is the Integrated Global Radiosonde Archive (IGRA, formerly the Comprehensive Aerological Radiosonde Dataset or CARDS) dataset currently served from the NOMADS GDS into the LAS. Unfortunately the LAS is unable to handle GDS sequence data without construction of a dataset-specific driver. To overcome this limitation, the dataset was placed into a "MySQL" database and used an existing LAS ocean dataset driver to serve the data. This dataset remains in experimental phase, however, while the development of a process to keep the MySQL database updated while developing advanced LAS display capabilities which are developed primarily for the needs of the ocean community. The PMEL LAS software developers are working on a generic in-situ dataset driver and it is hoped that the NCDC NOMADS efforts can assist in this development.

5. NOMADS ARCHITECTURE

The hardware architecture in use by the NOMADS project at NCDC is a classic file and application server model. Within this architecture the Network File System (NFS) distributed file system is deployed over a private gigabit Ethernet Local Area Network (LAN). At present, a pair of Dell PowerEdge™ 2850 servers running Red Hat® Enterprise Linux® WS3 act as file servers. Between these servers a total of approximately 20 TB of direct attached

Redundant Array of Independent Disks (RAID) level-5 storage is shared with the application server via read-only NFS. Additionally, these servers carry out an ingest function for the data feeds served by NOMADS.

On the application server side, a single Dell PowerEdge 2850 server running Red Hat Enterprise Linux WS3 is deployed. This server currently hosts the many applications needed for public access to the NOMADS data holdings. These applications include an Apache Web server, two GrADS Data Servers (GDS), a Live Access Server (LAS), and several smaller applications used for sub-setting and index file generation. On average, this system currently handles a million hits per month and serves 2-4 Terabytes (TB) of data per month. Given the load this system handles, we have experienced relatively few problems to date. However, with the steadily increasing volume of data being served from NOMADS as well as a user base which continues to grow, there is a need to address the potential issues of scalability, reliability, and availability. The main problem with the current architecture is the fact that the data path and each server are single points of failure for the system. To address the problems on the application server side a new architecture is being implemented to use a high availability (HA) cluster system. This new architecture removes one of the single points of failure from the current system and provides for both redundancy and load balancing. Addressing redundancy for the file server side has proven to be more of a challenge due to the very high data volume involved. NOMADS is investigating several possible solutions such as data mirroring, multipath input/output (MPIO), and moving to a storage area network (SAN).

6. DEALING WITH THE VOLUME

In May 2005, NOMADS moved away from the NOAAPort and Local Data Manager (LDM) ingest systems which was in place to obtain near-real time thinned, low resolution numerical model grids, toward a direct data feed of the entire model from NCEP. This change was in part due to staffing requirements to maintain these additional ingest feeds. NOMADS, an unfunded pilot, could no longer afford the staff time required to maintain these ever-changing feeds. The new direct connection to the mainframe at NCEP has resulted in a significant increase in the data volume from approximately 75 megabytes per month to approximately 450 megabytes per month. The elimination of the NOAAPort system has also resulted in the loss of the Rapid Update Cycle (RUC) model. However, NCDC is in discussions with the NWS to receive direct ingest of the high resolution RUC-13km output.

This volume increase has had several implications to the overall NOMADS system at NCDC. In its current state, the NOMADS system can keep approximately seven months of high-resolution model data on-line, as opposed to several years of NOAAPort model data. This change has implications to users with slower internet connections downloading these larger volumes of data, even after sub-setting. The new ingest typically has a two- to three-cycle delay from real-time, however, as these data first pass through the NCDC ingest process.

Although delayed from real-time, these new high-resolution models provide a complete archive of the North American Mesoscale (NAM) and Global Forecast System (GFS), including all parameters, vertical levels, and forecast time steps.

NOMADS primary OPeNDAP server is the GrADS Data Server. As the number of data files on NOMADS increases, the time it takes for the GDS to refresh its inventory (or “reload”) also increases. When the NOAAPort archive approached 2 years of on-line archive, the reload time began to become a problem. It took up to 10 minutes every half hour. The reload time depends more on the number of files than the file sizes. Moving away from the NOAAPort feed to high-resolution NWP alleviated this issue slightly since there was a reduction in the ingest and archive of the many different grid projections as found on NOAAPort. Another counter measure to improve reload performance was to set up two different GDS instances: one for datasets with frequent cycles (and frequent reloads), and another that holds more static datasets such as reanalysis, upper air data, and climate models. With these measures in place, reload times are minimized, which equates to less server downtime. Further improvements as discussed in the preceding section for load balancing as well as additional instances of GDS serving individual datasets are in progress.

7. DATA ACCESS METHODS

The NOMADS system currently has 20TB of on-line disk space on its data server. However, the NCDC-NOMADS archive far exceeds this amount. To gain access to the off-line archive, NOMADS developed access to the NCDC archive through an interface to the High Performance Storage System (HPSS), the deep archive in use at NCDC. This interface is required to allow users access to the entire period of record of the NOMADS archive. If user selected data is not available on-line, an automatic off-line ordering system is available to move the off-line data in the HPSS to the on-line NOMADS data

servers for sub-setting, Web browsing, or animation. Only the most recent data are kept on-line while older data are migrated into the deep archive. Before the data are migrated, however, a copy of the inventory and support files for each data file are created and retained on-line. These inventory files are used by the applications server to build and display an off-line data availability range. The users can then select from this off-line data range and place a request through the NOMADS Web interface to stage data from the deep archive to the on-line data server. Submitting an order initiates a sequence of processes that request the data from the HPSS, uncompress the files, and stage it so that the data appears identical in form to on-line data. This staged data can then be acted upon by the NOMADS sub-setting and access and display scripts by the user. After the user selects off-line data for staging, the system provides an e-mail notification to alert the user when the order has finished processing and is ready for access. The current system has some limitation with respect to the available space for staging (approximately 1TB). If many requests for high volume data from either a single or multiple users are received and processed, this space can quickly be entirely utilized. This space is freed-up on a regular basis, after the user has obtained the requested data. In addition, the order processing system is single-channel. This means if many large requests are placed in a short period, they all need to be processed linearly in the order they are received. This, on occasion, may take so long that orders expire before they are completed. Additional staging disks are being purchased and efforts are currently underway to increase the number of independent channels thereby increasing the number of concurrent requests to the HPSS system that can be performed.

Methods are also available for advanced access for users that need individual model variables or atmospheric levels of raw GRIdded Binary Code FM-92 (WMO GRIB, 2001) files. A series of PERL scripts were developed in the NOMADS collaboration between NCEP and NCDC to allow partial Hyper-text Transfer Protocol (HTTP) transfers. To achieve this capability, NOMADS regularly creates inventory files that list the byte position of all the GRIB messages with corresponding variable and atmospheric level information. This provides the user with the capability to access a single GRIB message (variable or level) from a file that contains hundreds of messages. This process, depending on how many parameters and levels are requested, reduces the amount of data that needs to be transferred. A new comprehensive NOMADS Users Guide has been developed to help users understand and use these advanced capabilities and can be found at: http://nomads.ncdc.noaa.gov/guide/. This method is currently available for the NARR, NAM, and GFS products.

8. NEW DATASETS AND SERVICES

The NCDC NOMADS is expanding the data types it serves. Along with new model output from NCEP, the NCDC NOMADS now provides limited subsets of the latest Intergovernmental Panel of Climate Change (IPCC) GCM Climate Model (CM) output developed and served primarily by the GFDL NOMADS data portal (see http://nomads.gfdl.noaa.gov/). NCDC also has several indices derived from the long-term CM model trends as developed by GFDL of precipitation, heat wave duration, and drought.

NCDC NOMADS is the primary archive for the North American Regional Reanalysis (NARR), a high volume (4.3TB) dataset which consists of reanalysis of the North American Mesoscale (formerly ETA) model from 1979 to the present at 32km resolution. This dataset is updated monthly. Users are becoming more adept and aware with the capability that an entire dataset is not required in many research applications. NOMADS offers a wealth of options to “slice and dice” large data sets for the most efficient use of available band width and computer resources. The NARR data set is a good example of how NOMADS client applications work with the server to subset and organize needed data sets. These client applications include HTTP, FTP, GrADS Data Server, and an Interactive Web Interface. NARR data have been aggregated in three ways: daily, monthly, and a few select variables for the entire dataset. These aggregates were primarily created to prevent the need for users to write scripts to loop over hundreds of daily files just to get a subset of a single parameter.

NOMADS recently made available the NCDC Integrated Global Radiosonde Archive (IGRA) reference quality dataset under the GDS and a database powered Web interface for fast access to in-situ observations. This was accomplished by using scripts to place the entire archive into a database table and then creating Computer Generated Interface (CGI) scripts and interface system to collect input from the user, form the appropriate queries and then format the results into a user-friendly format for access or display. Another set of scripts were used to move the data from the database into binary files compatible with the GDS OPeNDAP server.

Other new datasets now available on the NCDC NOMADS system include the Smith-Reynolds Optimum Interpolation (O/I) Monthly Extended Reconstructed Sea-Surface

Temperatures (ERSST). The NCDC Global Historical Climate Network (GHCN) is also provided in OPeNDAP accessible form.

NCDC NOMADS has recently provided public access to several test datasets under the new Unidata THREDDS Data Server (TDS). Datasets available under the TDS include the International Satellite Cloud Climatology Project (ISCCP) Satellite dataset; and two experimental NCDC datasets including the High-resolution Infrared Radiation Sounders (HIRS) Clear Air Satellite dataset, and a Blended Ocean Wind dataset. The TDS and these new datasets can be accessed from the main NCDC NOMADS Web page.

NASA's Global Modeling and Assimilation Office are planning a major reanalysis that will span the "modern era" of satellite observations. The goal of the Modern Era Retrospective-analysis for Research and Applications (MERRA) effort is to generate physically consistent research datasets for the scientific community, with special emphasis on the study of the hydrological cycle. The effort will produce tens to hundreds of terabytes of model output. Goddard's Software Integration and Visualization Office, working with the local Distributed Active Archive Center, is investigating new methodologies for rapid, seamless distribution of the data that will allow inter-comparisons with model output from other agencies. A prototype MERRA – NOMADS application is currently being developed using the NOMADS GDS and LAS interfaces for real-time inter-comparison to NOAA NWP and for near real-time access. NOAA NWS and NCDC are also involved in new reanalysis efforts and NOMADS has been requested to serve many of these new reanalysis products when available.

Finally, NCDC has a new mirror backup server capability for the NCEP Real-Time (R/T) server. This server will contain a mirror copy of the NCEP server for access to real-time data. This server will also act as an ingest mechanism to many new NOMADS available datasets to include the Coupled Climate Forecast System (CFS), high and low resolution Ensembles, and the Global Data Assimilation System (GDAS) model input data. See http://nomads6.ncdc.noaa.gov/ for the current location of the NCDC installation of the NCEP R/T NOMADS.

9. FILE FORMAT COMPATIBILITY

NCDC NOMADS developed a capability for access and plotting to the new NWS National Digital Forecast Database (NDFD). NDFD uses the relatively new GRIB Edition-2 (GRIB-2) as its file format. The backbone of the NOMADS distributed data access capability is the GDS and the GDS is dependant upon the Grid Analysis and Display Systems (GrADS) (Doty, 2001) package to read the data on the server, subset, and return the results in a format neutral format. Currently, GrADS supports common data file standards such as GRIB-1, HDF, ASCII, NetCDF, and to a limited extent the Binary Universal Form for the Representation of meteorological data (BUFR), a World Meteorological Organization (WMO) standard file format. GrADS does not have support for GRIB-2, certain forms of BUFR, and text data. For these formats, conversion or pre-processing is being performed. GrADS currently uses GRIB1 (Edition-1). Recently, decoders have emerged to convert GRIB-2 into GRIB-1 and a GrADS binary file. Converting the GRIB-2 data was initially unsuccessful due to an incompatibility in the converter with the packing method used for NDFD. Once a patch was developed by NCEP for this purpose, NOMADS can now successfully display NDFD on user demand. GRIB2 has the ability to store data using various packing methods, any encoder or decoder for this data must be able to account for any of these. Until decoders and display software packages such as GrADS improve their GRIB-2 support, NOMADS will have to deal with the setback of having to convert files from GRIB-2 to an alternate format.

In many ways GRIB-2 is superior to GRIB-1. GRIB-2 can improve data set compression. Presently GRIB-2 inhibits NOMADS sub-setting capabilities. In GRIB-1, the data was packed in such a way that a program could utilize random access to jump to a specific position in a file and extract a single point of data from an entire grid. This was indispensable for extracting a time series from a single point location for months or years for hourly GRIB-1 data. NOMADS software efficiently pulls the data points from the GRIB-1 records without having to unpack the entire grid. With GRIB-2 complex packing methods data sub-setting under the current NOMADS architecture are not available because decoders will need to unpack the entire grid just to get at the single points of data that are desired. Although this sub-setting capability will still be available in such cases, the computational overhead is the same as unpacking the entire grid just to get one data point. This raises concerns about how useful this will be when very large files are involved. The data may be left in the GRIB-like simple packing; however, this gives no added benefit to NOMADS other than improved metadata and compression capabilities. NCDC NOMADS has already initiated an effort to develop GRIB-2

decoders for a new open-source version of the GrADS application package.

BUFR has the ability to represent nearly any form of data, from soundings to aircraft reports, or even grids of points. It also utilizes two forms of compression. Because of such a broad scope of possibilities, creating a BUFR decoder that works with all types of data becomes extremely difficult. For instance, while attempting to decode upper-air collectives from BUFR messages as received at NCDC using the decoder supplied within GrADS, errors occurred since the BUFR messages which contained more then five subsets were compressed using the special "BUFR compression" which the GrADS decoder did not support. To overcome this limitation, an alternative decoder was developed at NCDC. Using this new BUFR decoder the data were then successfully decoded, and subsequently GrADS "Station Data" files were generated.

The shortcomings of the NOMADS system being dependant on GDS (or the LAS which uses the GDS to access GRIB and BUFR datasets) have become evident. If GDS does not directly support a dataset's file format, then an extra step must be invoked to convert the dataset into a format in which GrADS is compatible. During this process, there is always the possibility to introduce data or metadata corruption or loss and inter-comparison of datasets that are currently not begin performed. Data in the NCDC deep archive however are always stored as they are received and therefore original data are being preserved. Tests comparing original versus "converted" data have shown no degradation to the original data, however this is an area of concern.

10. FUTURE STEPS

NOMADS has become an international resource for hundreds of users downloading TB of model and other data every month. NOMADS will continue to provide this service and develop new partnerships to increase inter-operability for various services now being developed throughout the geo-science community. NOMADS will actively engage with the CLASS data access system as a partner and provider to that overarching NOAA system and develop generic interfaces coupling NOMADS and CLASS data systems. NOMADS is actively advancing new Web-services based distributed access in emerging technologies such as the Open-GIS and OPeNDAP-G. OPeNDAP-G is a grid enabled OPeNDAP service for use on the Grid. NOMADS is an active collaborator within the Committee on Earth Observation Satellites (CEOS) CEOS-Grid activity.

The NOMADS team is also involved into serving various forms of satellite and radar data from NCDC's Remote Sensing and Applications Division (RSAD). This will be a huge increase in volume and will likely require additional servers and disks. A third GDS may also be needed once the added load is accessed. The already robust Web interface will continue to be upgraded to service different types of data. Planned features include binary format conversion (NetCDF, and GRIB output), time-height cross sections, skew-T Log-P diagrams for station and gridded data, multi-panel plots, land/ocean masks, and improved support for GrADS station and BUFR datasets. To lessen NOMADS dependence on GrADS, other forms of display and analysis tools are to be explored. As noted earlier, the LAS is functional and tested with promising results. Other OPeNDAP enabled servers are to be investigated, including the new THREDDS Data Server form Unidata. The NOMADS Web-service philosophy will continue to include model data inter-comparison capability and model to in-situ is to be expanded.

NOMADS is cited as a pilot Web-service to promote under NOAA's new Data Management Integration team (DMIT) DRAFT Global Earth Observations Integrated Data Environment (GEO-IDE) plan. GEO-IDE outlines a long-term NOAA vision to develop loosely coupled Web-services using standards to build a NOAA "system of systems" compatible with the US-Group on Earth Observations (US-GEO) a US National effort that also cites NOMADS as a candidate project for interoperable access to data. Given the above, NOMADS will continue to harden its access services and develop a true operational system for the community. It is expected that the emerging NOMADS-CLASS relationship is a staring point from which the goals of the NOAA GEO-IDE framework, readily applicable to NOMADS, can be institutionalized throughout NOAA.

11. REFERENCES

Davis, E. R., J. Gallagher, J., 1999: Using OPENDAP to Access and Deliver Remote Data. *Proceedings of the 15th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, Amer. Meteor. Soc., Dallas, TX. 571-573.

Doty, B.E., Wielgosz, J., Gallagher, J., Holloway, D., 2001: GrADS and OPENDAP. *Proceedings of the 17th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, Amer. Meteor. Soc., Albuquerque, NM. 385-387.

Gallagher, J., and George Milkowski, 1995: Data Transport within the Distributed Oceanographic Data System. *Fourth International World Wide Web Conference*, December, 1995. [http://www.w3.org/Conferences/WWW4/].

Hankin, S., J. Callahan, and J. Sirott, 2001: The Live Access Server and OPeNDAP: Web visualization and data fusion for distributed holdings. *17th Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, Amer. Meteor. Soc., Albuquerque, NM. 380-383.

Rutledge, G. K., J. Alpert, W. Ebuisaki , 2005: NOMADS A Climate and Weather Model Archive at the National Oceanic and Atmospheric Administration. Accepted for publication *Bulletin Amer. Metero. 2005.*

WMO, 2001: Manual on Codes, FM 92- XII Part I BUFR (General Regularly- Distributed Information) in Binary Form, Vol. 1.2, 306.

WMO, 2001: Manual on Codes, FM 92-XII Part II GRIB (General Regularly- Distributed Information) in Binary Form, Vol. 1.2, 306.