When running UC virtualized, the UC application does not follow an appliance model (where management of the underlying hardware is provided through the application).

When running virtualized, be sure to complete the following tasks manually (outside of the application itself).

Contents

- 1 Upgrading ESXi
- 2 Upgrading the UC applications
- 3 Monitor Your Hardware Health
  - 3.1 Out-of-Band Hardware Monitoring
    - 3.1.1 CIMC
    - 3.1.2 IPMI over LAN (ipmitool, ipmiutil)
  - 3.2 In-Band Hardware Monitoring
    - 3.2.1 ESXi Access
    - 3.2.2 vSphere Client
    - 3.2.3 VMware API
    - 3.2.4 CIM, WBEM and SMASH
    - 3.2.5 wbemcli
    - 3.2.6 IBM Director
    - 3.2.7 LSI MegaRAID Storage Manager (MSM)
    - 3.2.8 syslog
- 4 Upgrading Firmware on your TRCs
- 5 Backup, Restore, and Server Recovery

Upgrading ESXi

Multiple versions of ESXi are supported by UC applications.

To upgrade the ESXi software on a host, you must power off all virtual machines or migrate the virtual machine to a different host.

Instructions for installing or upgrading to a specific ESXi release are available in the release notes of the ESXi version you install. For example, here are the **VMware ESX 4.1 Update 1 Release Notes.**

After upgrading ESXi, or whever moving a virtual machine to a different host that is running a different version of ESXi, VMware Tools must be upgraded on the UC applications so that the tools versions shows "Up to Date" in the vSphere Client (see **VMware Tools**).

It is OPTIONAL to upgrade the VMV (Virtual Machine Version) version of the virtual machine to the newer VMV version. If you do upgrade to the newer VMV version, note that your virtual machine cannot be

changed back to the previous VMV version, and that the new VMV version will NOT run on your older version of ESXi. For example, if you are running a vmv7 virtual machine on ESXi 4.1, and you upgrade to ESXi 5.0, you can optionally upgrade the vm to vmv8. Once you upgrade to vmv8, your virtual machine can only run on ESXi 5.0 and later. Once upgraded to vmv8, the virtual machine cannot be converted back to vmv7. For further information see VMware documentation.

Upgrading the UC applications

When running UC virtualized, it is necessary that the virtual machine configuration be set properly for the specific release of the UC application that you are running.

When doing a fresh install, it is required to deploy the OVA for the application that matches the release of the UC application itself. This ensures that the fresh install has the proper virtual machine configuration for the release being installed.

When a UC application is upgraded (eg. upgrading CUCM from 8.5.1 to 8.6.2), it is necessary to go to www.cisco.com and look at the Readme file for the OVA of the release you are upgrading to. For example: go to www.cisco.com and go to Support -> Downloads. From there Products -> Voice and Unified Communications -> IP Telephony -> Call Control -> Cisco Unified Communications Manager (CallManager) -> Cisco Unified Communications Manager Version 8.6 -> Unified Communications Manager Virtual Machine Templates-8.6(1)

When you're at the appropriate OVA, select "Download Now" and you will be taken to a page where you can click to view the Readme file. In the Readme file, you will see the virtual machine settings for the release. You must edit settings on the virtual machine you have upgraded to match the setting in the Readme file.

You will need to make the changes (edit settings) while the virtual machine is powered off. After you are booted up on the new software version, you must gracefully shut down the virtual machine, change the settings, and boot it back up.

There may be cases where there are no required changes. It will depend on the specific "from" and "to" versions of the upgrade.

Monitor Your Hardware Health

When deployed in a virtualized environment, the UC applications do not monitor the hardware. Hardware must be monitored independently from the UC applications. There are multiple ways to do this:

- Out-of-Band Hardware Monitoring
- In-Band Hardware Monitoring

Each method is explained in this page.

Out-of-Band Hardware Monitoring

Out-of-Band Hardware Monitoring uses Intelligent Platform Management Interface (IPMI) to communicate with the Baseboard Management Controller (BMC). Common IPMI interfaces are HP Integrated Lights Out (iLO) and Cisco Integrated Management Controller (CIMC). IPMI allows the user to to inspect system sensors, power cycle the chassis, obtain remote console and manipulate virtual media. UCS C-series hardware provide IPMI over LAN for use by tools such as ipmitool.

CIMC

CIMC provides GUI and CLI interfaces to hardware sensors and other status. This information is available via ipmitool using IPMI over LAN. The login screen of CIMC shows hardware status:

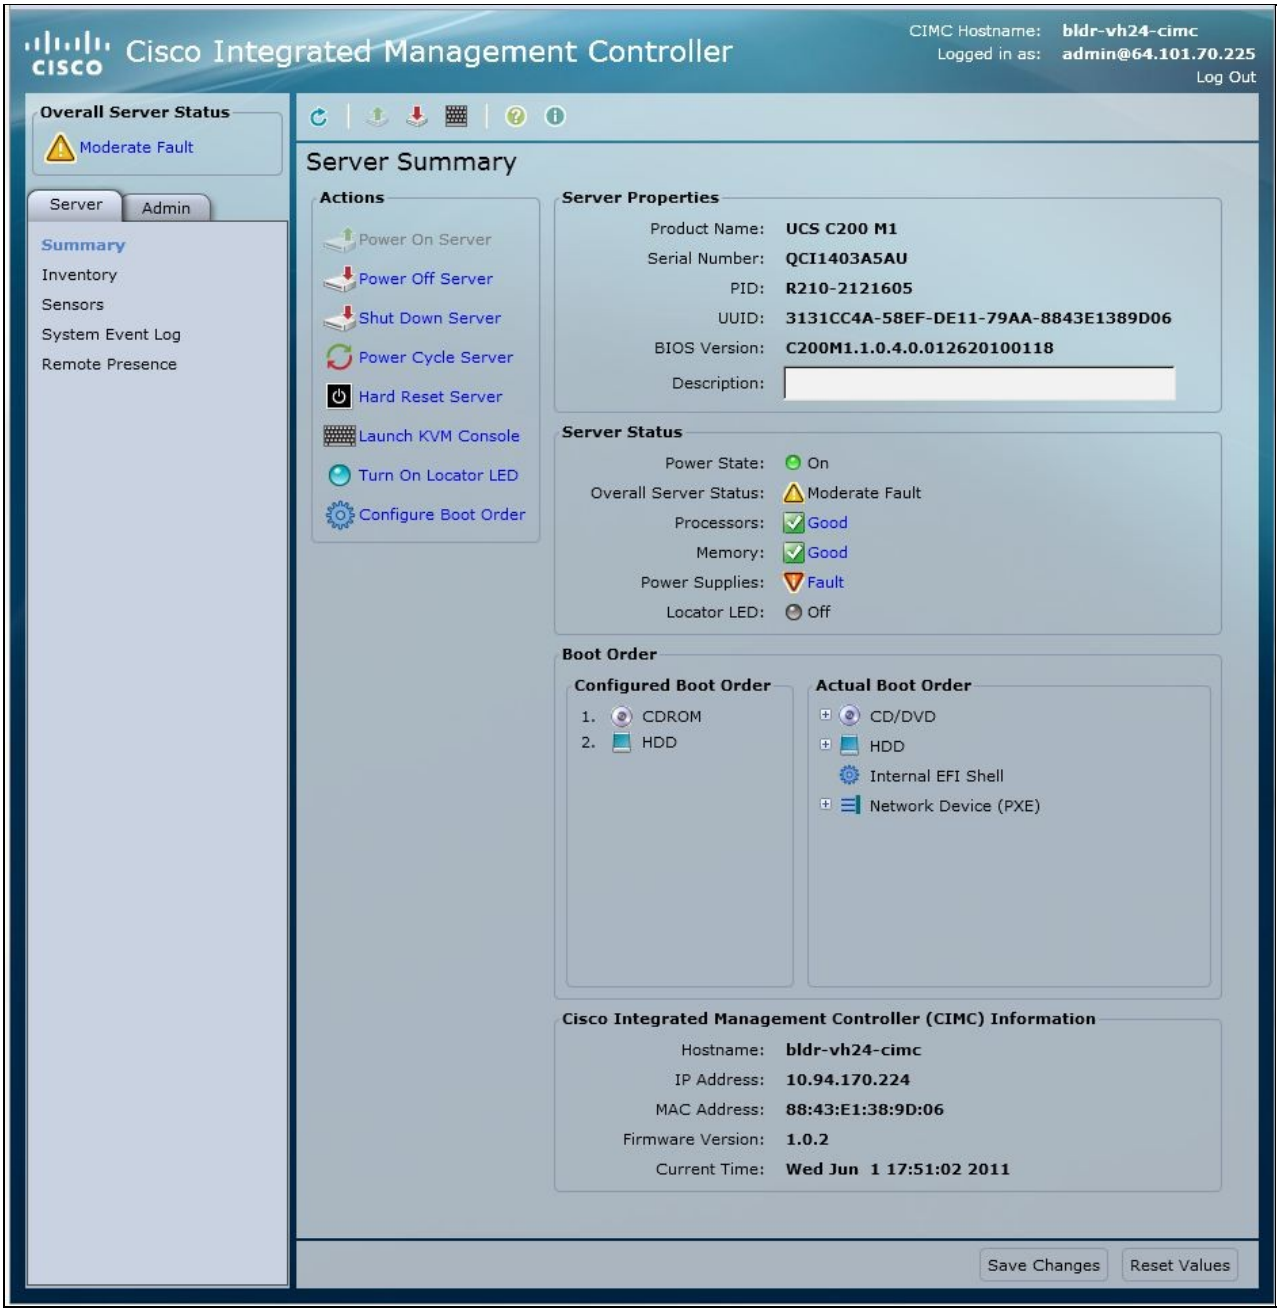

See the selections for hardware sensors and inventory in the following image:

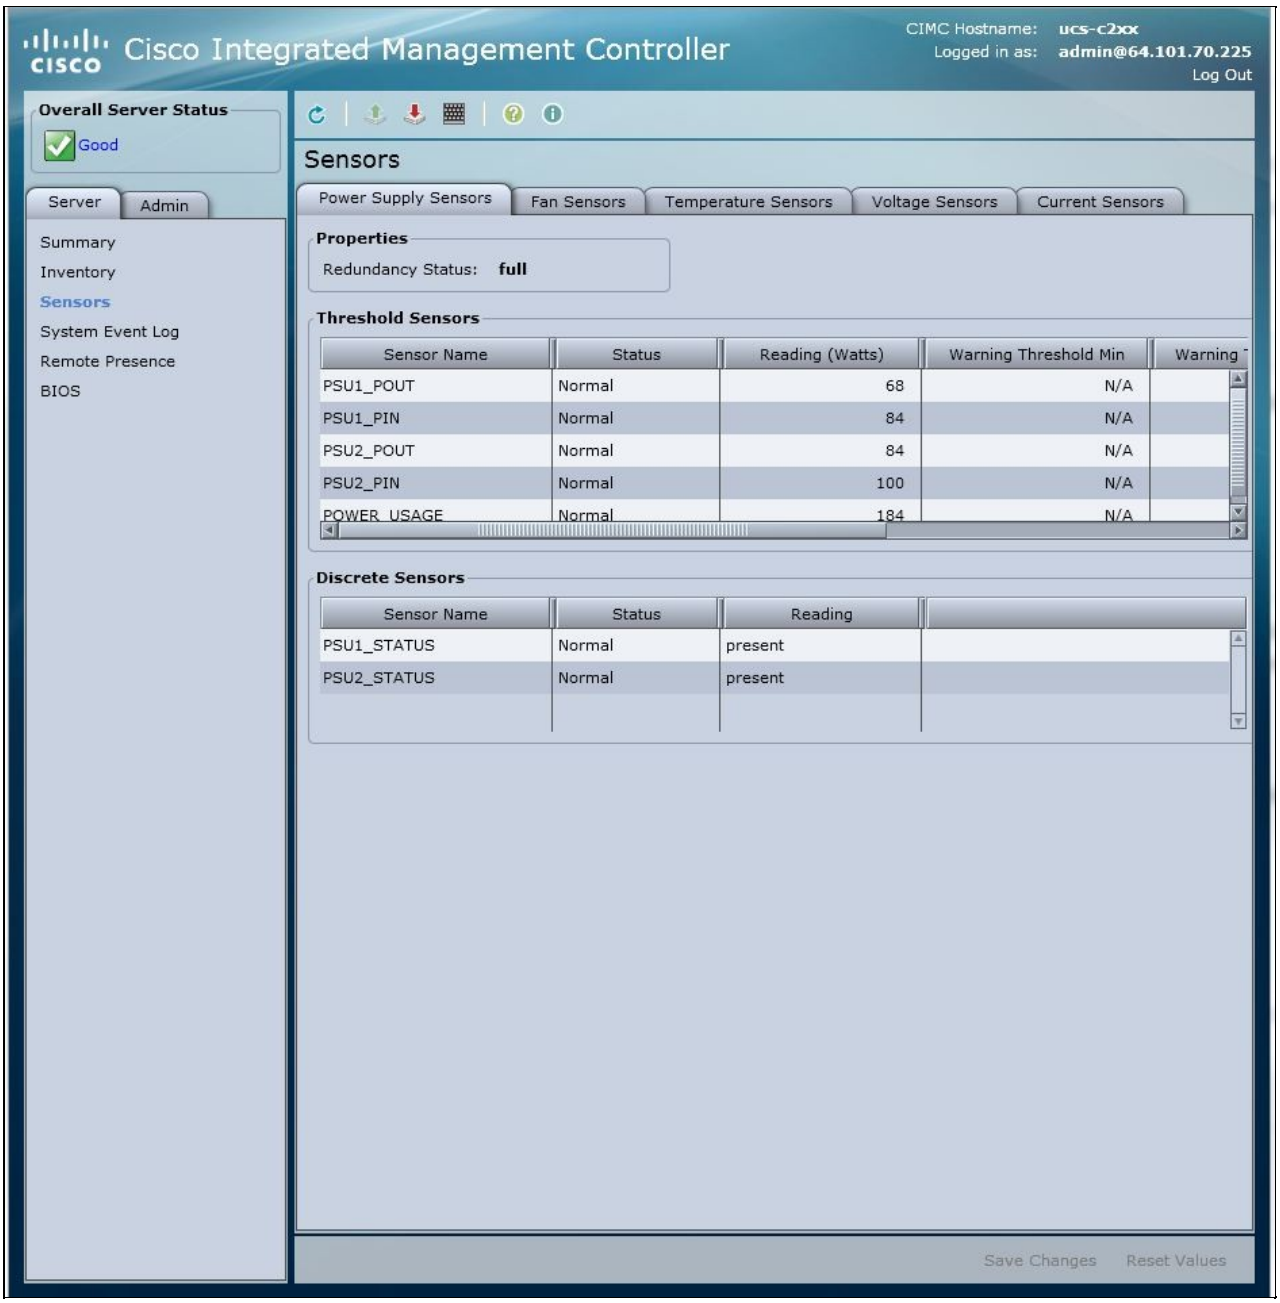

IPMI over LAN (ipmitool, ipmiutil)

ipmiutil connects to the CIMC over LAN UDP port 623 (ASF remote management). There are a huge number of options to this command; some of which can be used to set up Platform Event Filtering (pef) and alerting. We have not set up alerting from ipmitool.

Commands to dump the event log and sensor data:

```
export IPMI_PASSWORD=<password>
ipmitool -E -H <iLO, CIMC or IMM IP address> -U <userid> sel list
ipmitool -E -H <iLO, CIMC or IMM IP address> -U <userid> sensor list
```

ipmitool sel list dumps the system event log and can generate thousands of lines of output. Here is a small sample:

```
6a8 | Pre-Init Time-stamp | Processor #0x20 | Predictive Failure Deasserted
```

```
6a9 | Pre-Init Time-stamp | Power Supply #0x60 | Presence detected | Asserted
6aa | Pre-Init Time-stamp | Power Supply #0x61 | Presence detected | Asserted
6ab | Pre-Init Time-stamp | Power Supply #0x61 | Power Supply AC lost | Asserted
6ac | Pre-Init Time-stamp | Entity Presence #0xb3 | Device Present
6ad | Pre-Init Time-stamp | Entity Presence #0x66 | Device Present
6ae | Pre-Init Time-stamp | Platform Alert #0x59 |
6af | Pre-Init Time-stamp | Platform Alert #0xb0 |
6b0 | Pre-Init Time-stamp | Entity Presence #0x64 | Device Absent
6b1 | Pre-Init Time-stamp | Entity Presence #0x65 | Device Absent
6b2 | Pre-Init Time-stamp | Entity Presence #0x67 | Device Absent
```

ipmitool sensor list generates lots of output; here is a sample:

```
PSU1_VOUT | 12.000 | Volts | ok | na | na | na | na | 13.000 | 15.000
PSU1_IOUT | 13.000 | Amps | ok | na | na | na | 53.000 | 58.000 | 60.000
PSU1_TEMP_1 | 40.000 | degrees C | ok | na | na | na | 60.000 | 65.000 | 70.000
PSU1_FAN_1 | 2944.000 | RPM | ok | na | na | na | na | na | na
PSU1_POUT | 164.000 | Watts | ok | na | na | na | 652.000 | 680.000 | 700.000
PSU1_PIN | 184.000 | Watts | ok | na | na | na | 652.000 | 680.000 | 700.000
PSU2_VOUT | 0.000 | Volts | ok | na | na | na | na | 13.000 | 15.000
PSU2_IOUT | 0.000 | Amps | ok | na | na | na | 53.000 | 58.000 | 60.000
PSU2_TEMP_1 | 31.000 | degrees C | ok | na | na | na | 60.000 | 65.000 | 70.000
PSU2_FAN_1 | 0.000 | RPM | ok | na | na | na | na | na | na
PSU2_POUT | 0.000 | Watts | ok | na | na | na | 652.000 | 680.000 | 700.000
PSU2_PIN | 0.000 | Watts | ok | na | na | na | 652.000 | 680.000 | 700.000
```

In-Band Hardware Monitoring

In-Band hardware monitoring talks to an agent or provider on the host operating system via the Common Information Model (CIM). CIM defines standard XML definitions of computer equipment, including RAM, CPU and peripherals such as disks, RAID controllers, and other options.

CIM data is monitored by:

- Third-party tools like IBM Director
- Linux tools such as wbemcli
- vCenter

ESXi Access

ESXi is a closed platform. The following network ports are provided:

- TCP 22 (ssh) Same as "unsupported" login, only if enabled
- TCP 80 (http)
- TCP 443 (vSphere, VMware API)
- TCP 902 (older VMware API)
- UDP 427 (SLP)
- TCP 5989 (WBEM)

vSphere Client

vShpere Client, when logged directly into an ESXi host, provides a **Health and Status** selection under the **Configuration** tab, from vCenter, this information is available from the Hardware tab when you've selected a host in the Hosts and Clusters view. Health and Status gives an overview of individual physical drives and logical drive groups (RAID arrays). This information is transferred via the following CIM classes in the root/cimv2 namespace:

- VMware_HHRCDiskDrive
  Information about each physical drive including
- VMware_HHRCStorageVolume
  Information about the logical volumes (RAID groups), including

  ◊ Member drives (<drive number>e<enclosure number>)
  ◊ Array state (Optimal, degraded, etc)

Here is an image of the vSphere Client showing sensor data by connecting the viClient directly to the ESXi host.

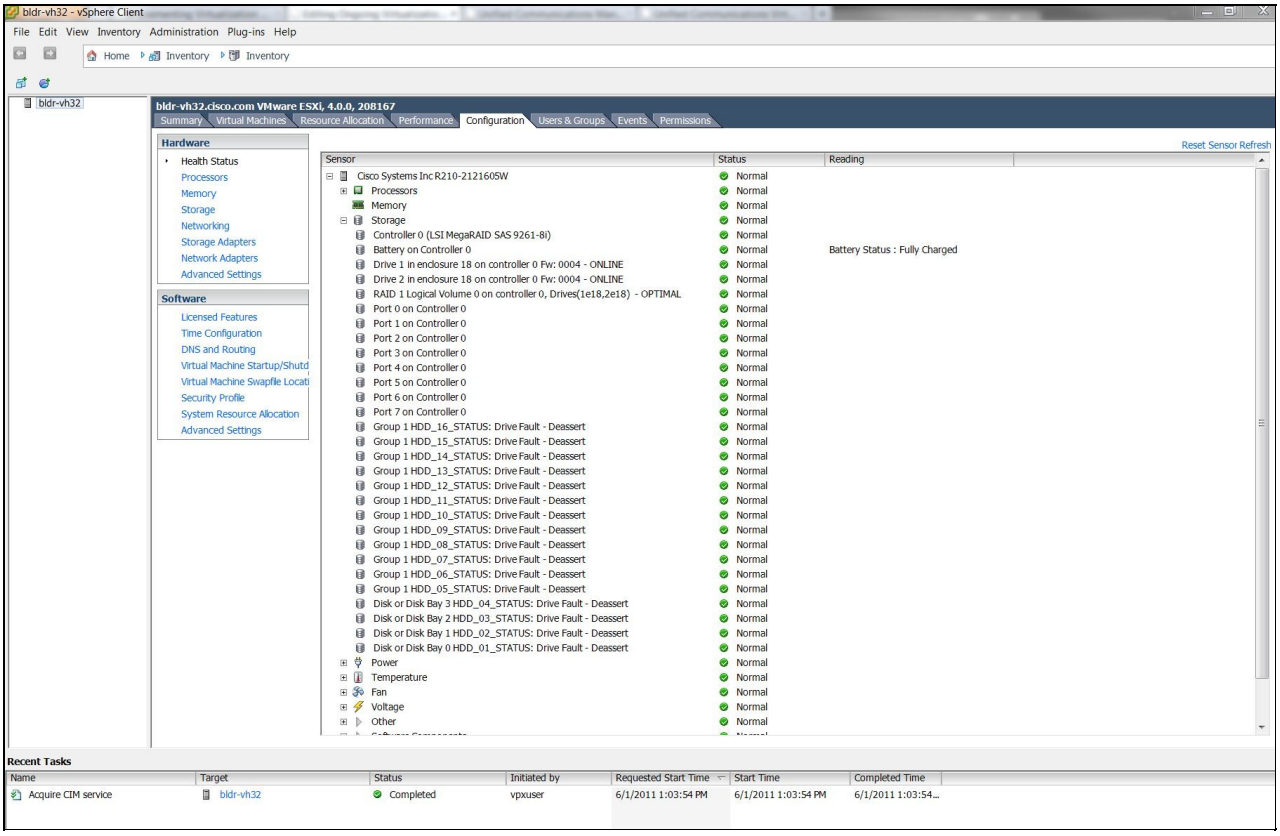

VMware API

The VMware API is a set of perl and java libraries and classes that talks SOAP over port 443 to the ESXi host to implement functions used by vSphere Client and virtual Center.

The VMware API is available from vmware.com and comes with the vSphere Manager Assistant (vMA). The vMA OVF can be deployed from **http://download3.vmware.com/software/vma/vMA-ovf-4.0.0-161993.ovf** and installed in a virtual machine.

Everything that can be done from vSphere Client can be done from perl scripts using the **VMware Perl SDK,** including virtual machine creation, configuration and management, performance statistics collection, virtual switch administration and virtual host administration.

Here is a simple VM query:

```
[root@bldr-vcm9 apps]# export VI_PASSWORD=cisco123
[root@bldr-vcm9 apps]# export VI_USERNAME=root
```

```
[root@bldr-vcm9 apps]# vm/vminfo.pl --server bldr-vh29 --vmname sjc-rfd-pub-1

Information of Virtual Machine sjc-rfd-pub-1

Name: sjc-rfd-pub-1
No. of CPU(s): 2
Memory Size: 6144
Virtual Disks: 2
Template: 0
vmPathName: [vh29_RAID5_8] sjc-rfd-pub-1/sjc-rfd-pub-1.vmx
Guest OS: Red Hat Enterprise Linux 4 (32-bit)
guestId: rhel4Guest
Host name: sjc-rfd-pub-1
IP Address: 10.9.10.4
VMware Tools: VMware Tools is running, but the version is not current
Cpu usage: 243 MHz
Host memory usage: 6202 MB
Guest memory usage: 552 MB
Overall Status: The entity is OK
```

Here's a command to display network counters:

```
performance/viperformance.pl --url https://bldr-vh29/sdk/vimService --username root --password cis
```

CIM, WBEM and SMASH

WBEM and SMASH use SLP discovery to locate manageable hosts.
Web Enterprise Management and Common Information Model; WBEM and CIM

```
slptool findattrs service:wbem:https://10.94.150.229:5989
```

returns nothing, and bldr-vh29 is not discoverable by LSI MSM. That's because slptool is using SLP and multicasts to svrloc are not answered by 10.94.150.229.

```
11:11:03.967177 IP bldr-ccm39.cisco.com.32829 > 239.255.255.253.svrloc: UDP, length 56
```

```
slptool findattrs service:wbem:https://10.94.150.228:5989
```

returns among other things the different namespaces, in bold, below, that can be queried with wbemcli:

(template-type=wbem),(template-version=1.0),(template-description=This template describes the attributes used for advertising WBEM
Servers.),(template-url-syntax=https://10.94.150.228:5989),(service-hi-name=Small Footprint CIM Broker 1.3.0),(service-hi-description=Small Footprint CIM Broker 1.3.0),(service-id=sfcb:10.94.150.228),(CommunicationMechanism=CIM-XML),(InteropSchemaNamespace=root/interop Read,Association Traversal),(FunctionalProfileDescriptions=Basic Read,Association Traversal),(MultipleOperationsSupported=true),(AuthenticationMechanismsSupported=Basic),(**Namespace=qlogic/cimv2** Hardware RAID Controller:Block Services,ANSI:Host Hardware RAID Controller:DA Target Ports,ANSI:Host Hardware RAID Controller:Disk Drive Lite,ANSI:Host Hardware RAID Controller:Drive Sparing,ANSI:Host Hardware RAID Controller:Extent Composition,ANSI:Host Hardware RAID Controller:Fan,ANSI:Host Hardware RAID Controller:Generic Initiator Ports,ANSI:Host Hardware RAID Controller:Indications,ANSI:Host Hardware RAID Controller:Job Control,ANSI:Host Hardware RAID Controller:Physical Asset,ANSI:Host Hardware RAID Controller:Physical Package,ANSI:Host Hardware RAID Controller:Power Supply,ANSI:Host Hardware RAID Controller:SAS/SATA Initiator Ports,ANSI:Host Hardware RAID Controller:Software Identity,ANSI:Host Hardware RAID Controller:Software Update,ANSI:Host Hardware RAID Controller:Storage Enclosure,ANSI:Host Hardware RAID,Other:LSI Support:StoreLib Cmd,DMTF:RAID Controller Alarm,ANSI:Server,SNIA:FC

HBA,SNIA:Server,SNIA:HDR,DMTF:Software Inventory,Other:IPMI OEM Extension,DMTF:Sensors,DMTF:Profile Registration,DMTF:System Memory,DMTF:Record Log,DMTF:CPU,DMTF:Fan,DMTF:Base Server,DMTF:Battery,DMTF:Physical Asset,DMTF:Power Supply,DMTF:Power State Management,SNIA:FC Initiator Ports,SNIA:Indication)

For some reason 10.94.150.228 is responding to the multicast:

tcpdump port 427 on bldr-ccm142, running on bldr-vh29, **does not** see the multicast traffic.
tcpdump port 427 on bldr-ccm144, running on bldr-vh29, **does** see the multicast traffic:

```
03:18:37.236162 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 56
03:18:37.985780 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 56
03:18:38.985759 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 72
03:18:43.986543 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 85
03:18:44.986499 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 85
03:18:54.003788 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 56
03:18:54.753108 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 56
03:18:55.753082 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 72
03:18:56.503037 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 72
03:18:57.502985 IP 10.94.150.39.32830 > 239.255.255.253.svrloc: UDP, length 72
```

Perhaps the multicast group registrations via IGMP aren't being seen or processed by our lab switches, and the multicast traffic isn't routed to the management LAN of some of these ESXi hosts. At any rate, SLP discovery has been unreliable to ESXi hosts, at least in our lab.

wbemcli

wbemcli is a command line interface to CIM servers. This command is able to dump information about
The LSI CIM namespace is lsi/lsimr12. This command enumerates all classes in that namespace:

```
wbemcli ecn -nl -noverify 'https://root:<password>@<machine IP>:5989/<namespace>'
```

<namespace> can be root/cimv2, root/interop or any of the namespaces returned by slptool, above.

This command sees a drive in that is being rebuilt and probably uses the same namespace and class that Virtual Center uses:

```
wbemcli ei -noverify 'https://root:<password>@<ESXi Host IP>:5989/root/cimv2:VMware_HHRCDiskDrive'
```

This command is able to dump storage volumes (RAID arrays) and tell whether they are OPTIMAL or DEGRADED:

```
wbemcli ei -noverify 'https://root:<password>@<ESXi Host IP>:5989/root/cimv2:VMware_HHRCStorageVol
```

We can query the state of the RAID controller's write cache battery backup via:

```
wbemcli ei -noverify 'https://root:<password>@<ESXi Host IP>:5989/root/cimv2:VMware_HHRCBattery'
```

IBM Director

We can provision a host's IP address directly in IBM director and so don't rely on SLP. We are able to monitor drive removal events in director:

- Click on **System Discovery** under **Discovery Manager** on the Welcome to IBM Director start page.
- Enter the host's IP address and click the **Discover** button

The resulting discovery process takes about 1/2 hour to complete. There are certainly faster ways to do this but this works. Once the system is discovered:

- Click on **Navigate Resources** located on the left.
- Click on **All Systems.**
- Click on this host's entry under the **Access** column and request access.

Once granted access, collect the inventory of this host's hardware:

- Click on this hosts in the table.
- Click on the **Collect Inventory** button.
- Select **Run Now** (default) and run the job.

Individual drives, DIMMS, sensors and such will be displayed under the inventory for this host, after inventory collection has completed. This typically takes a minute or two.

The host's Event Log is also visible from this window, under the Event Log tab.

- Open on **System Status and Health** on the left hand side of Director.
- Click on **Event Log**.

Events for all machines come to this log. Events for a particular machine can be viewed by:

- Click on **Navigate Resources**.
- Click on **All Systems**.
- Click on the host's **Operating System** entry (not the Server entry, which has an empty Event Log).
- Click on the **Event Log** tab.

We pulled drives on two machines, bldr-vh27 and bldr-vh28, to generate events.

This is the event log for bldr-vh27 after pulling a drive:

We did the following to rebuild the RAID arrays on these 2 machines:

- Reboot and hit CTRL-Y to enter the Preboot CLI
- Reinsert the drive and mark it ready for removal

```
-pdinfo -physdrv [252:5] -a0 (verify the drive was missing and reinsert it)
```

-pdmakegood -physdrv [252:5] -a0 (mark the drive good) -pdprprmv -physdrv [252:5] -a0 (prepare the drive for removal) -pdlocate -start [252:5] -a0 (turn on the locator LED)

- Remove and reinsert the drive
- The RAID array will be rebuilt.

During rebuild, we see many rebuild progress events in syslog (event code 0x0067):

```
May 20 13:26:18 vmkernel: 0:00:22:48.165 cpu5:4101)<6>megasas_service_aen[4]: aen received
May 20 13:26:18 vmkernel: 0:00:22:48.165 cpu0:4337)<6>megasas_hotplug_work[4]: event code 0x0067
May 20 13:26:18 vmkernel: 0:00:22:48.175 cpu0:4337)<6>megasas_hotplug_work[4]: aen registered
```

Then we see the final syslog entries indicating rebuild is complete and the RAID array is OPTIMAL (event code 0x00f9):

```
May 20 13:26:21 vmkernel: 0:00:22:50.694 cpu5:4101)<6>megasas_service_aen[4]: aen received
May 20 13:26:21 vmkernel: 0:00:22:50.694 cpu14:4328)<6>megasas_hotplug_work[4]: event code 0x00f9
```

While the RAID array is being rebuilt we see the following events in the Event Log:

LSI MegaRAID Storage Manager (MSM)

The LSI MSM Release 3.6 (downloaded as 2.91) is able to monitor, configure and repair RAID arrays while ESXi is active. This LSI MSM depends on SLP via multicast to find servers. Multicast to an ESXi host appears to be unreliable, and server discovery of ESXi hosts suffers as a result. This Release of MSM does not support alerting on specified events.

LSI MSM 6.9 does appear to support email alerting. However it still depends on SLP discovery and so must be installed on a VM that shares vSwitch0 with its ESXi host. Once the server is discovered:

- LSI MSM opens wbem-https port (5989) to ESXi host
- ESXi host provides CIM indications to LSI MSM over wbem-exp-https port (5990)
- LSI MSM allows configuration and rebuilding of RAID arrays in place without any ESXi outage.
- LSI MSM provides real time display of drive rebuild and completion events
- LSI MSM does not yet provide alerting via email or other actions.

Here is a screen shot of LSI MSM showing the logical drives:

syslog

Direct inspection of ESXi syslog requires you log into the CIMC KVM console:

- Press ALT-F1
- Enter unsupported followed by the root password
- vi /var/log/messages

It is possible to enable remote syslog in ESXi by specifying the host name or IP address of a syslog server under the Configuration tab and Advanced Settings. We used our CUCM publisher and it worked fine.

The following syslog messages are generated for RAID drive rebuild progress indication:

```
May 20 13:04:06 vmkernel: 0:00:00:36.227 cpu0:4115)<6>megasas_service_aen[4]: aen received
May 20 13:04:06 vmkernel: 0:00:00:36.227 cpu14:4335)<6>megasas_hotplug_work[4]: event code 0x0067
```

These syslog messages indicate the rebuild is complete:

```
May 20 13:26:19 vmkernel: 0:00:22:49.412 cpu5:4101)<6>megasas_service_aen[4]: aen received
May 20 13:26:19 vmkernel: 0:00:22:49.412 cpu8:4330)<6>megasas_hotplug_work[4]: event code 0x0063
May 20 13:26:20 vmkernel: 0:00:22:49.643 cpu8:4330)<6>megasas_hotplug_work[4]: aen registered
May 20 13:26:20 vmkernel: 0:00:22:49.644 cpu5:4101)<6>megasas_service_aen[4]: aen received
May 20 13:26:20 vmkernel: 0:00:22:49.644 cpu9:4329)<6>megasas_hotplug_work[4]: event code 0x0064
```

These syslog messages indicate the RAID array is restored to OPTIMAL state after the rebuild:

```
May 20 13:26:20 vmkernel: 0:00:22:49.644 cpu12:4331)<6>megasas_hotplug_work[4]: event code 0x0072
May 20 13:26:20 vmkernel: 0:00:22:49.644 cpu12:4331)<6>megasas_hotplug_work[4]: aen registered
May 20 13:26:20 vmkernel: 0:00:22:49.644 cpu5:5403)<6>megasas_service_aen[4]: aen received
May 20 13:26:20 vmkernel: 0:00:22:49.644 cpu14:4332)<6>megasas_hotplug_work[4]: event code 0x0051
May 20 13:26:21 vmkernel: 0:00:22:50.694 cpu12:4332)<6>megasas_hotplug_work[4]: aen registered
May 20 13:26:21 vmkernel: 0:00:22:50.694 cpu5:4101)<6>megasas_service_aen[4]: aen received
```

```
May 20 13:26:21 vmkernel: 0:00:22:50.694 cpu14:4328)<6>megasas_hotplug_work[4]: event code 0x00f9
```

Upgrading Firmware on your TRCs

When running virtualized, UC applications do not manage the firmware on the physical server (host).

The customer must manage the firwmare manually. The customer must monitor new releases of firmware published by the hardware vendor and upgrade when necessary based upon the recommendations of the hardware vendor and VMware.

For deployments on Cisco UC hardware, instructions for upgrading the firmware are available in the Release Notes for each version of posted firmware. It is important to understand that the upgrade procedure may vary between releases of firmware. For example, see **firmware release for the UCS C-series** for details on the rackmount servers.

For installation and configuration information on UCS servers, refer to the documentation roadmap for either the B-Series or C-Series servers:

- **Cisco UCS B-Series Servers Documentation Roadmap**
- **Cisco UCS C-Series Servers Documentation Roadmap**
- **Cisco UCS C-Series Integrated Management Controller Documentation**
- **Cisco UCS Manager Documentation**

When deploying on Cisco UCS servers, you can sign up for notifications of new releases at **http://www.cisco.com/cisco/support/notifications.html**.

Backup, Restore, and Server Recovery

Disaster recovery for Cisco Unified Communications application virtual machines supports the same in-host techniques as Cisco Unified Communications applications on physical servers: the same backup options are available with Cisco Unified Communications running on ESXi as on physical servers.

Other backup, restore techniques available in a virtualizaed environment are currently not supported.

See UC Applications-Specific Virtualization Information for information specific to each UC application.

Back to: Unified Communications in a Virtualized Environment