CMHS High-Availability Supercomputing Environment

"The Hive"

Carlos Gross Jones

1 Introduction

1.1 Educational Objectives

- Virtualization - probably the most important recent trend in IT, the practice of running servers as virtual machines under a hypervisor, rather than directly on the physical hosts, allows more robust, flexible, and scalable allocation of computing resources.
- Server hardware
- Web development
- Parallel processing & clustering - One of the greatest challenges in largescale computation and "big data" analysis is developing algorithms and protocols to ensure that computers can effectively work in parallel.
- Networking - While the physical networking required to implement the cluster itself is relatively simple, the ability to create large and complex systems of virtual hosts, virtual routers, and virtual networks means that network administration (e.g., CCNA training) can be taught with far less resources than are required to buy actual networking hardware for a teaching lab.

2 Software

The compute nodes operate as a virtualization cluster. Each host runs ESXi, and is managed by vCenter Server. Thus, rather than running a VM on a particular host, each node contributes its resources to the common pool. VM images are stored on the fileserver, and instantiated by vCenter with whatever computational resources the application requires.

Each of the three storage servers runs on bare metal. The fileserver acts as bulk data storage, for media, documents, VM images, etc. The database server runs MySQL server, and potentially other types of database server in the future. Both the fileserver and database server replicate their data onto the backup server, which is designed to be a high-capacity, high-latency data store.

Virtualization allows maximum utilization of computational resources. For example, for the purpose of rendering animations, a VM disk image could be made with Windows and the render engine. When the time comes, the hypervisor could be instructed to instantiate a number of rendering VMs, each cloned from the stored disk image. These VMs would be spread among the physical hosts, and could use up to all of the resources available in the cluster. Once the render is complete, the rendering VMs are destroyed, and their resources are re-allocated to lower-priority VMs. The process of creating virtual machines takes seconds, as opposed to shutting down a server and booting into a different OS. Thus, the same cluster can be used for rendering, intensive simulation, etc., without any need to reconfigure or reboot physical machines.

In addition, a virtualized environment allows better allocation and segregation of development server. For students learning web development, systems administration, and so on, it is desirable to give each student their own server, so that mistakes are contained. However, it would be both impractical and wasteful to allocate one physical server per student. However, it is very easy to create a practically unlimited number of small VMs, each with minimal vCPU and memory resources needed for learning. Furthermore, VM disk images could be created with commonly used configurations (e.g., Django, LAMP stack, DokuWiki) already installed. When a student wants to experiment with a technology, a VM can be quickly and easily cloned from the image and used for development.

3 Hardware

While finding a sponsor willing to fun a large server cluster for a high school would be preferable, it is expected that the cluster will start in a fairly minimal configuration (see below). This is acceptable, and in

fact intrinsically supported by the architecture. By using the cluster nodes as virtual machine hosts, rather than running each as a standalone server, capacity can be easily and seamlessly added. For example, if another server is purchased and added to the cluster, instead of migrating a resource-intensive application to the new server, the hypervisor merely has additional computational resources to allocate to the task.

In fact, the BOM presented for the minimal configuration is probably conservative. It would probably be fairly easy to find a company willing to donate a used server rack to a high school, and a used or refurbished UPS could be found that is nonetheless up to the task.

Table 1: Minimum Configuration Hardware

Item	**Description**	**Price Per**	**Quantity**	**Line**
Router	Cisco 2610 (Already owned by CHMS)	$0.00	1	$0.00
Switch	Cisco Catalyst 2950 (Already owned by CHMS)	$0.00	1	$0.00
Compute Node	Dell Poweredge 1950 (Est. preowned)	$100.00	2	$200.00
Storage Node	Dell Poweredge 2950 (Already owned by Carlos)	$0.00	1	$0.00
UPS	CyberPower OR700LCDRM1U	$195.95	1	$195.95
Rack	42 U Open-Frame 4 Post Rack (Est. preowned)	$200.00	1	$200.00
			SUM	$595.95

Table 2: Maximum Configuration Hardware

Item	**Description**	**Price Per**	**Quantity**	**Line**
Router	Cisco 2801 (Est. preowned)	$100.00	1	$100.00
Switch	Cisco Catalyst 3750G-48TS (Est. preowned)	$200.00	1	$200.00
Patch Panel	ICC Cat5e Feedthrough Patch Panel	$195.51	1	$195.51
Environment Monitor			1	$0.00
Compute Node	Dell PowerEdge R610 (dual HC) (Est. preowned)	$600.00	17	$10,200.00
Console Drawer	Dell PowerEdge Rack Console (Est. preowned)	$100.00	1	$100.00
KVM Switch	16 Port USB KVM Switch	$80.00	2	$160.00
Storage Node	Dell PowerEdge R710 (dual QC) (Est. preowned)	$600.00	3	$1,800.00
UPS	APC Smart-UPS SMT1000RM2U	$464.23	2	$928.46
Rack	42U Enclosed Rack (Est. preowned)	$500.00	1	$500.00
			SUM	$14,183.97

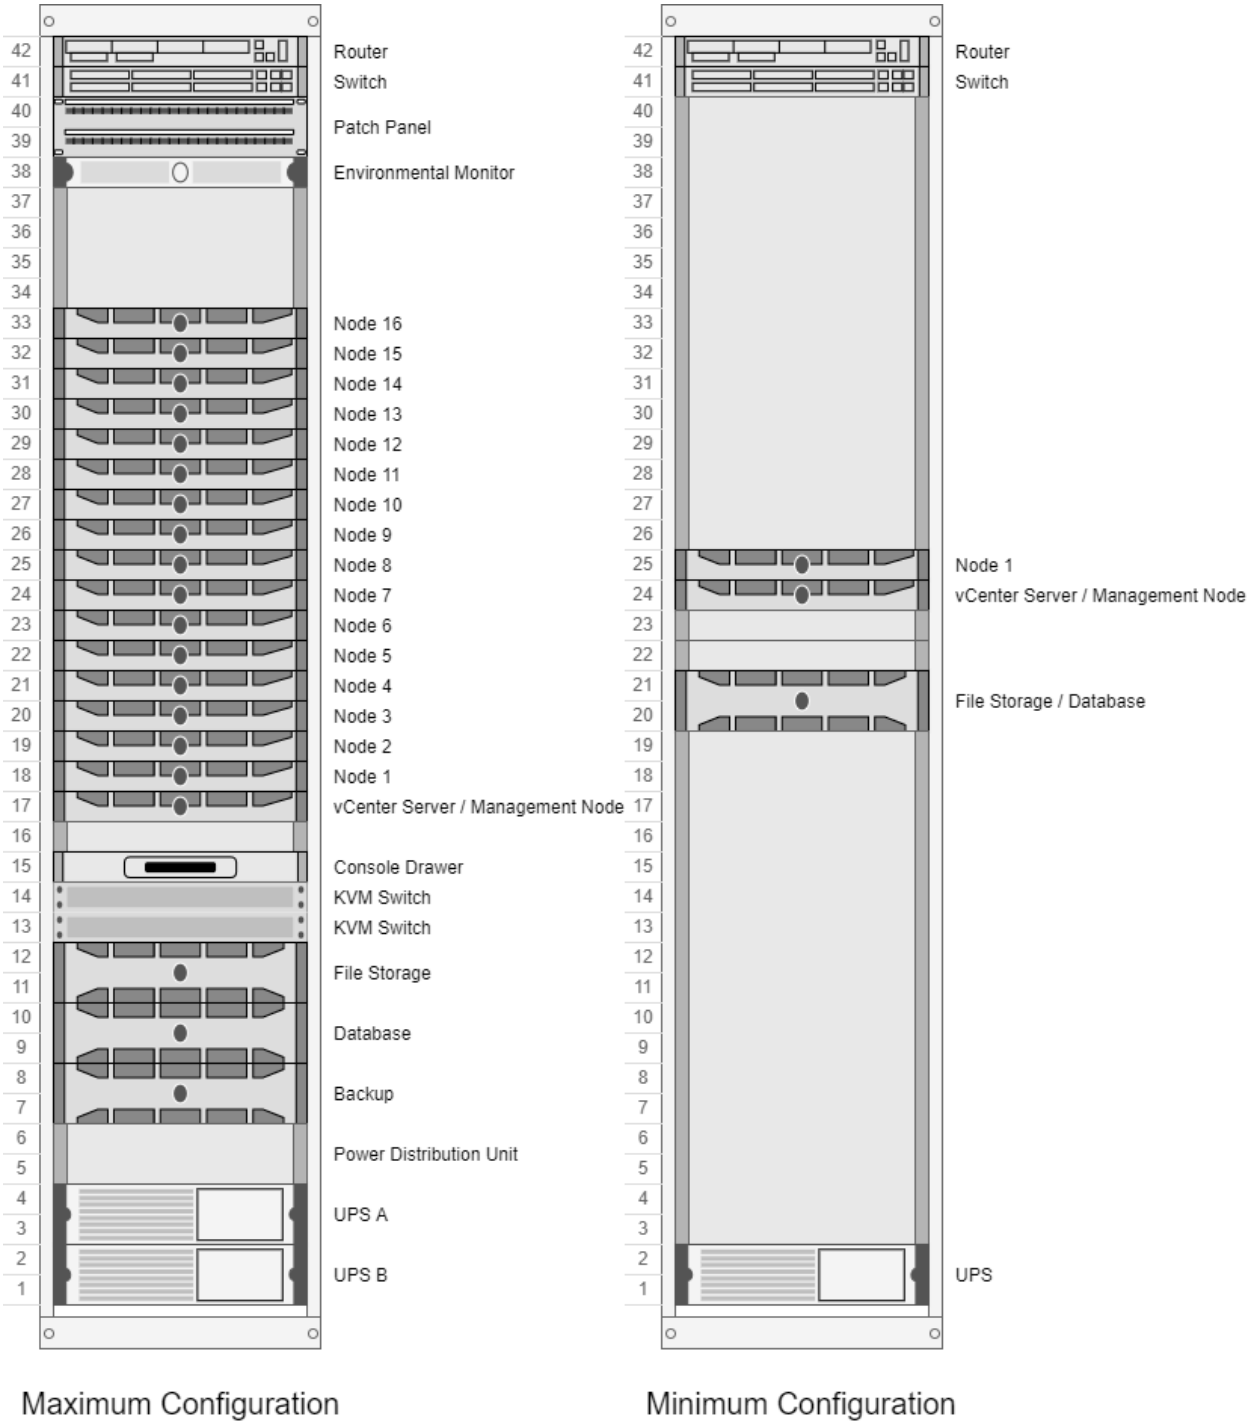

Figure 1: Minimum and maximum (based on power requirements) cluster configurations