Upgrading the Penn State IDD Relay for the Next Generation
(Final Report for 2018-Unidata Equipment Awards)

Arthur Person and George Young
Department of Meteorology and Atmospheric Science
The Pennsylvania State University

Abstract

Weather information is the lifeblood of any organization tasked with making forecasts or analyzing weather threats. As technology advances, so does the quantity of weather information, and this requires hardware upgrades to handle additional data loads, additional users, and to maintain system reliability. The award for this project provided $11,910 to replace our previous generation of Unidata IDD relay hardware, which is now seven years old, with newer and expandable hardware to provide continued support for relaying the next generation of Unidata IDD weather information.

The Penn State IDD Configuration

The Penn State IDD relay consists of a three-server LVS direct routing cluster (refer to http://www.linuxvirtualserver.org/VS-DRouting.html). Using a custom configuration, two of the servers act as primary and backup "directors" while also functioning as "real servers" for data delivery. The third server acts only as a data delivery "real server". Data ingest is performed on the primary director system, while the other two servers feed from this ingest machine. If the primary director/ingest machine goes down, the cluster will automatically reconfigure to allow the secondary director/server to take over the director and ingest function. The third system will automatically switch ingesting to the secondary director/server when it detects the primary has failed. In the case where either the secondary director/server or the third server go down but the primary remains up, they simply drop out of the cluster and the cluster continues to function normally. In the unlikely case where two servers go down, the cluster will automatically reconfigure to the remaining server unless only the third server remains. In that case, the third server can be manually reconfigured to handle the full load.

The selection of a three-node versus a two-node or four-or-more node cluster, in our environment, seems to be optimal. With three nodes, the loss of one system still leaves two systems to handle the work load and to provide redundancy should a second system fail. In addition, with three nodes, one node can be removed for testing or upgrades without overloading the cluster function since two systems remain. Although four nodes would provide additional capacity and redundancy, the capacity is currently not needed for our workload, and the likelihood of three systems failing at the same time is very small. In our three-node configuration, the worst case would be the failure of two systems. In that case, a single system can handle the current workload although data latencies might increase since 128 GB of memory may be insufficient to handle both the data queue and all the processes associated with our downstream feeds.

The retired systems replaced by this project ran dual Xeon E5606 2.13 ghz 4-core processors with 72 GB of memory and could not be expanded beyond 96 GB. With a full complement of Unidata feeds being ingested at an average rate of 63 GB/hour (1.5 TB/day) with a peak of 86 GB/hour, these systems were not able to store a minimum of 1 hour of data in their queues. The new systems utilize a single Xeon 4110 2.10 ghz 8-core processor with 128 GB of memory, and easily handle the full complement of IDD data maintaining at least 1 hour of data in the queues. Included with the new systems is a 240 GB SSD for the operating system, a 2 TB hard drive for the LDM queue, and a 10 Gbps optical network interconnect. The systems are expandable to two processors and up to 1 TB of memory.

Penn State IDD Data

Penn State ingests and relays all of the Unidata feeds and is a top-level distributer of NCEP CONDUIT data. Currently, the total amount of data ingested per day is 1.5 TB. The largest contributors to these data feeds are gridded model data (36%), GOES-16 and GOES-17 data (33%), and Nexrad level II data (14%). The remainder is comprised of MRMS, Nexrad level III, and other text and image products. Timely relaying of these data are possible thanks to Penn State's 100 Gbps connection to the Chicago, IL Internet 2 router, and soon, a direct 100 Gbps connection to the Ashburn, VA Internet 2 router will be installed. This will make Penn State's proximity to the Washington D.C. area and NCEP much closer network-wise.

Network capacity is important since data sources continue to grow and be added to the Unidata IDD offerings. Over the last two years, GOES-16 and GOES-17 data have been added to the IDD offerings, and the new GFS FV3 model will soon replace the current GFS spectral model. The GFS FV3 will include pressure level grib files that are 50 percent larger, and forecasts for 240 to 384 hours that are at 3-hour intervals instead of 12-hour intervals. Additional data sources like these will continue to place increased demands on network bandwidth and relay capacity, and the upgraded Penn State relay is poised to handle these demands.

The Penn State IDD Service

The Penn State IDD relay service currently serves about 23 downstream sites. Most are educational sites including George Mason, MIT, NACSE, NC State, Ohio State, Oswego, Rutgers, Texas A&M, U. at Albany, U. Kentucky, U. Michigan, U. Nebraska, U. Northern Iowa, U. Oklahoma, U. Quebec, U. South Alabama and Valparaiso, but also includes a couple of government-related sites (CPTEC in Brazil, and NOAA) and a couple of miscellaneous sites.

Internally, Penn State is also a heavy consumer of Unidata IDD data with the primary outlet being the Penn State electronic map wall, or e-Wall (http://www.meteo.psu.edu/ewall). Products available on this site are produced from Unidata IDD products ranging from text-based observational data to GOES satellite imagery and model grib data. The e-Wall web site is used within the Penn State Department of Meteorology and Atmospheric Science in weather forecasting classes, in weather discussion seminars, in the production of the Weather World TV

show (http://weatherworld.psu.edu/), and by the Campus Weather Service (https://campusweatherservice.com/). It is also used broadly by educational, government, industry and private users worldwide.

Conclusion

The upgraded Unidata IDD relay cluster at Penn State is now installed and operational, and is providing reliable, timely data delivery to the Unidata community. As additional users and data expand beyond the capacity of the upgraded system, it can be expanded by adding processors and memory at minimal cost, thus providing continued service for years to come.