

2021 Unidata Community Equipment Award

Final Report

The Next Generation of NEXLAB

—

Server Upgrade for College of DuPage's Meteorology Lab

Name: Paul L. Sirvatka
Title: Professor of Meteorology
Institution Department of Earth Science, College of DuPage

Name: Ron Stenz
Title: Associate Professor of Meteorology
Institution Department of Earth Science, College of DuPage

Name: Evan Anderson
Title: Web Developer & Lab Manager
Institution Department of Earth Science, College of DuPage

Name: Mike Zuranski
Title: Meteorology Support Analyst
Institution Information Technology, College of DuPage

Name: Gilbert Sebenste
Title: Meteorology Support Analyst
Institution Information Technology, College of DuPage

A. Project Summary

The College of DuPage received a 2021 Unidata Community Equipment Award of \$20,000 to purchase a replacement server for our GOES ingest and visualization processing. Given the previous server has reached the end of its warranty period, and the fact that our satellite imagery made with that server is our most popular product suite, it was imperative that we replace that equipment before any hardware issues were encountered. The difference in hardware also allows for continued development and product offerings; two additional Geostationary Lightning Mapper (GLM) products have been made available since the new server went into production.

Based on our experience with satellite processing on the previous server, this time we decided to use Proxmox to turn the new server into a virtual machine host instead of installing everything on the bare metal. This allows us to capitalize on the hardware as much as possible, without having one set of processing adversely impact another.

A few challenges were seen with both the acquisition of the server as well as getting it into production, not to mention the untimely transition of key staff members, so it took longer than originally expected to put everything to work. Now that the new server has been driving our satellite processing for several months we are able to take a look back and review how the project went, and what doors this equipment has opened for the COD Meteorology program.

B. Equipment Purchased

This award has allowed us to purchase a HPE DL380 Gen10 rack server (detailed specifications shown in Table 1). The server was installed in May of 2022 and was put into production that August. It is located in the college's climate-controlled server room, which uses redundant generator-backed power, and is staffed & monitored 24/7 by the college's IT department.



Figure 1. The new server (top) and the 2017 Unidata Equipment Award server (bottom)

Table 1. System Specifications	
HPE DL380 Gen10 2U Rack Server	
Chassis	
HD Bays	8 – 2.5 inch
Power Supply	2 – 1600W Redundant Power
System Configuration	
CPU	2 - Intel Xeon-Gold 6248R 3.0GHz 24-core (96 threads total)
RAM	8 - 32GB DDR4-2933 (256GB)
SSD	2 - HPE 800GB SAS SSD, ZFS RAID-1
HDD	4 - HPE 1.2TB SAS 10K, RAID-5

C. Project Description

Since we began visualizing GOES-16 data in 2017, satellite imagery has remained the most viewed set of products on our site by far. We currently offer GOES East and West imagery from all 16 ABI bands, eight multispectral (RGB) products, and 14 GOES derived overlays including four GLM products across over a hundred CONUS, Full-Disk and Mesoscale domains. Most of this work was developed using the previous server, though this upgrade has allowed us to visualize two additional GLM products: Minimum Flash Area (MFA) and Total Optical Energy (TOE). More RGB products will also be possible pending a new

C-1. Using Proxmox as a Virtual Machine Host

Unlike our previous server where the operating system was installed directly onto the machine, this time we decided to use Proxmox to turn this server into a VM host. The decision to do this came about because we noticed satellite processing at the scale we do doesn't utilize the full power of the CPUs, but adding additional work could still cause problems if one set doesn't finish before the next set comes in, thus creating a snowball effect that could disrupt the visualizations if not addressed. Even though this hardware includes faster processors and a greater number of threads in total, being able to separate processing into different VMs allows for much greater flexibility.

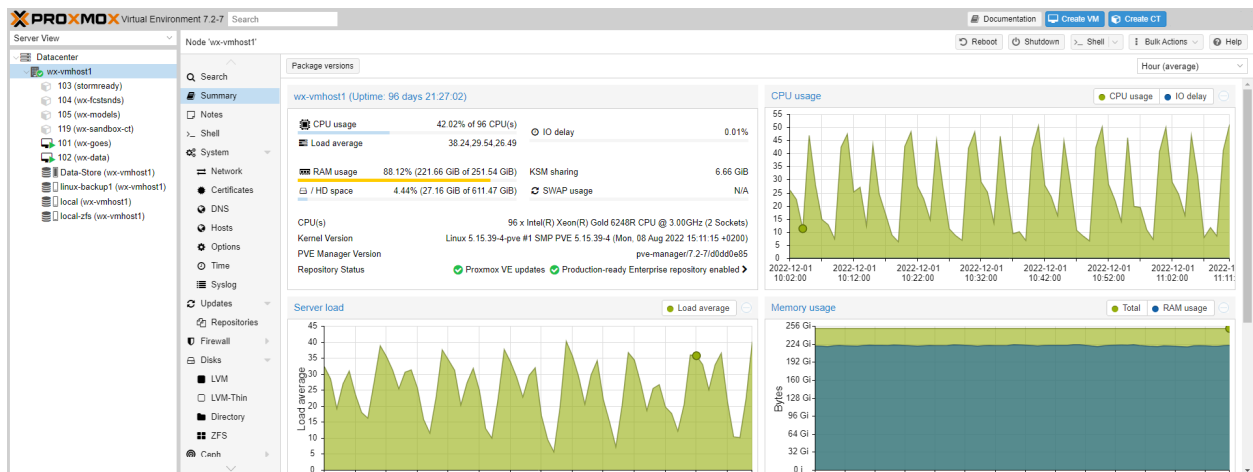


Figure 2. Proxmox interface showing the summary page for the new server

There are presently two VMs running on the new server, one for satellite processing and another for image storage. Previously the images were stored on the processing server, which was problematic for the site & users if the server needed to be rebooted. By having the imagery stored separately, the processing VM can go down for maintenance without impacts to the end users.

Other planned VMs include one for forecast model visualizations, one for forecast sounding generation and one for StormReady / weather alert scripting. While this hasn't been done operationally, during testing we visualized the GEFS model on a VM alongside the satellite processing VM with no apparent impacts to either. This VM host also offers opportunities for testing and development, as being able to spin up VMs of different operating systems and specifications is trivial.

C-2. RAID Settings and Performance

This server has two sets of drives: two SSDs and four larger HDDs. The two SSDs make up the "OS Partition" volume. Defined during the Proxmox installation, the SSDs are configured for ZFS RAID-1 to allow for balanced performance as well as the protections offered by the ZFS file system. It is this volume where all VMs and containers are defined & stored, and the space available for these is ample.

The four larger HDDs are combined to make a single storage volume for all output imagery. Initially, we created the storage volume within Proxmox as a RAIDZ-1 volume, the ZFS equivalent to hardware RAID-5. Before long, however, the CPU overhead for the ZFS layer became noticeable, and I/O wait on the processing VM was too much for the amount of processing we do to be sustainable. Because the storage volume contains nothing but imagery, the benefits gained by using ZFS here weren't enough to offset the CPU cost, so we remade the storage volume a hardware-defined RAID-5. Proxmox sees this as a single drive, making its inclusion into Proxmox that much simpler, and we retain the protections offered by a battery-backed hardware RAID-5 array.

C-3 Challenges and Delays

When we were approved for the Equipment Award towards the end of 2021, we were hoping to purchase and receive the server within a few months, as that had been the typical turn-around for this type of server purchase before. However, due to processor and other part shortages, as well as delays during the purchasing process itself, the server didn't arrive at the college until January 2022.

This server, as well as a number of others recently purchased by the college, requires 220-volt outlets instead of 110. While we knew this well before the server arrived at the college, there were sizable delays with getting the 220v outlets installed and the server wasn't powered in the rack until May, approximately four months after the server was delivered.

While Proxmox is free to use, it's strongly recommended to use their Enterprise Repository for any machines in production use and that requires a Proxmox subscription. We decided not to put the new server into production until we had access to the Enterprise Repository, which took another month or so to purchase and receive. Proxmox needs to be installed on the hardware to be used before a subscription can be purchased, as a server ID number is required and that is presented through the Proxmox interface. While we were able to configure and test most things during this time, it did present an additional delay for getting everything into production.

The last unplanned challenge we faced was the untimely departure of our IT support analyst, Mike Zuranski, who left the college to become (coincidentally) Unidata's new Data Engineer. Our position has since been staffed by Gilbert Sebenste, but the transition has delayed further development on the server and additional VMs. The expectation is development will resume as Mr. Sebenste settles into his new role.

D. Looking Ahead

With the additional headroom this server now provides, we are able to offer additional products that we weren't able to before. We have already added the MFA and TOE GLM products, and we also plan on expanding the GLM coverage to hopefully all GOES-East and GOES-West sectors. The upcoming Unidata v2021 release of McIDAS-X adds a feature for GAMMA correction, a parameter required by numerous RGB products. While the software update would have come anyway, we should be able to add some of these RGB products to our suite that we wouldn't have been able to otherwise. And with the ability to spin up VMs for other purposes, we now have the potential to balance processing workloads from other machines. We are excited to be able to offer these new products and explore new possibilities thanks to this Unidata Equipment Award.