




Canadian Light Source
Centre canadien de rayonnement synchrotron



Migrating Control System Servers to Virtual Machines

Decisions and experiences at the Canadian Light Source in 2008-2009

Glen Wright

ICALEPCS 2009

What is a Virtual Machine?

A Virtual Machine (VM) is a complete abstraction of a computer including processor, memory, storage resources and peripheral devices.

Virtual Machine technology has been available since the 1960's.

What is a Virtual System Server?

- A Virtual System Server is the Hardware (a real physical computer) and the Software (a Virtual Machine Monitor, or hypervisor) that can run one or more virtual machines.

The CLS Control System

The CLS control system computer hardware is composed primarily of:

Display Consoles for Operators

Rack mount “generic” PC’s for EPICS Services

- Mixture of Linux and Windows

Numerous small footprint IOC’s (EROC’s & Moxa’s)

A facility-wide Ethernet network using VLAN’s

VM History at the CLS

- VMware used on SUN systems in the late 1990's to provide MS Windows applications
- Slow, especially when switching between VM and Sun O/S
 - Extra layer confusing to casual computer users
 - Not used as part of the control system

VM History at the CLS

DOSemu supports an old application required for facility equipment monitoring.

- Installed on a PC/104 Linux system
- Uses the serial line to communicate
- Allows staff to remotely log in rather than travelling to the install location

Am I Doing Cloud Computing?

Does this give me Cloud Computing?

No


What Problem needed Solving?

- The increasing number of computer systems necessary to facility operation increased risk of down-time
- There were a number of computer maintenance contracts that were reaching end-of-term
- This included a computer that used 802.1Q directly, and computers that were still running RedHat 7.2

What Were the Options?

1. Do Nothing, Buy More Spares
2. Replace with equivalent current hardware
3. Replace with a Blade Server
4. Replace with a Virtual System Server

Option 1: Do Nothing




Pro

- Systems are running
- Cheap (on paper, short term)

Con

- Difficult to find compatible hardware for legacy system spares
- Failure requires reconfiguring a standby system to match the failed system

Option 2: Replace




Pro

- Continue working with systems the way we know how
- Can improve reliability using systems with redundant components

Con

- Strongly encouraged to update software to newer O/S releases
- Systems with higher redundancy have higher cost

Option 3: Blade Server




Pro

- Less Rack Space
- Less Power
- Possibly fewer network drops

Con

- Direct cost comparison indicated we might not reach a break-even point
- Still have individual system configurations

Option 4: VM Server




Pro

- Redundancy at the server level allows failover for the Guest O/S's
- Legacy systems easily supported into foreseeable future
- IEEE 802.1Q connection to network reduces the number of network drops

Con

- Some legacy systems (e.g. Fiber Optic Link to VME crate) not supported.



What did we do?

Purchased a Virtual System Server

- VMware ESX software
- “Cluster” of 2 DELL 2950 Servers
 - 32 Gig Memory
 - Dual 4-core processors
- iSCSI SAN with 3.6 Terabytes of RAID Storage

What Didn't This Solve?

Not all computers can be virtualized.

Any system requiring direct physical connection to a physical device other than video, keyboard, mouse, Ethernet, or hard disk cannot be virtualized.


New Problem – Naming!

The CLS uses a standardized prefix, a location number, and a sequence number within a location for labeling equipment.


This had already been a minor issue when servers had been moved from one location to another.

Short term solution- give a name as if the system was physically installed, and then add the prefix **VM**.


New Opportunities



- Inter-department support and knowledge exchange
 - IT department involvement early on means expertise not unnecessarily duplicated
- Centralized hardware monitoring














IMPLEMENTATION




- Allows multiple VM's
- Provides virtual disk space
- Provides virtual network switches – each VM can have up to 4 virtual ports on the virtual network switches




-  RH 7.2 Development System (VMRH72)
-  SL4.7 x64 Epics Gateway
-  VMCSS2400-105
-  VMCSS2400-106
-  VMIOC1126-001
-  VMIOC1607-101 (vespers Science Studio 5)
-  VMIOC2400-109 Channel Archiver
-  vmsscstudio-01
-  vmsscstudio-02
-  vmsrvxtal-01
-  vmsrv-yawl-01

Base Systems




- Lower priority – basis for templates
- Allow standardized part of OS configuration to be found in one location

Production Systems



- Higher resource pool priority
- Customized (if necessary) for different tasks

Example System 1



- loc2400-106 – runs many EPICS IOC applications, most of which communicate via TCP/IP with MODICON PLC's
- O/S updated from 32bit RedHat 7.2 to 64bit Scientific Linux 4.7
- Uses the VMware network switch to provide access to Controls VLAN
- No changes for applications running on the IOC

Example System 2

- Epics Gateway Host
- Runs one instance of the EPICS Gateway software per connected VLAN
- Has 19 virtual NICs configured by the guest O/S (uses Linux VLAN driver)
- O/S is Scientific Linux 4.7, 64-bit installation

File Edit View Inventory Administration Plugins Help

Inventory Scheduled Tasks Events Administration Maps

Hosts & Clusters

- CID
 - Main Cluster
 - srv-vmc-01.clsi.ca
 - srv-vmc-02.clsi.ca
 - Base
 - SL4.7 i386 Server Base (VM4732server)
 - SL4.7 i386 Workstation Base (VMSL4732)
 - vm47test (VMTEST01)
 - vm47testx64 (VMtestx64)
 - Production
 - RH 7.2 Development System (VMRH72)
 - SL4.7 x64 Epics Gateway
 - VMCSS2400-105
 - VMCSS2400-106
 - VMIOC1126-001
 - VMIOC1607-101 (vepsers Science Studio S...
 - VMIOC2400-109 Channel Archiver
 - vmcstudio-01
 - vmcstudio-02
 - vmrvxtal-01
 - vmsrv-yawl-01

SL4.7 x64 Epics Gateway

Getting Started Summary Performance Tasks & Events Alarms Console Permissions Maps

What is a Virtual Machine?

A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. An operating system installed on a virtual machine is called a guest operating system.

Because every virtual machine is an isolated computing environment, you can use virtual machines as desktop or workstation environments, as testing environments, or to consolidate server applications.

In VirtualCenter Server, virtual machines run on hosts or clusters. The same host can run many virtual machines.

Basic Tasks

- Power off the virtual machine
- Suspend the virtual machine
- Edit virtual machine settings

Explore Further

- Learn more about virtual machines
- Learn how to install an operating system

Diagram labels: VI Client, VirtualCenter Server, Datacenter, Cluster, Host, Virtual Machines.

Recent Tasks

Name	Target	Status	Initiated by	Time	Start Time	Complete Time
Scan Entity	srv-vmc-01.clsi...	The task was ...	CLSI\kingtut06	24/03/2009 11:04:11 A	24/03/2009 11:04:43 ...	24/03/2009 11:04:43 ...
Remediate Entity	srv-vmc-01.clsi...	The task was ...	CLSI\kingtut06	24/03/2009 11:01:52 A	24/03/2009 11:04:51 ...	24/03/2009 11:04:51 ...

VMware Infrastructure Client

SL4.7 x64 Epics Gateway



Getting Started Summary Performance Tasks & Events Alarms Console Permissions Maps

General	Resources
Guest OS: Red Hat Enterprise Linux 4 (64-bit)	CPU usage: 526 MHz

General

Resources

CPU usage: **509 MHz**
 Host memory usage: **1.63 GB**
 Guest memory usage: **522.00 MB**

Datstore		Capacity	Free
 OS 1		499.75 GB	106.08 GB
 ISOs		99.75 GB	80.88 GB

Network

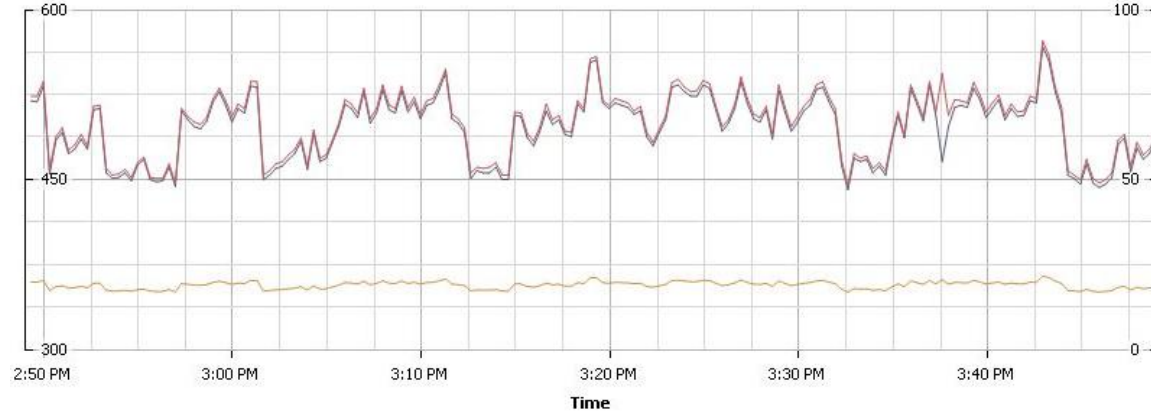
 VGT

SL4.7 x64 Epics Gateway

Getting Started Summary Performance Tasks & Events Alarms Console Permissions Maps

J/Real-time, 02/10/2009 2:49:02 PM - 02/10/2009 3:49:02 PM Change Chart Options.. Switch to: CPU

h refreshes every 20 seconds



Performance Chart Legend


Object	Measurement	Units	Latest	Maximum	Minimum	Average
0	CPU Usage in MHz (Average)	MHz	484	568	441	495.74
SL4.7 x64 Epics Gateway	CPU Usage in MHz (Average)	MHz	488	573	445	499.96
SL4.7 x64 Epics Gateway	CPU Usage (Average)	Percent	18.35	21.54	16.76	18.81



SL4.7 x64 Epics Gateway

Getting Started Summary Performance Tasks & Events Alarms Console Permissions Maps

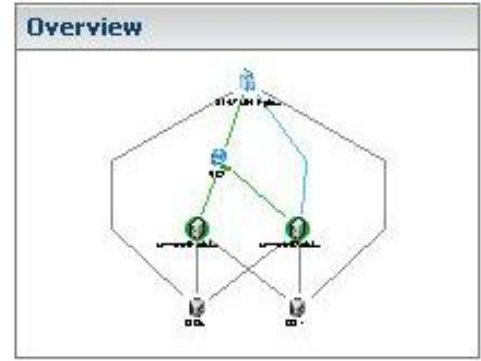
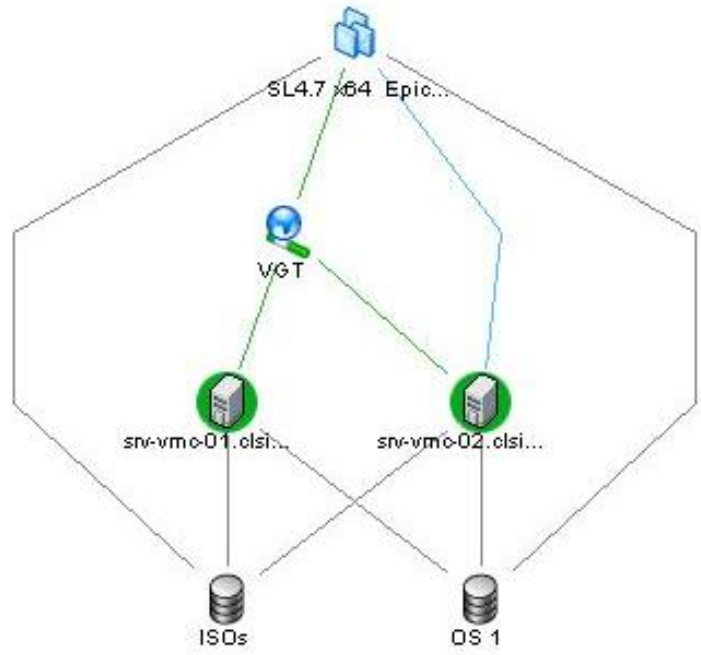
```
Sep 15 14:45:19 PU Gateway Version 2.0.2.1 [Sep 18 2007 16:59:05]
EPICS 3.14.9 PID=12848 ServerPID=12847
EPICS_CA_ADDR_LIST=10.50.255.255
EPICS_CA_AUTO_ADDR_LIST=NO
EPICS_CA_SERVER_PORT=Not specified
EPICS_CA_MAX_ARRAY_BYTES=Not specified
EPICS_CAS_INTF_ADDR_LIST=10.52.12.254
EPICS_CAS_SERVER_PORT=Not specified
EPICS_CAS_IGNORE_ADDR_LIST=10.52.12.254
Running as user control on host EpicsGateway
Statistics PU prefix is gw643
Sep 23 10:53:29 Warning: Virtual circuit disconnect vmIOC2400-105.cs.clsi.ca:380
04
Oct 01 15:44:15 gateServer::exCB: Channel Access Exception:
  Channel Name: Unavailable
  Native Type: Unavailable
  Native Count: 0
  Access: Unavailable
  IOC: Unavailable
  Message: Identical process variable names on multiple servers
  Context: Channel: "TM1606-4-B10-01", Connecting to: vmIOC2400-105.cs.clsi.ca:4
2569, Ignored: vmIOC2400-105.cs.clsi.ca:32857
  Requested Type: TYPENOTCONN
  Requested Count: 0
control@EpicsGateway:67 >_
```




SL4.7 x64 Epics Gateway

- Getting Started
- Summary
- Performance
- Tasks & Events
- Alarms
- Console
- Permissions
- Maps

[Refresh](#)



Retrieval error. Try again later.




THE FUTURE

Other Servers

There are still a number of Servers that can be moved when the hardware they're on is no longer acceptable for use.

The IT department is evaluating another VM server system.

VM Server upgrades



The memory usage is very close to, and has on occasion exceeded, the total that would be available in the event of a single server failure. The controls group is planning on adding a third server to ensure continued smooth operation.

Acknowledgements

- Dell Professional Services for their assistance in designing, configuring, and installing the VM Servers