

P2.39 ARCHITECTURE OF MADIS DATA PROCESSING AND DISTRIBUTION AT FSL

Chris H. MacDermaid*, Robert C. Lipschutz*, Patrick Hildreth*,
Richard A. Ryan*, Amenda B. Stanley*, Michael F. Barth,
and Patricia A. Miller

NOAA Research—Forecast Systems Laboratory
Boulder, Colorado

*In collaboration with the Cooperative Institute for Research in the
Atmosphere (CIARA), Colorado State University, Fort Collins, Colorado

1. INTRODUCTION

The NOAA Forecast Systems Laboratory's (FSL) Meteorological Assimilation Data Ingest System (MADIS) (Miller, et al. 2005) supplies quality controlled (QC) observational data to a large and growing segment of the meteorological community, including operational National Weather Service modeling and forecasting users. The requirements for this system have dramatically expanded to accommodate the increasing volume and types of data handled, as well as the growing user base. Today, the system ingests and merges data from dozens of sources, and supplies the resultant quality controlled products to hundreds of users via several standard protocols. Both real-time and saved MADIS datasets are made available. While it needs to be highly reliable, the system must be sized for the CPU-intensive QC analyses, and must provide controlled access to the data so that proprietary datasets are only available to appropriate users. Methods have also been developed for keeping track of the large, diverse user community and consistently responding to specific user needs, for example, notification of significant events.

To ensure reliability of MADIS, the system runs within the FSL Central Facility, a controlled environment maintained by FSL's Information and Technology Services, which provides system administration, networking, configuration management, data acquisition/distribution, and monitoring support (Lipschutz and MacDermaid, 2005). MADIS comprises a distributed architecture for ingest, processing, and data distribution functions. In addition, in a recent architectural advance, the various functional hosts are arranged in pairs using High-Availability (HA) Linux to provide auto-

mated failover in case of system failure. Real-time data distribution methods include FTP, Local Data Manager (LDM), and the Web-based OPeNDAP (OPen source project for Network Data Access Protocol). Special consideration for protecting proprietary datasets has added an extra layer of complexity, but is well handled by the configuration methods implemented on the LDM and FTP servers.

In this paper, we review the requirements and architecture of the operational implementation of MADIS at FSL. We describe our use of High-Availability Linux toward achieving high system reliability, and consider our implementation of the data distribution methods. Finally, we also discuss our user support infrastructure.

2. SYSTEM REQUIREMENTS

MADIS has a number of requirements that the Central Facility needs to meet. Among the requirements are acquiring data from many sources, moving the data between processing nodes, and distributing to MADIS's many clients.

MADIS routinely acquires mesonet data from several dozen network providers representing over 14,000 stations. This translates into over 30,000 station reports each hour. Several systems running both inside and outside the firewall acquire the data. These data are sent to the Central Facility's Advanced Weather Interactive Processing System (AWIPS) data server for preprocessing and conversion to a common format, NetCDF. The NetCDF files are then transferred to the MADIS compute nodes. These distributions are accomplished using LDM, which writes the data to local disk for access by the MADIS applications.

MADIS also requires access to other surface and upper air data sets. These are acquired through many methods including FTP, LDM, dedicated networks, and satellites. The data are processed using ODS into NetCDF and made avail-

Corresponding author address: Chris MacDermaid,
NOAA/OAR/FSL R/FST, 325 Broadway, Boulder,
CO 80305-3328; e-mail:
Chris.MacDermaid@noaa.gov

able to MADIS through the Central Facility's NFS attached fileserver.

MADIS data files are made available to FSL users on the Central Facility's fileserver that is widely mounted via NFS throughout FSL. For distribution outside FSL, files are distributed to the Central Facility's LDM server for real-time distribution and to an FTP server and OPeNDAP server for real-time and archived distribution. Additionally, several files are routinely distributed to the National Weather Service Telecommunications Gateway (NWSTG) for distribution to the National Weather Service Offices for operational backup. In

addition, QC'ed mesonet data are transferred to the National Centers for Environmental Prediction (NCEP).

Finally, MADIS requires archiving of all observations. This archive is maintained on the Central Facility's Mass Store System (MSS) and on-line.

3. ARCHITECTURE

MADIS comprises a large set of machines for its numerous functions. As shown in Fig. 1, data are transported from input and preprocessing systems to the central compute systems and then out to other hosts for storage and distribution.

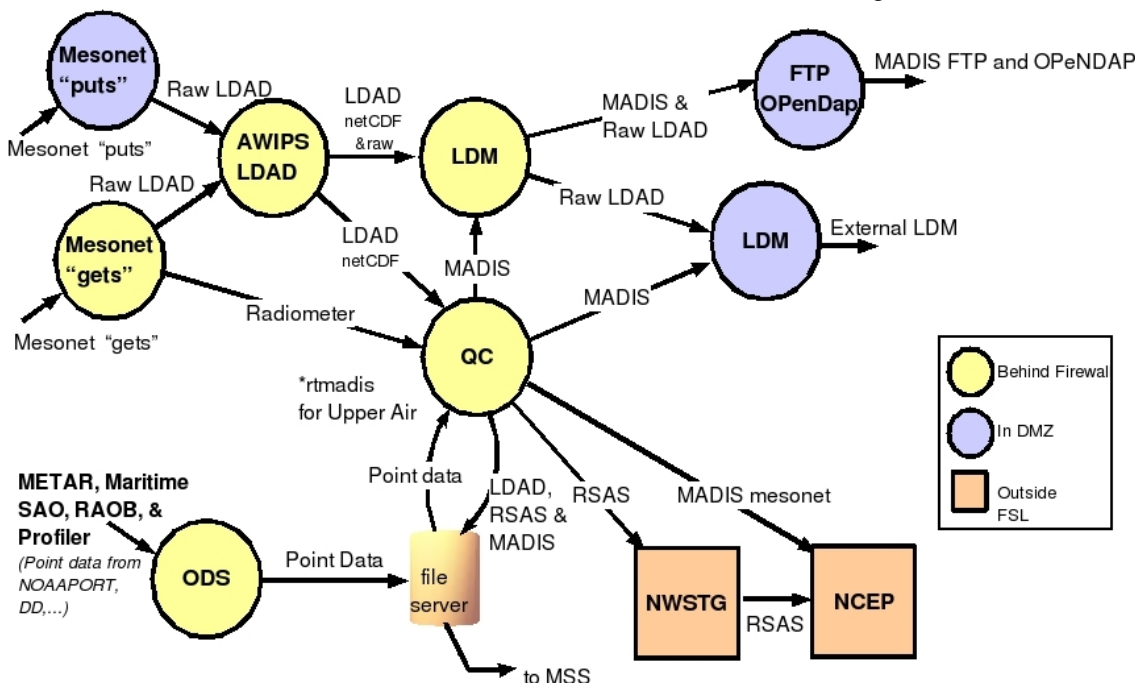


Fig. 1. MADIS Flow

The MADIS compute platforms comprise Intel-based servers running the RedHat Enterprise Linux operating system. Many of these computers are configured using High-Availability Linux clustering. This supports a two-node primary/secondary model with a serial heartbeat and IP address takeover. This is configured so that the secondary node is a hot backup for the primary node and in the event a hardware failure occurs, the secondary node automatically takes over processing. Configurations and system images are managed for these machines using the open source tools CVS and SystemImager (Lipschutz and MacDermaid, 2005).

Starting from the left side of Fig. 1, the mesonet data ingest function is shared across two hosts. On one host, scripts are configured to get data via FTP and HTTP from mesonet providers. The other host is running FTP and LDM servers to

which mesonet providers are enabled to put data.

These data are sent through LDM to the HA-configured AWIPS Local Data Acquisition and Dissemination (LDAD) server. Preprocessing to get these data into a comma-separated-value (CSV) format is performed. The AWIPS LDAD processing system then converts these into NetCDF files.

These data are then sent to the compute nodes where they are written to local disk for QC processing. Additional surface and upper air data are ingested into the Central Facility and processed using ODS. These additional data are made available to MADIS through the Central Facility's NFS attached file server.

The QC processing nodes are also configured in a HA cluster; the hardware for this system is the Dell 1750 dual 3.2 GHZ Xeon with 4 GB of RAM,

and a directly attached Dell PowerVault 220 storage device with 8 15,000 RPM 36 GB disks in a RAID 0+1 configuration. The RAID 0+1 configuration was chosen because it enables high I/O rates thanks to its multiple stripe segments. The local I/O required by this processing is currently about 100MB/S.

Since some MADIS data are considered proprietary by the provider, access to these data need to be controlled. To accommodate this, MADIS data are split into six different versions based on the level of access allowed by the data provider. A user accesses these data based on their access level. Distribution to FSL users is through the Central Facility's NFS mounted fileserver and MSS. Distribution to other clients is accomplished using FTP, LDM, and OPeNDAP. Multiple LDM servers serve this data out to about 100 users. FTP and OPeNDAP servers run on a shared host to distribute data to over 250 users.

4. USER SUPPORT

An important element of the real-time system for MADIS is the user support function. In particular, with the very large number of clients dependent on a reliable data stream, mechanisms for monitoring the system and notifying users about data outage events are critical. Users gain access to this data through a request through the MADIS web page, http://www-sdd.fsl.noaa.gov/MADIS/data_application.html. Central Facility operators monitor MADIS through the Facility Information and Control System (FICS). See Fig. 2 for a view of the FICS interface. FICS is a web-based monitor that is used to monitor all FSL's real-time Central Facility processing including MADIS ingest, processing and distribution. Through the use of green, yellow, and red icons, operators can quickly determine MADIS's state. Through links on the web page, operators are able to find troubleshooting information, log the troubleshooting steps taken, and send outage notices to users.

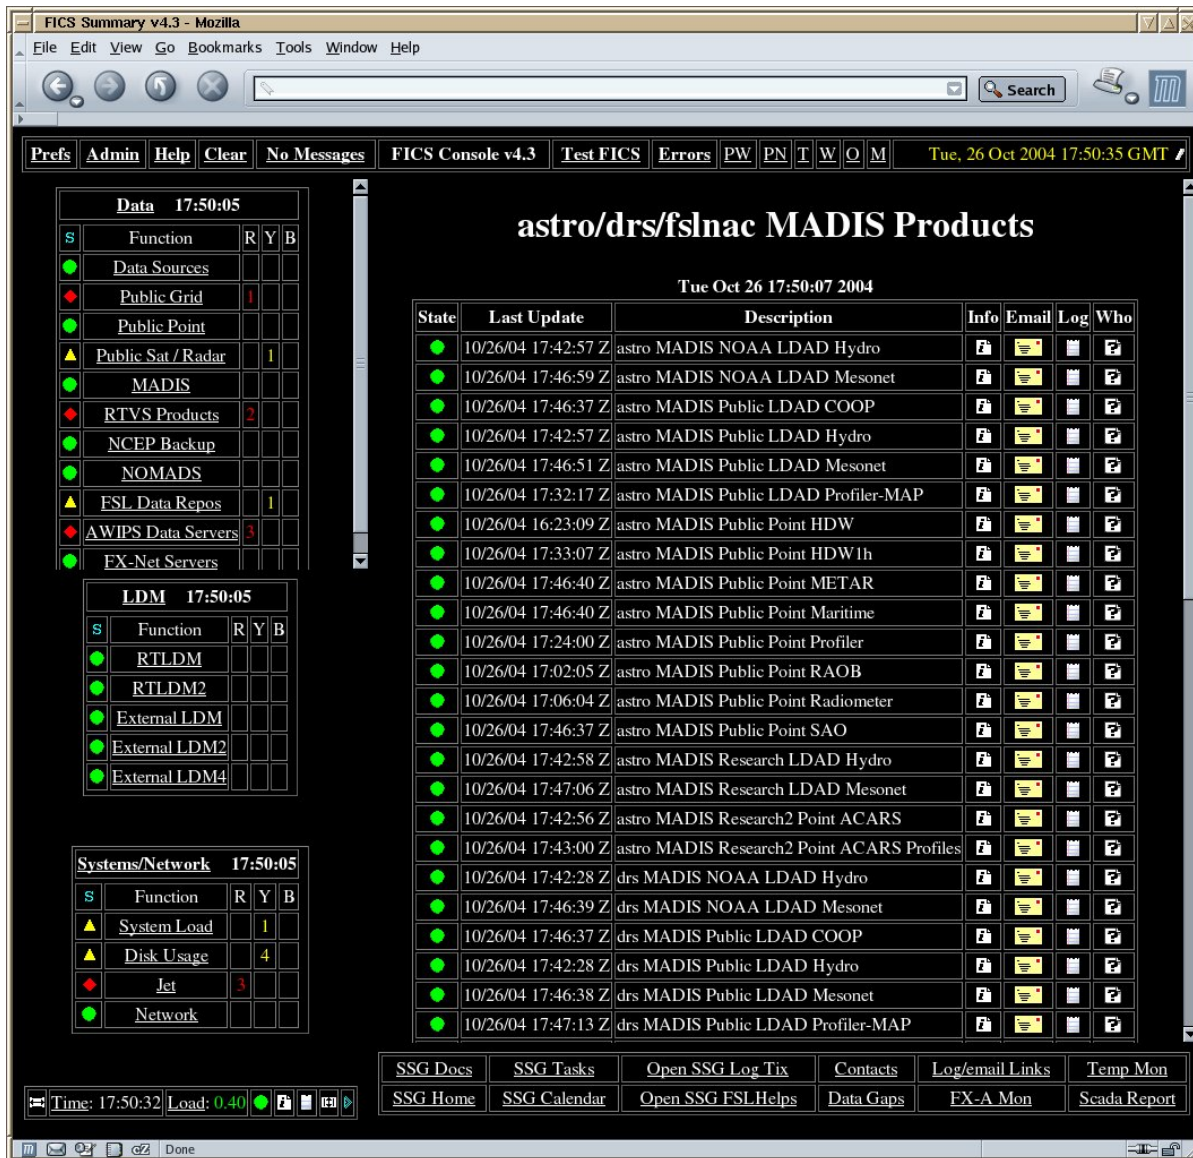


Fig. 2. A Typical FICS Display

In addition to FICS, computer resources are monitored by a customized event monitoring system based on Nagios (<http://www.nagios.org>). Network monitoring is done by the Central Facility's network administration using SolarWinds (<http://www.solarwinds.net>).

Another essential aspect of user support is tracking client contact information. Mailing lists of the MADIS clients' contact information and data types requested are maintained to notify

the users of data outages. Currently, contact and account information is maintained in text files in several different locations. This is being moved to an Open Lightweight Directory Access Protocol (OpenLDAP) directory approach to better maintain information. A directory is a database optimized for reading, browsing, and searching. The directory interface is shown in Fig. 3.

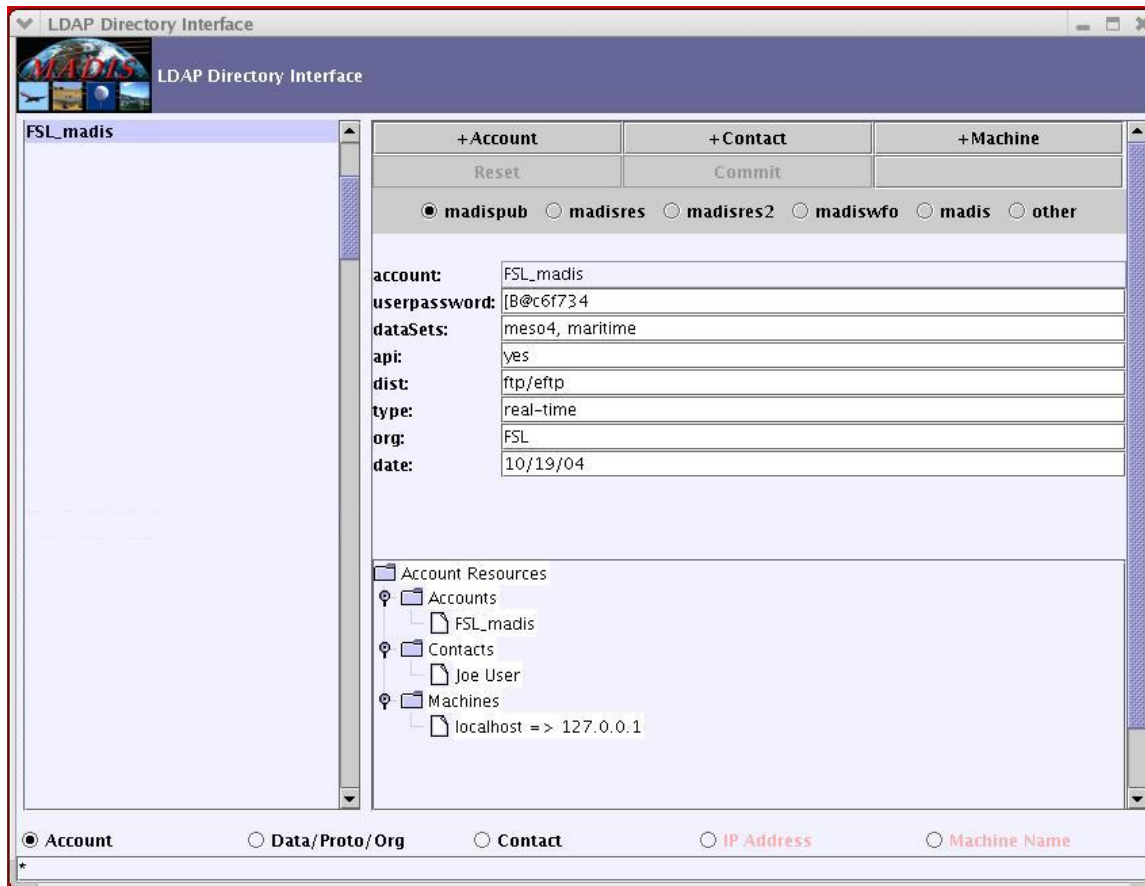


Fig. 3. User Directory Interface

5. CONCLUDING REMARKS

FSL's Central Facility will provide support for the ever-increasing MADIS data processing, ingest, and distribution. To increase reliability, the transition of the rest of MADIS to a HA-configuration will continue. Ongoing research into other clustering technologies such as Single System Image Clusters for Linux (OpenSSI) (<http://openssi.org>) will continue.

6. REFERENCES

- Miller P. A., M. F. Barth, and L. A. Benjamin, 2005: An update on MADIS observation ingest, integration, quality control, and distribution capabilities. Preprints, 21st International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, San Diego, CA. Amer. Meteor. Soc.
- Lipschutz R. C. and, C. H. MacDermaid, 2005: Recent advances in the FSL Central Facility Data Systems. Preprints, 21st International Conference on Interactive Information and Processing Systems for Meteorology,

Oceanography, and Hydrology, San Diego, CA. Amer. Meteor. Soc.

- Hamer, P., 2005: FSL Central Facility Data System concepts. Preprints, 21st International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, San Diego, CA. Amer. Meteor. Soc.