IDD Ingest Relay Replacement in Support of the THREDDS/RAMADDA Server System at Penn State

Final Report for UCAR Sub-Award No. Z11-90985
2011 Unidata Community Equipment Awards

Mr. Arthur Person, The Pennsylvania State University
Dr. Charles Pavloski, The Pennsylvania State University
30 April 2012

The Unidata Local Data Manager (LDM) package provides a stable mechanism for the dissemination of meteorological data from various sources to the greater Unidata community over the Unidata Internet Data Distribution (IDD) network. The Unidata THREDDS and RAMADDA platforms provide simple access methods to these data. The Department of Meteorology at The Pennsylvania State University (Penn State) has participated as an IDD level-two relay data distribution site since 1998, a top-level relay of CONDUIT data since 2010, and a THREDDS/RAMADDA data provider since 2011. Recently, the equipment providing our IDD relay service (3 servers with Intel Xeon 3.0 MHz processors/8 GB of RAM/10Krpm drives) had become frail being over 6 years old and of marginal capacity. Using funds provided by the Unidata Equipment Grants program, a trio of quad-processor Xeon E5606 2.13 GHz 64-bit machines running the LDM and Linux Virtual Server (LVS) in a virtual machine configuration is now providing IDD relay service originating from Penn State. In addition, the storage capacity of our THREDDS/RAMADDA data service was expanded to allow more historical Bufkit data to be available to the Unidata community.

Unidata's Role at Penn State and Elsewhere

Unidata's IDD feed and analysis/display software packages such as GEMPAK/NAWIPS, IDV, McIDAS-X and the coming AWIPS II are considered vital tools for research, instruction and outreach at Penn State. For example, the GEMPAK/NAWIPS suite of software allows our students to explore current and past weather scenarios as part of upper-level undergraduate meteorology courses. The GEMPAK/NAWIPS software is also used for the generation of graphics for the popular Penn State electronic map wall (e-Wall) available on the World Wide Web at:

http://www.meteo.psu.edu/ewall

The Unidata LDM feed and GEMPAK/NAWIPS software is also used on our 36-panel electronic display wall in the Penn State Weather Station:

http://ploneprod.met.psu.edu/facilities/teaching-facilities/mapwall.pdf/view

For research purposes, students, instructors and faculty use our real time and archived data from the IDD for a significant number of research initiatives, for example, accessing historical Bufkit data via the THREDDS data service (http://tds.meteo.psu.edu:8080) or RAMADDA repository (http://tds.meteo.psu.edu:8080/repository). As the scope of products and the volume of data increase via IDD, so does the importance of Unidata products to our educational, outreach and research programs.

Currently 13 non-PSU, non-Unidata sites connect to our IDD data relay including:

cascade.atmos.albany.edu, vortex.esc.brockport.edu, shu.calu.edu
wx.gmu.edu, flightrisk.meas.ncsu.edu, measwx.meas.ncsu.edu
lightning.msrc.sunysb.edu, thunder.msrc.sunysb.edu, coriolis.met.tamu.edu
idd.unl.edu, kepler.sca.uqam.ca, fujita.valpo.edu, omega.lsc.vsc.edu

System Configuration and Rationale

The upgraded hardware consists of 3 servers each configured with Intel Xeon 5606 quad-core processors running at 2.13 GHz, 24 GB of memory and a 450 GB 10Krpm hard drive. In order to leverage the capability of the servers, virtual host hypervisors were installed on each server and then virtual machines (VMs) were configured on each physical server. These were then used to configure a high-availability Linux Virtual Server (LVS) environment.

In the LVS environment, a director machine is used to field incoming LDM connection requests for data and these requests are then passed on to "real" servers which then handle the actual data transfer (see http://www.linuxvirtualserver.org/VS-DRouting.html). In order to increase reliability, two directors are implemented in a failover configuration so that if one director goes down, the other immediately takes over. In our virtual environment, the first and second physical servers are each configured with one director and two real servers. The third physical server is configured with just two real servers. The director VMs are configured with 1 GB of dedicated memory to prevent swapping and a dedicated processor. The two real servers each have 9.5 GB of memory dedicated to them and a dedicated processor. The remaining 4 GB of memory and the one processor is shared by the VMs as needed.

To populate data among the machines, one real server acts as the primary ingest machine obtaining all LDM ingest data from the Internet. The other real server instances obtain their feeds either from this primary ingest machine or from a peer on the same physical machine. Should the primary ingest machine malfunction, one real server is configured to act as a failover to replace the function of the primary ingest machine.

The trio of new servers configured for the high-availability Linux Virtual Server (LVS) environment and LDM Ingest Data at Penn State.

In operation, one director (the other director is not used unless the first one fails) fields all LDM connection requests and forwards them to a real server using an LVS load-balancing algorithm. As connections increase, each of the six virtual real servers on the three physical machines will be assigned a connection in a load-balanced manner thus spreading-out the load between all the virtual real servers. If the first physical server with the active director fails, the second physical server's standby director takes over. The net loss is two virtual real servers handling data, but four real servers remain on the two remaining physical servers. Thus high-availability is achieved.

In our previous environment without the use of virtualization, five machines were required to operate two directors and three real servers. In the new environment, only three machines are required thus saving costs for hardware. In addition, efficiencies found through the sharing of resources using an efficient hypervisor further leverage hardware resources.

Caveats and Future Expansion

In the current configuration, we are attempting to take the most advantage of the hardware resources available by configuring two directors and six real servers on three physical machines. The physical machines each have 24 GB of memory, so each real server (two per physical machine) has less than 12 GB of memory to work with. However, the LDM queue size must be about 14 GB or larger to maintain one hour of LDM data in the queue during peak data flow. Since the queue is larger than the available memory, some of it must be swapped out to disk. During normal operation this is okay, but if there were an extended network outage requiring downstream sites to reach back to the tail of the queue, a potential for thrashing could arise resulting in degraded behavior. If this were to actually occur in a problematic way, the configuration could be modified to allow for only one real server per physical machine thus making nearly 24 GB of memory available for use. Or, the systems could be upgraded by adding a second processor and an additional 24 GB of memory. Unless this problem actually occurs, however, we will maintain the current leveraged state and monitor operation.

About the Bufkit Historical Data

With the addition of 12 TB of storage to our existing THREDDS/RAMADDA data server, we have been able to retrieve and restore Bufkit data back through 2007. Our archives extend back through 2004 and the process of restoring these data to the THREDDS/RAMADA server continues.

More information on the Bufkit program can be found at http://wdtb.noaa.gov/tools/BUFKIT/