MIGRATING CONTROL SERVERS AND APPLICATIONS TO VIRTUAL MACHINES

G. Wright, C. Angel, C. Finlay, E. Matias, CLS, Saskatoon, Canada

Abstract

The Canadian Light Source has ever-increasing numbers of distributed Controls applications that can run on generic networked computers. This has led to additional servers when segregation by operating system, by LAN, or by system load has been required. Additional concerns of maintenance of older computer hardware, driver support for older O/S's on new hardware, and growing Virtual LAN issues have led to an adoption of moving applications to Virtual Machines (VM). Our implementation using VMware Infrastructure provides a high reliability environment, with centralized monitoring of performance and simplified expandability. The distributed 'average' reliability hardware has been replaced by a single high-reliability system with built-in redundancies. A new virtual machine can be started in a matter of minutes from an existing pre-configured template, and can be joined to up to four VLANs with a simple software configuration.

INTRODUCTION

In 2008, CLS was evaluating upgrade paths for a number of controls computer systems. This provided an opportunity to consider options other than simple one-for-one replacements. The computers to be replaced were standard rack-mount systems providing EPICS IOC functions, shared standard development environments, or network services. Of the systems providing IOC functions, some required dedicated hardware, some simply used network communications with their devices.

VM TECHNOLOGY DETAILS

- A Virtual System Server can provide one or more Virtual Machines
- Each VM has its own virtual disk partition. To the server, this is simply a large file
- Each VM can be migrated from one server to another
- Each VM provides support for running a Guest O/S
- The Guest O/S doesn't need to know the details of the server, or even that it is running on a VM server!
- A Guest O/S can be installed on a Virtual Machine from installation CD's in the same manner as a regular computer O/S installation occurs. Additionally, an installation can be performed by *cloning* an existing VM and Guest O/S.

THE PROBLEM

The problem being faced at the CLS is the same problem faced at many facilities: there are always more computer systems being added running mission-critical functions. As the number of computers increase, the probability of a failure within the group increases.

EVALUATION CRITERIA

The three major hardware options considered were:

- Purchase newer versions of the existing hardware
- Purchase a blade server
- Purchase a Virtual Machine server

The Issues

There were a number of specialty systems in the systems list. For instance, one of the computers was configured as an EPICS Gateway between the many VLANs set up at the CLS. This system required a network port that allowed IEEE 802.1Q communications. Another computer provided archiving for Channel Access data, and required a reasonably large amount of immediately connected data storage. Some systems were required to continue running RedHat 7.2 as their O/S.

As always, the dollar cost of any solution is important.

The Analysis - Update Existing Systems

Simply replacing existing systems would do little to reduce the overall system failure risk. Failure of a system would require manual effort to configure an available standby and become fully operational again.

Additional computers would require more rack space, more power, more network ports, more available fail-over systems, and more system administration. Newer computers would take more work to have driver support available for any legacy OS.

But, this would be simple from the point of view of continuing to do the same thing.

The Analysis - Replacing existing systems with a blade server

The major advantage of blade servers is reduced rack space and reduced cost for a fully populated blade server as compared to the equivalent number of rack-mount computers.

The blade computers would still require dedicated network drops (although possibly fewer than independent computers), and available hardware for fail-over. Also, they would be unable to communicate using any hardware other than Ethernet (PCI slots are not provided).

The Analysis – Replacing existing systems with a Virtual Machine Environment

A VM server system has a number of advantages over standard rack-mount systems and blade systems:

- Hardware failover for the server allows failover for every Virtual Machine
- Shared CPU resources provides easier matching of computing power to requirements
- Shared network resources reduce network cabling and network infrastructure requirements
- Shared power resources reduce power cabling
- The cost per virtual machine can be lower, as the resources can be allocated more efficiently and effectively

There are disadvantages. The most obvious are

- Any system requiring additional PCI bus communications cannot be easily virtualized
- The initial server cost (hardware, software, and training) is high; a reasonable starting point requires that a number of systems must be ready to migrate to the VM server
- The VM server is one more piece of software to support
- Multiple VM's requiring lots of resources may be in immediate contention with each other
- Directly connected disk storage is more expensive than off-the-shelf SATA disks

HARDWARE AND SOFTWARE

The complete server system was purchased from Dell Canada, and included training and support.

Hardware

The hardware system consists of:

- 2 Dell 2950 Servers
  - 2 4-core 2.66 GHz processors
  - 32 Gig memory
- iSCSI SAN with 4 Terabytes storage

Software

The VM software consists of:

- VMware ESX
- VMware Infrastructure Client

THE MIGRATION PROCESS

Moving applications from a physical machine to a virtual machine occurred during a planned facility outage in March and April of 2009. Although an application is available to take an image of an existing O/S installation from a hard disk and make it a Guest O/S on a VM Server, it was instead decided to use the opportunity to install a newer O/S version which could then be configured and cloned. The migration process was then reduced to configuring the destination system to run the applications that had previously been run on the physical machine.

VIRTUAL SYSTEMS

Figure 1 lists the VM's that are currently used. Of these, 5 were migrated from existing systems, and the rest were new installations. The systems listed under *Base* have been configured either for testing or for easy cloning. The systems listed under *Production* are those systems in active use. The division into two sections allows easy resource division, ensuring Production systems priority over Base systems for CPU and memory.

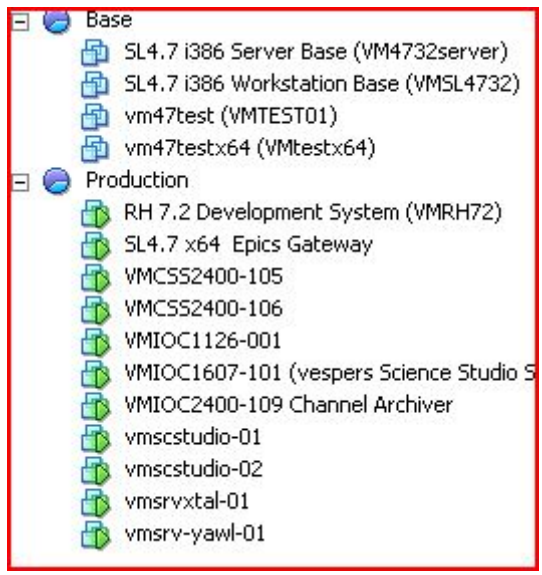

Figure 1: List of virtual machines available.

CSS2400-105, CSS2400-106 - These systems are an easy and natural fit for virtualization. The previous hardware was memory limited, and this limited the maximum number of threads that could run. Each runs a large number of EPICS servers that communicated over Ethernet with Modicon PLCs. As well, they support a number of software-only applications that control higher level configuration of the CLS hardware.

CSS2400-109 (Channel Archiver) - This system monitors and records ~23,000 Process Variables. The needs for this system are simply a network connection and all available free disk space the facility can provide.

CSS2400-103 (Application Development) - A RedHat 7.2 system used by controls developers for applications requiring legacy O/S support.

Epics Gateway – This system runs Scientific Linux 4.7 and uses IEEE 802.1Q on a network switch trunk line for direct access to 19 VLANs.

VMIOC1126-001 – A Scientific Linux 4.7 system with network support for a single beamline's VLAN and for the common control system VLAN, but with specific firewall rules for access to and from the control system VLAN.

VMIOC1607-101 – A system on a single beamline for software testing.

VMSCSTUDIO-01, VMSCSTUDIO-02, VMSRVXTAL-01, VMSRV-YAWL-01 – Development systems.

At this time, there are plans to move more systems from a physical system to a virtual system.

RISK REDUCTION

Computer Hardware Failure

With our previous configuration, a computer hardware failure required that:

- the failed system be taken off-line,
- a standby system would have its network information updated to be on the same VLAN as the failed system,
- any required system configuration update would be made (e.g. add a script to /etc/init.d)

If all the software configuration files are available within the software configuration management database then the task is not too time consuming. The reality is that not all configuration files are part of the software configuration management, which can lengthen the recovery time.

With the VM servers, the failure of one of the servers would require that any VM on the failed server be reassigned to run on the alternate server. This is done automatically through the VMware Control Station.

SAN Failure

The VM setup, while removing many hardware points of failure, does add new ones. A SAN failure would stop all the VM's running on the servers. However, as the SAN has dual power supplies, dual control intelligence, runs RAID 5, and has a hot spare drive, the reliability is expected to be high.

SO, WHAT CAN THE VM SERVER REALLY DO?

Figures 2 and 3 show resource usage.

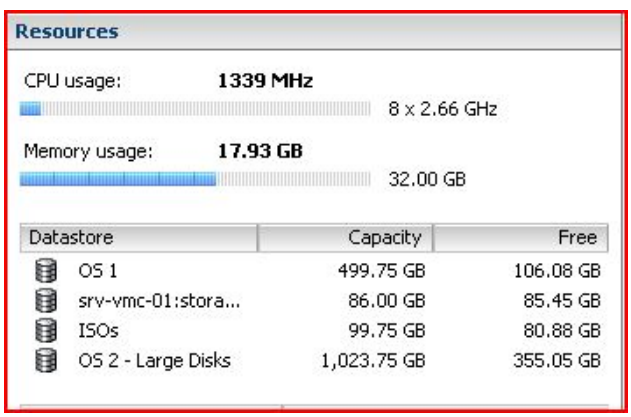

Figure 2: CPU, memory, and disk usage of the first VM server.

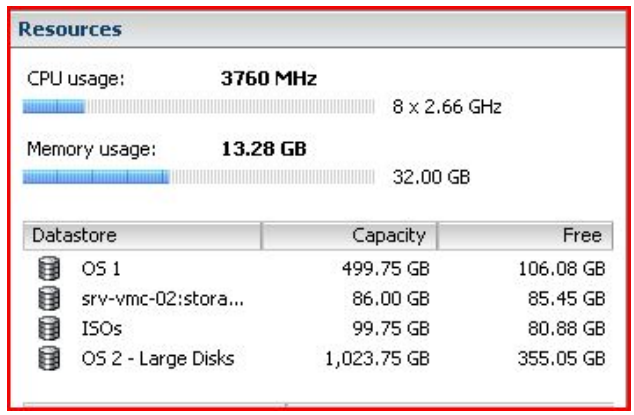

Figure 3: CPU, memory, and disk usage for the second VM server.

Although the above figures look like there is still a lot of capacity available, it is important to remember that both of the above systems act as a backup of the other.

OTHER BENEFITS OF VIRTUAL SYSTEMS

The important consideration when looking at Virtual Systems must be the contribution to keeping the facility operational. There are other benefits to Virtual Systems that can make life in the trenches easier.

- Virtual Console – any circumstance that prevents a machine from connecting to the network requires access to the console to diagnose and repair. Having a virtual console that is available on the network even when the Guest O/S isn't on the network is timesaving.
- Quick Deployment – once a decision is made to deploy a new system, it is no longer necessary to ensure rack space, network drops, or power (it is necessary to ensure available CPU and memory capacity)
- Quick Removal – in the case of a temporary system, or any system no longer required, the Virtual Machine can be shut down and removed immediately making pooled resources available for other projects.
- Quick Capacity Expansion – it is fairly easy to add additional servers to the cluster.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the assistance of Dell Professional Services, who helped design, configure and install the VM Servers.