VICTR: VIRTUALIZED INFRASTRUCTURE WITH CYBER AND TESTBED RESOURCES

A. Lohr, J. Thebarge, M. Kwong, M. Yun, M. Ellabidy

Wentworth Institute of Technology (UNITED STATES)

Abstract

Virtualization software and platforms have become increasingly popular over the past few decades. Virtualization allows companies and universities to save money and resources while maximizing the potential of their equipment to host all the services and functionalities they need. Virtualized Infrastructure with Cyber and Testbed Resources (VICTR) is an ESXi based virtualization platform that provides Wentworth Institute of Technology (WIT) students and faculty with access to a virtual machine environment that is hosted on WIT's internal network resources. By using VICTR, students and faculty can access multiple virtual machines to conduct labs, testing, and research without sacrificing the computing resources on their personal computers. In this paper, we present how we designed VICTR and provide the virtual switch and network configurations in detail. We demonstrate the cyber range by displaying communication between Metasploitable virtual machines and Kali virtual machines. Furthermore, we configured 21 virtual machines to test resource load with simultaneous connections to prove the capability of handling typical class environment at WIT. From the testing results, we showed our system can support a typical classroom environment. Whether it be used for a system administration course consisting of 20 or so CentOS virtual machines or utilization of the cyber range practicing exploits or penetration techniques on Metasploitable virtual machines, VITCR can aide in the application of the materials learned in the classroom environment.

Keywords: Virtualized Infrastructure, Virtual Machines, Testbed Platform.

1 INTRODUCTION

Virtualization, especially virtual machines have become a key component in businesses and collegiate universities. The concept of virtual machines has been around since the early 1960s but were widely adopted about forty years later in the early 2000s. Virtualization became so popular due to the ability it gave companies and institutions to partition servers and run legacy applications on multiple operating system types and version [1]. This advancement enabled organizations to lower their overall spending costs on purchasing the equipment, set up, cooling, and maintenance [2].

Virtual machines are made by installing software, known as a hypervisor, which separates the physical resources into pieces that the virtual environments can use. There are two types of hypervisors – Type 1, "Bare Metal" and Type 2, "Hosted". "Bare metal" means that the hypervisor acts as a lightweight operating system that runs directly onto the physical host. Type 2 runs similarly to a computer application; it is a software layer on top of an operating system. The most popular type of hypervisor is Type 1, which allows companies and institutions to meet their needs for data-center's [3].

Using virtual machines allow collegiate universities to create essentially a computer within a computer. Our project initiated to have hundreds of virtual machines hosted on our school's internal network at Wentworth Institute of Technology (WIT). The goal of this project is to provide students and faculty a resource to complete coursework or individual research without using their personal computing resources. Personal computers and laptops have limitations on the amount of processing power and research subjects such as cyber-attacks and defences. We developed a virtualization platform, Virtualized Infrastructure with Cyber and Testbed Resources (VICTR), that provides WIT students and faculty with access to a virtual machine environment. By using VICTR, students and faculty can access multiple virtual machines to conduct labs, testing, and research without sacrificing the computing resources on their personal computers. VICTR is built with 3 Dell PowerEdge R720 servers, 3 Dell EMC Isilon Storage Arrays, one Cisco 2960G Switch, and one Mellanox IS5024 InfiniBand Switch. The Isilon systems and the Dell servers were externally attached to the Cisco switch which was attached to WIT campus network. In addition, the Isilon systems were internally networked together through the Mellanox InfiniBand switch.

The rest of this paper is organized as the following: Section 2 presents the design requirements and decisions including the pre-build requirements for VICTR and the decisions that were considered.

Section 3 provides the physical network topology and virtual network design of VICTR including Isilon storage cluster design and ESXi server design. Section 4 depicts the testing results and our findings. Finally, Section 6 presents our conclusions and final thoughts.

2 DESIGN REQUIREMENTS AND DECISIONS

We designed a cluster of servers that support upwards of several hundred virtual machines. Our system can be used for several different academic courses, a cyber range, and a testbed for research and classroom activities. In addition to this, there was a need to implement automation for rapid deployment and configuration of virtual machines.

The existing hardware that we had to work with was: 11 36-bay Isilon IQ108NL storage array with 3TB hard drives in each bay, 7 Dell PowerEdge R720 servers, 2-3 Dell PowerEdge servers with various models, one Cisco 2960G 24-port switch, and one Mellanox IS5024 InfiniBand switch.

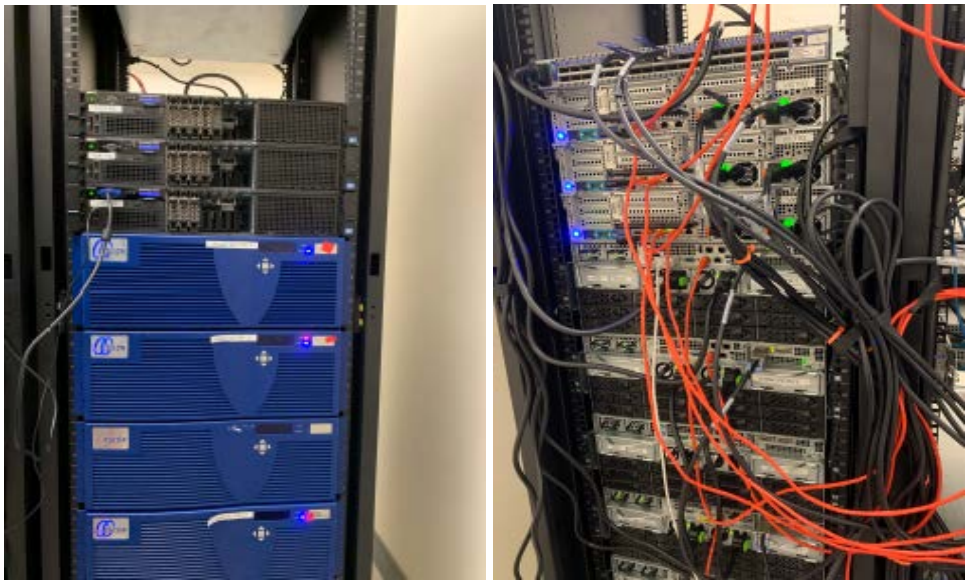

Figure 1. (Left) Front View and (Right) Rear View of Server Cabinet.

There were many design decisions that were made and taken into consideration. The first step of this project was to evaluate the existing hardware and to see if any parts needed to be ordered. Out of the 7 Dell PowerEdge servers that were available to us, there was only 800GB RAM total. Out of the 800GB, there was 40 Dual In-Line Memory Module (DIMMs) at 16GB and then 20 DIMMS at 8GB. However, the servers did not come with any internal storage at all. This presented a question of how the servers would be able to boot ESXi (Bare Metal hypervisor). We decided to use three 32GB USB drives to have the servers externally boot ESXi. The next step was to evaluate the Isilon storage systems; all the systems came with 36 3TB hard drives. The next limitations to evaluate were the power supplies available and available rack space. In our server room there was an existing PDU that we could plug 1 rack mounted APC into. In addition to the PDU, there was also a 30-AMP overhead power supply that had 2 open outlets if we needed. In terms of the rack limitations, we only had access to one 42-U server cabinet; the Dell PowerEdge servers were 2U, the Isilon storage bays were 4U, the Maddox Switch was 1U. Taking all these factors into account we decided to utilize 3-Isilon storage arrays (yields up to 300TB of storage; 12U of rack space), 3 Dell PowerEdge R720's (yields up to 800GB memory; 6U of Rack Space) , one Mellanox IS5024 InfiniBand switch (1U of rack space), one Cisco 2960G 24-port switch (2U of rack space), one rack mounted PDU that will be attached to the existing APC. All these physical components fit into the 42U server cabinet with room to spare for future expansion as shown in Fig. 1. Only approximately 21U of space was utilized from the server cabinet, about ½ of the total capacity. The three Dell PowerEdge servers are on the top, then the Isilon storage arrays are at the bottom of the rack. With already having an existing networking architecture, the red cabling in Fig. 1, is the Ethernet connection to the Cisco switch and WIT's internal network. For the software component, we determined that ESXi and the vCenter Suite was the correct way to go considering all the features and potentials [4]. As students of Wentworth, we were granted access to the Enterprise Plus version of vCenter and vSphere for free.

3 VICTR

3.1 Physical Network Design

3.1.1 Physical Network Topology

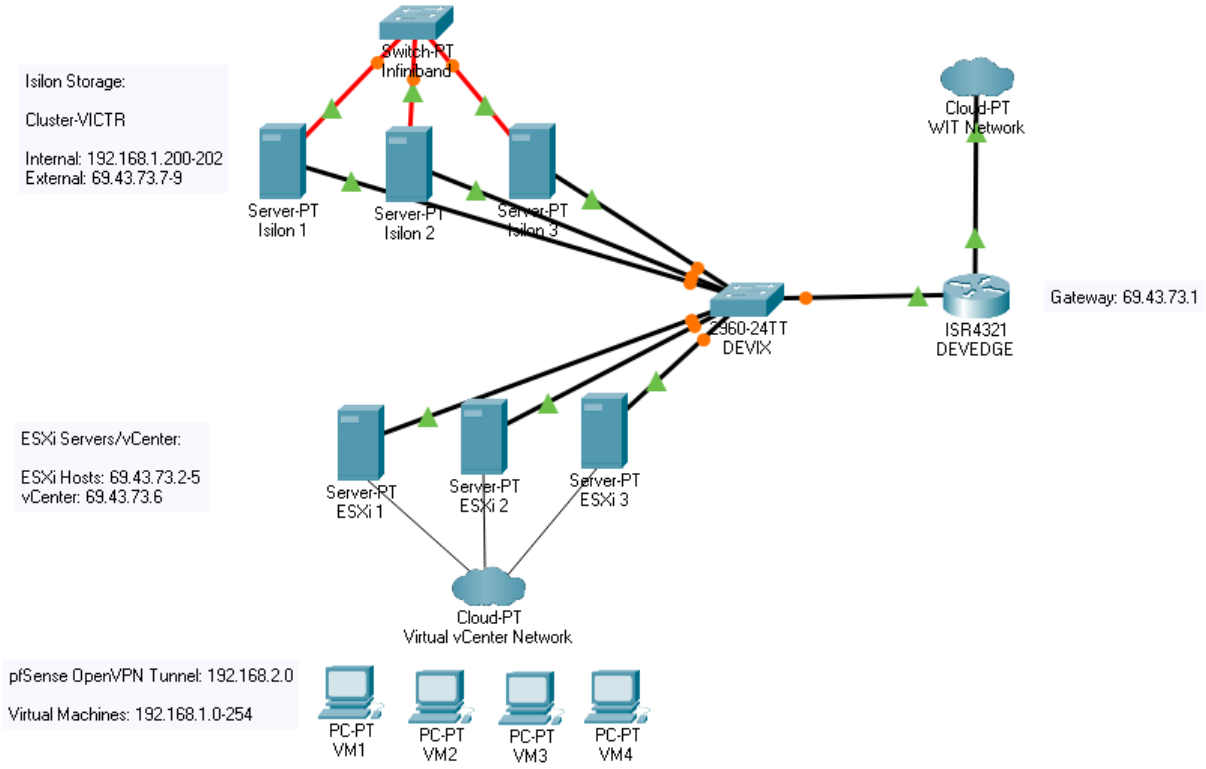

Figure 2. Physical Network Diagram.

The physical network topology of VICTR is shown in Fig. 2. On the right-hand side, the existing network includes the DEVIX (Cisco 2960G) Switch and the DEVEDGE Router; both devices were already configured to work inside the WIT Internal Network. The subnet available for use was 69.43.73.0/24 with a gateway IP address of 69.43.73.1. The DEVIX switch was adapted a little bit to accommodate our new infrastructure. A virtual local area network (VLAN) was created on the DEVIX switch to isolate our new infrastructure. Ports 1-19 were added to VLAN 70.

Attached to the DEVIX switch the three Dell EMC Isilon storage arrays and the three Dell PowerEdge R720 servers. The Isilon storage arrays were assigned static IP addresses of 69.43.73.7-9 and there was a cluster management address of 69.43.73.14. The Isilon arrays were also assigned internal addresses of 192.168.1.200-202.These internal addresses only help the Isilon arrays form and maintain the cluster. The Mellanox InfiniBand switch use these internal addresses for the purpose of clustering. The Dell PowerEdge servers was assigned static IP addresses of 69.43.73.2-4. The vCenter virtual machine, which is the central management for the ESXi Cluster, was assigned a static IP address of 69.43.73.6.

3.1.2 Storage Cluster Design

Fig. 3 shows out storage cluster overview from OneFS. The Isilon storage arrays were assigned internal IP addresses of 192.168.1.200-202 and external addresses of 69.43.73.7-9. The internal addresses were strictly for the creation of the storage cluster. The cluster was created with the metric shown in Table 1.

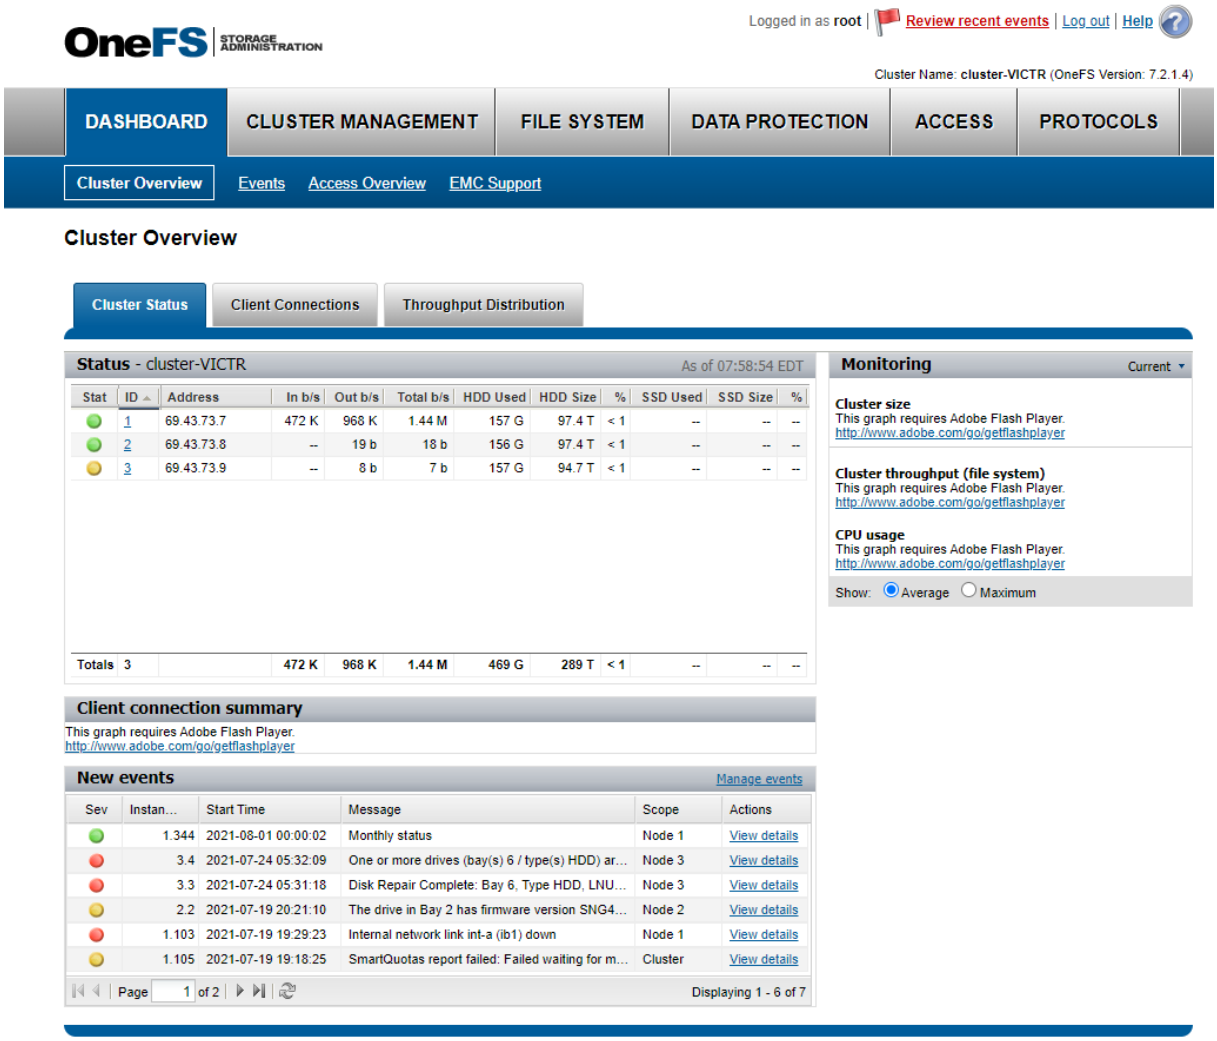

Figure 3. OneFS Cluster Overview.

Table 1. Isilon Cluster Metrics.

Configuration Parameters	Values
Configuration Name	Cluster-VICTR
Cluster Encoding	UTF-8
Interface Int-A Netmask (Internal)	255.255.255.0
Interface Int-A IP Range (Internal)	192.168.1.200-202
Interface Ext-1 Netmask (External)	255.255.255.0
MTU	1500
Interface Ext-1 IP Pool	69.43.73.7-13
Ext-1 Default Gateway	69.43.73.1
SmartConnect Zone name	mgmtVICTR
SmartConnect Service IP	69.43.73.14
DNS Servers	8.8.8.8
Cluster Time and Date	EST (Eastern Time Zone)
Cluster Join Mode	Default (Manual)

Once the cluster was created, confirmation of the cluster configuration and health was confirmed by going to the Web GUI at the address of 69.43.73.7 using the port 8080. Through this GUI we were able to confirm that the total available storage is about 300 TB. As shown in Fig 4, we can see that only 0.2% of the overall capacity of the cluster is being used for this project.

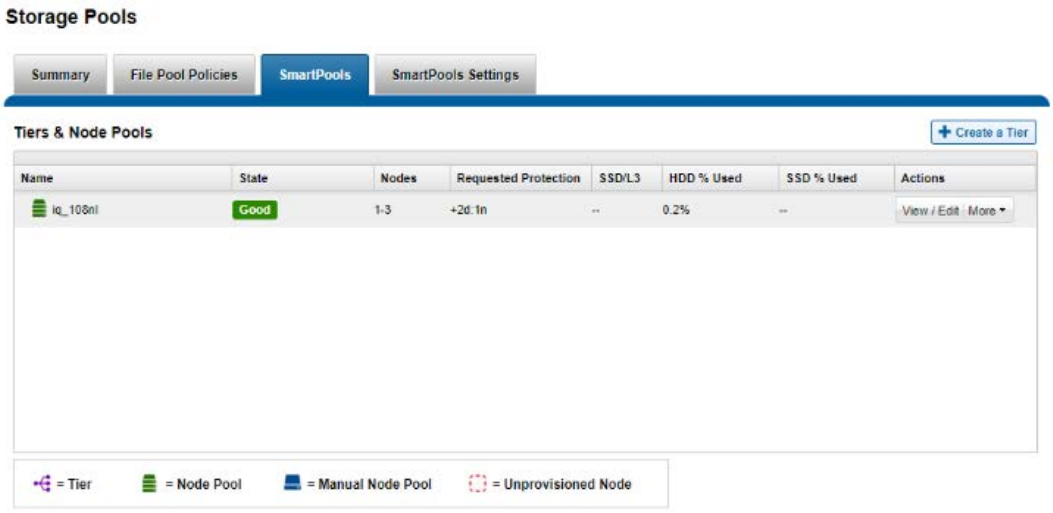

Figure 4. OneFS Storage Pools.

3.1.3 ESXi Server Design

As previously mentioned in Section II, the physical servers did not have any internal storage. Also, there was about 800 GB of RAM to split between three servers out of 7 Dell PowerEdge servers. Server 1 was given 384 GB of RAM (24 DIMMs at 16 GB), Server 2 was given 256 GB (16 DIMMs at 16 GB), and Server 3 was given 160GB (20 DIMMs at 8 GB).

Each Server booted from an external jump drive that has ESXi loaded allowing the servers to boot from the jump drives. The servers are given static IP addresses in the range of 69.43.73.2-4 with a gateway of 69.43.73.1 and DNS server of 8.8.8.8. Once the servers finished installing ESXi, the configuration was confirmed through the online GUI portal by navigating to each IP address in an internet browser, shown in Fig. 5. The next step was to install vCenter. VCenter is the central management software for individual ESXi hosts. The vCenter virtual machine was given a static IP address of 69.43.73.6. Once the installation was complete, we were able to login to the GUI through a web browser at the address of 69.43.73.6. The configuration of vSphere will be discussed in Section 3.2.

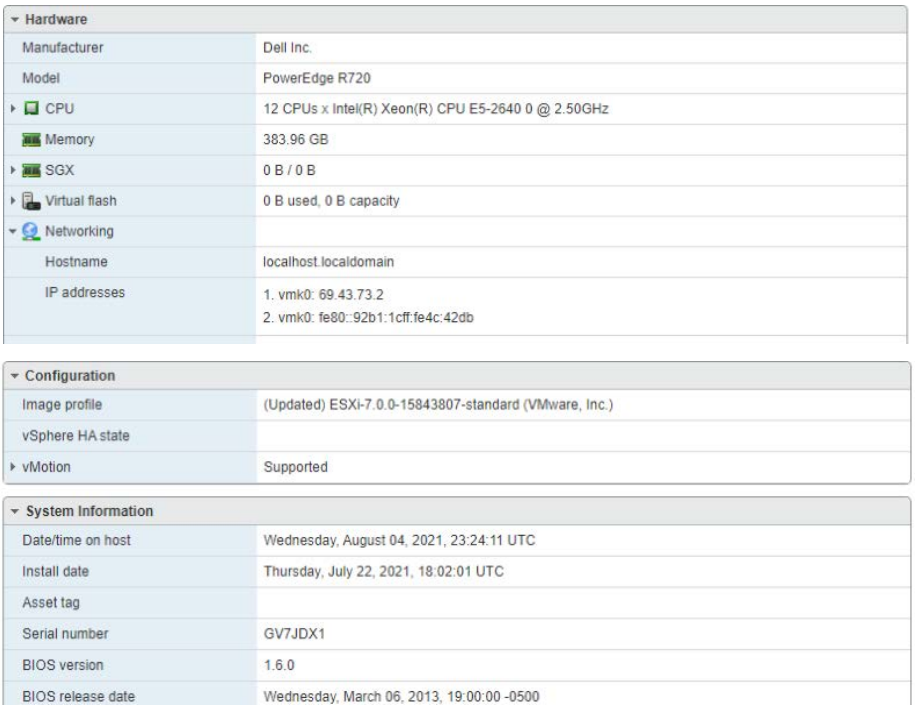

Figure 5. ESXi Server Configuration

3.2 Virtual Network Design

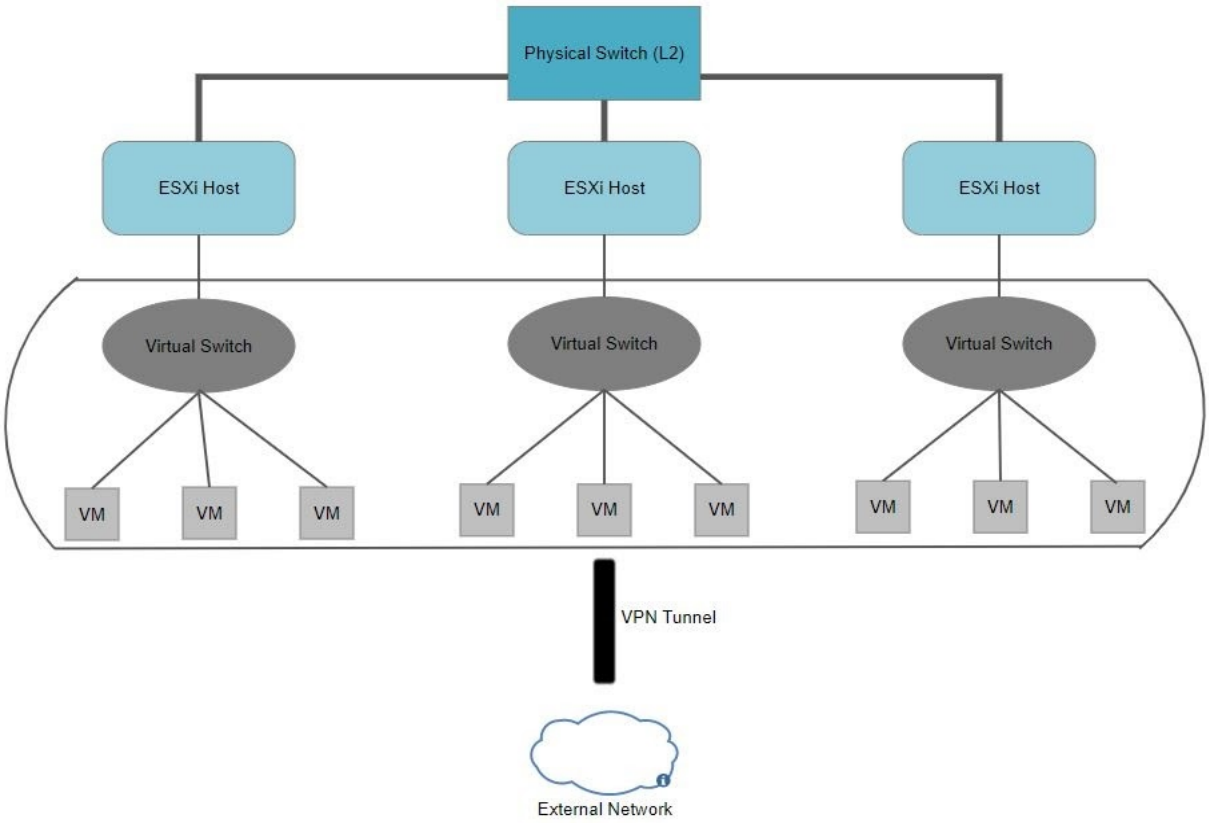

Figure 6. Virtual Network Diagram.

As shown in Fig. 6, VICTR has 3 ESXi hosts, and they contain 3 separate entities. One host is used for academic research. This contains a virtual machine running CentOS version 8 and the purpose of this machine is to provide an environment for testing and research. The second host is used for a system administration course for junior students in WIT. This used a test bed for students and professors for hands-on activities and projects in the system administration course. Finally, the third host is used for the VICTR cyber range. By have a controlled and interactive test bed, our cybersecurity undergraduates can learn how to detect and mitigate cyber-attacks.

We implemented virtual switches to connect these virtual machines (VMs) to the ESXi server host. Virtual switches are using the same protocols that are used over physical switches. This also helped with system performance because the system can pay the lowest possible cost in complexity and demand [5]. This all leads to the virtual private network (VPN) tunnel. The VPN was set up through using a Pfsense VM installed on the ESXi host. The VMs can be accessible to outside computers through a Windows remote desktop protocol (RDP) session with this VPN configuration. A prospective user needs to download the OpenVPN software and import our corresponding key to initiate a session to the VPN tunnel. Once the session is validated, then the user can access the VM through RDP.

4 TESTING

For initial testing, our focus was to determine if our VICTR could perform up to 20 connections and active use of the VMs [6]. This is representative of a typical class size at WIT. We connected 20 sessions through our openVPN and generate network traffic internal and external with 80-byte Internet Control Message Protocol (ICMP) packets being sent from each VM.

Fig. 7 and Fig. 8 shows the network usage during this testing phase. Fig. 7 is the network usage before the test is initiated and Fig. 8 shows the network utilization during the test. From the graphs we can see that we capped out at about 20KBps, and for our dual gigabit connections connected to both ESXI server this was no problem. We also checked our total resource utilization during this testing session and found that as far as computational power goes, there is enough to support the 20 simultaneous connections.

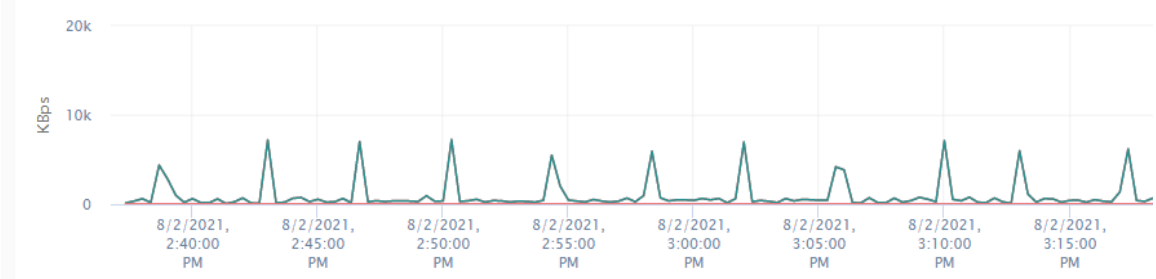

Figure 7. Network Usage Graph Before Testing.

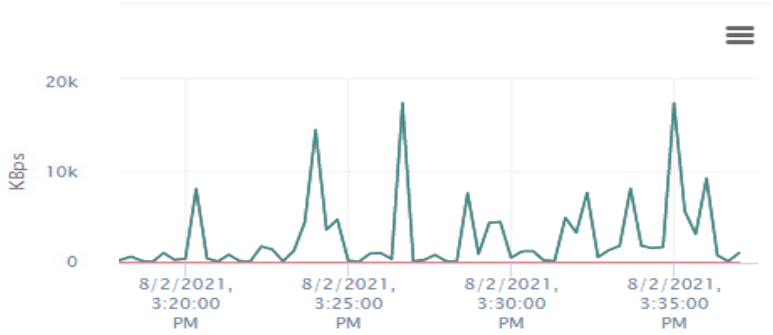

Figure 8. Network Usage Graph During Testing.

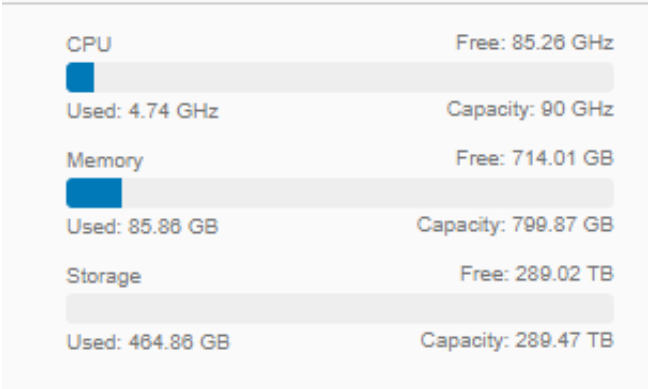

Figure 9. Research Utilization Graph.

As shown in Fig. 9, the test is only using about ten percent of the usable memory, and 5 percent of the CPU. We also tested the cyber range to ensure that penetration testing techniques could be run successfully on the other VMs, and we were able to get positive results back. For example, we simulated a lab from one of WIT's academic courses. In this type of lab students would be asked to conduct exploits from a Metasploitable virtual machine to another host. We successfully used a Metasploitable virtual machine and were able to gain access to another virtual machine hosted on the vCenter cluster. To simulate a System Administration Course, we performed some basic Linux commands and scripts to make sure that students could create and execute scripts. The System Administration Course only relies on one CentOS VM per student to conduct their lab environment, so their course would be the easiest to implement on the VICTR system. Lastly, we used basic network connectivity tests such as ping, nmap, traceroute to ensure that our network was fully functional, and the virtual machines were completely operational.

From the testing results, we concluded that our system can support a typical classroom environment. Whether it be used for a system administration course consisting of 20 or so CentOS VMs or utilization of the cyber range practicing exploits or penetration techniques on Metasploitable VMs, we have no doubt VITCR can aide in the application of the materials learned in the classroom environment.

5 CONCLUSIONS AND FUTURE WORK

VICTR is an ESXi based virtualization platform that provides WIT students and faculty with access to a virtual machine environment. By using VICTR, students and faculty can access multiple virtual machines to conduct labs, testing, and research. In this paper, we successfully presented how we designed VICTR and provided the virtual switch and network configurations in detail. We demonstrated the capabilities of VICTR in the classroom environment through the multiple testing.

VICTR is considered usable, but not production ready. We want to see VM management and deployment come into play prior to undergraduate classes utilizing this resource. The setup is completely manual outside of having common VM templates created, and for the team and or administrator continuing the deployment of this system they will need a means for VM deployment to assist in the ease of managing the system altogether. Also, there is about 300TB worth of storage readily available and capable of supporting any active use, but the storage arrays are not set up for failover. Prior to VICTR being ready for production, we plan to add two more storage arrays into the cluster so that Redundant Array of Independent Disks (RAID) 5 and smart partitioning can be enabled to have structure and control over how the file system is being used.

However, even though VICTR is not production ready, it is still a usable system. We created a working networking infrastructure that has great potential. During this project 21 virtual machines were created on top of the ESXi hosts and lab environments were simulated. The storage capabilities alone, will allow students to gain access to a resource that they would never have on their own. VICTR will allow students and faculty to access multiple virtual machines to conduct academic coursework and research. With a few tweaks this system that were outlined in this paper, VICTR can be a powerful environment.

REFERENCES

[1] C. Chen, Y. Du, S. Chen and W. Wang, "Partitioning and Placing Virtual Machine Clusters on Cloud Environment," 2018 1st International Cognitive Cities Conference (IC3), 2018, pp. 268-270, doi: 10.1109/IC3.2018.000-2

[2] J. E. Smith and Ravi Nair, "The architecture of virtual machines," in *Computer*, vol. 38, no. 5, pp. 32-38, May 2005.

[3] R. Morabito, J. Kjällman and M. Komu, "Hypervisors vs. Lightweight Virtualization: A Performance Comparison," *2015 IEEE International Conference on Cloud Engineering*, pp. 386-393, 2015

[4] N. Suradkar and S. Lomte, "VMware ESXi: Virtual Web Server performance evaluation with weighttp Benchmark," 2020 IEEE International Conference on Advent Trends in Multidisciplinary Research and Innovation (ICATMRI), pp. 1-4, 2020.

[5] VMware, “VMware Virtual Networking Concepts,”, VMware Information Guide, 2007. https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/virtual_networking_concepts.pdf

[6] I. Gordin, "Virtualizing real servers with unsupported OS by VMware vCenter Converter v6," 2015 14th RoEduNet International Conference - Networking in Education and Research (RoEduNet NER) , pp. 127-131, 2015.