Cluster Administration Report

Spring 2022

Written by

Andrew Capalbo

Supervised by

Dr. Marcus Hohlmann

Assisted by

Zackery Wihela

Objectives:

The main goal of my work in the HEP Grid Cluster was to repair the system to the point where it can act as a stand alone cluster with CMS software. This breaks down into 3 areas that most of my attention was focused on this semester: the Tripp-lite UPS, the CE (or compute element), and the SE (or storage element).

Introduction:

The USCMS Florida Tech Grid Cluster is comprised of 6 main parts: The CE, The SE, NAS-0, NAS-1 , compute nodes, and 3 UPS's. The CE, or compute element, acts as the front-end of the Cluster, and is where most of my work was done. The SE, or storage element, manages data transfers from CERN, and is currently not in use. The two NAS's act as the main storage device for the system, with NAS-0 housing the home directories for the Cluster. The 20 compute nodes contained in the cluster are responsible for running incoming batch jobs from the CE. The UPS's, or uninterpretable power supplies, supply power to cluster and contain a large battery pool to draw from in case of power loss.

In the beginning of this semester, 3 immediate problems presented themselves. The first was the bottom Trippe-lite UPS, which has not been able to power cycle after several months of non-use. The second was the SE, which also has not been able to power cycle after several months. The final and most pressing was the CE, which was not able to connect to the network. Without a functioning front-end, there is no cluster. This became my main focus over the course of the semester.

Problem 1: Tripp-lite UPS

The bottom rack UPS has been issue in the Cluster since the previous admin left in Spring of 2021. In the fall, Zack and I were able to find a corroded wire and burnt out wire cap inside the UPS. After replacing the wire and cap, the front LEDs lit up and there was a beeping but would only happen when connected to battery power, and would immediately shut off when supplied with AC.

Moving on to Spring 2022. Through further investigation we found that the fans were extremely rusted. We formed the hypothesis that it was booting off due to a safety feature with the faulty fans. After ordering replacements and installing them, they would begin spinning when connected to battery but would shut off with the LEDs when connected to AC. The issue has yet to be resolved and was put on the back burner over higher priority issues.

There is diagnostic software available for tripp-lite systems, and while informative on the system itself, did not shine any light on the issue regarding the boot cycle. Several self tests were ran that returned with no errors.

While still nonoperational, after some calculation we found that the two functioning UPS's output enough power for all the current cluster components(excluding the SE), so this is not detrimental to the overall function of the Cluster.

Problem 2: SE (Storage Element)

In the beginning of this semester a check was done on the SE and we found that it was unable to power on after pressing the power button. After pressing the button, the power supply fans would kick on and the LED panel would illuminate the HDD indictor light amber. As the HDD

indicator typically means a drive is being read, we took out all the drives to test if they had gone bad. With 4 drives in total, only one was able to read immediately by my laptop, and the others returned “wrong fs type, bad option, bad superblock on /dev/sda, missing codepage or helper program, or other error.” This is a fairly vague error message, so a self test was ran on the drives and showed no issues.

While checking on the drives, we opened the lid off the SE and saw that there were no RAM cards installed. After removing the RAM from NAS-0 and inserting it into the SE, there was no change. It is unclear what happened, and is likely that SE has integrated RAM rather than physical cards.

As the Cluster is now aimed to be a standalone system, the SE was put on the back burner and unfortunately no more progress was made during the semester. Once the other components are functioning correctly, I hope to come back to the SE and resolve this issue.

Problem 3: CE (Compute Element) and nodes

In fall of 2021, the CE began having issues connecting to the network. After several months of hitting a brick wall (into Spring of 2022), we decided the best course of action was to start from scratch and reinstall the operating system on the CE. We followed the installation guide written in 2019. After a successful installation, the system was unable to complete the initial boot. In the loading screen, hitting escape shows the operations being performed, and we noted that most of them resulted from the system being unable to mount the partitions made during installation.

After several weeks of research and hypothesizing, we found that the raid card within the CE was degraded. After removing it and connecting the drives directly to the motherboard, the system successfully booted. A side-effect of this is that there is now a desktop environment installed. After running **htop** to see how much system resources were being used to display the desktop, we decided it was inconsequential and left it installed. This has made networking much easier as the configuration files are easily located in settings.

Once the system was configured to the correct settings, we began the unification process with the nodes. This process begins with **insert-ethers**, which is how the CE initially recognizes the nodes, NAS’s, and other accessories. The program works by first creating a DHCP server that the node connects to, which gives it its IP address in the local network. The next step requires a TFTP, or trivial file transfer protocol, server. The CE, after pulling in the node, supplies a kickstart file to the node which in turn installs rocks onto the node.

Unfortunately, it appears there is an issue with the TFTP server. It is unclear if this is due to the CE or the node, and after trying to resolve it, we found that the kickstart file is simply just the base rocks configuration that we set up for the CE. Using this information, we attempted an installation on the node and successfully installed the operating system.

After successfully booting in, we found that the network was no longer working. This is odd, as the network must be operational to install rocks as it pulls the rolls from the rocks website. Once again, as the node gets it network through the CE, it is unclear if this is an issue with the node or the CE. By using **tcpdump**, I was able to confirm that the CE was receiving the pings sent out by the node, but was either not sending a reply or the node was not picking it up. When running the program on the node, no ping was detected from the CE, leading me to believe that it is an issue

with the node. More research and troubleshooting must be completed to definitively tell where the issue originates.

Conclusion:

While we were unable to complete the objectives given this semester, we have made great strides and have repaired the system to a usable point. CMS software such as HTCondor is installed and operational on the CE, and while there is an issue with the network on the node, it has this same software installed and (partially) operational. I am confident that coming into the next semester we will be able to build the Cluster back to its former glory.

The remaining tasks to complete the cluster are to finish the node unification and unify the NAS's as well. After this, the cluster will be completed as a working high throughput computer cluster. Additional tasks are to resolve the power cycling issues with SE and Tripp-lite UPS and re-upload the cluster website. These will be the first steps when coming into the Fall 2022 semester.