

Out of Band Management: an overview

Travis Campbell

hcoyote@gmail.com

Who am I?

- ✱ Senior Systems Engineer at a large HPC shop
- ✱ Working on lights-out datacenters for 6 years
- ✱ Haven't actually had to go into a data center for about 6 months now (yay OOB management!)

What is it?

- ✱ Dedicated management interface to systems
- ✱ Sometimes integrated by vendor
- ✱ Sometimes supplied by third party equipment
- ✱ AKA “Lights Out Management”

What does it do?

- * Remote power control
- * Remote serial or KVM access
- * Hardware monitoring and event notifications
- * Hardware inventory

Multiple Types

- ✱ Vendor specific (HP iLO, Dell DRAC, Sun ALOM)
- ✱ Industry standard (IPMI v1.5, IPMI v2.0)
- ✱ Third party add-ons (APC Remote PDUs, Cyclades/Avocent console access servers, etc.)

Why do you want this?

- ✱ Biggest reason: laziness.
- ✱ Reduce unnecessary trips to the data center
- ✱ Respond more quickly to physical system outages
- ✱ Crash carts suck and data centers are cold
- ✱ Fewer admins needed to handle large numbers of systems

Problems?

- * TIMTOWTDI (There is more than one way to do it)
- * It's standard! ... but which standard?
- * Every type has it's quirks
- * Sometimes you still have to go down and push buttons or hookup a crashcart.

IPMI

- * “Intelligent Platform Management Interface”
- * Open standard.
- * Only defines structure and format of interfaces, implementations vary
- * <http://www.intel.com/design/servers/ipmi>

IPMI 1.0

- * First stab at doing a standard
- * Had basic FRU information for the hardware
- * Firmware logs to track hardware events
- * No real remote management stuff
- * IPMI 1.5 way of the future!

IPMI 1.5

- ✱ New features to IPMI
 - ✱ IPMI over LAN, serial port sharing
 - ✱ Alerts via SNMP
 - ✱ New support for authenticating users and defining roles or privileges

IPMI 2.0

- * Even better
- * Serial over LAN
- * Remote power control
- * SMBUS access to query hardware sensors

IPMI tools

- ✱ Ipmitool
- ✱ Ipmiutil
- ✱ Openipmi (seems to be the most common one)
- ✱ GNU freeipmi
- ✱ See <http://en.wikipedia.org/wiki/IPMI>

Hewlett Packard

- * HP has iLO (“Integrated Lights Out”)
- * Three types of iLO
 - * iLO 2 (Good)
 - * iLO (Bad)
 - * Io100 (Ugly!)

HP iLO

- ✱ Generally in use on older mid- to high-end Proliant gear (300 series and higher)
- ✱ Standard features: Remote serial access, power control, scriptable via XML interface
- ✱ Advanced (licensed): remote graphical console, Windows Terminal Services access, virtual media
- ✱ Uses SMASH and DMTF CLP command line interface standards

HP iLO, page 2

- * Kind of quirky with Linux
 - * Use “REMCONS” to view bootup sequence
 - * Use “VSP” to have a usable serial console
- * Sometimes wouldn't power down systems, requiring hard power reset of system (ie, pull the plug)
- * iLO sometimes wedged, requiring hard power reset

HP iLO 2

- ✱ Next generation; actually works pretty well
- ✱ Better health monitoring, non-graphical KVM support, remote serial, shared network capability, power control and monitoring, still scriptable
- ✱ Still requires license for graphical KVM and virtual media (blade version includes this license)
- ✱ Uses DMTF CLP command line interface standards

HP iLO 2, page 2

- ✱ No real complaints
- ✱ REMCONS is no more; VSP actually works for full bootup sequence. Yay!
- ✱ Haven't encountered any power issues yet

HP Io100

- * “Lights Out 100”
- * Used on low-end Proliant 100 series.
- * IPMI based. It’s crap!
- * Remote serial, remote power. That’s about it.

HP iLO Scripting

- * HP provides a toolkit for managing and configuring the iLO series.
- * Reasonably well documented.
- * XML-based.
- * Can be used remotely over SSL connection or “in-band” with hponcfg util from HP Proliant Support Pack (if right kernel drivers loaded)

HP Scripting, page 2

- ✱ Download XML files: <http://tinyurl.com/6mdlh9>
- ✱ Documentation: <http://h18013.www1.hp.com/products/servers/management/ilo/documentation.html>

HP, Other info

- ✱ If you're working with Proliant gear, do yourself a favor and download (and install) the Proliant Support Pack. It's free!
 - ✱ <http://h18004.www1.hp.com/products/servers/management/psp/index.html>
- ✱ Integrates very well with Linux and Windows, as well as HP Systems Insight Manager.
- ✱ Provides better support for monitoring and management.

Sun Microsystems

- * Sun has two lights out management systems
- * ALOM (“Advanced Lights Out Management”)
- * iLOM (“Integrated Lights Out Management”)

Sun ALOM

- * Primarily for UltraSPARC-based hardware (Netra and Sun Fire V100/200/400 series)
- * Remote power control, monitoring, console access. Much like other LOMs.
- * Documentation: <http://docs.sun.com/source/819-2445-11/index.html>
- * (I don't have much experience with it)

Sun ILOM

- * The new hotness
- * Found on x64 and Coolthreads SPARC servers
- * IPMI 2.0 based
- * Remote KVM, serial access, remote power, virtual media support built in
- * Also has monitoring capabilities, fault management, and inventory

Sun ILOM, page 2

- * Because it's IPMI based, you can use standard ipmitool or OpenIPMI releases from sourceforge to interact
- * Integrates with Sun N1 Manager
- * Uses DMTF CLP command line interface standards
- * Documentation: <http://docs.sun.com/source/820-1188-11/>

DMTF CLP

- ✱ Distributed Management Task Force - Command Line Protocol (part of SMASH).
- ✱ DMTF is a working group that creates cross-platform standards for managing enterprise and Internet environments.
- ✱ Other standards: WBEM, CIM, SMBIOS
- ✱ <http://www.dmtf.org/standards/mgmt/smash/>

Third Party Stuff

- ✱ If your hardware doesn't support this stuff, you can “do it yourself”
- ✱ Serial console access via things like Cyclades ACS console servers or APC console servers
- ✱ Power control via Cyclades PM PDUs or APC Rack Switched PDUs

Cyclades/Avocent

- ✱ About half of our environment is Cyclades console servers.
- ✱ We don't use the Cyclades PM PDUs because they're not network-aware (last we looked).
- ✱ Terminal servers provide a web or ssh-based connection to a serial console.
- ✱ Easy to configure, no real problems. Kinda pricey.
- ✱ Maybe 3-4 firmware failures in 6 years of using them across hundreds of console servers

APC

- * Using APC PDUs pretty much everywhere, even when we have built in IPMI/ILO/etc.
- * Used for last ditch attempt to control power
- * Nifty SNMP access
- * We use cacti to graph our power usage
- * If management unit fails, PDU still “runs”, defaults to power on in our experience.

APC, page 2

- ✱ Haven't played with any of APC's other management gear to control all this.
- ✱ Built our own interface between the PDUs and our host database to manage controlling power over SNMP.
- ✱ Gotchas: configure your PDUs to turn ports on in sequence or you'll flip breakers when turning on all systems at once.
- ✱ We've used APC 7911 (208v 16 port 30A), 7901 (110v 8port 20A), and 7941 (208v 21port 30A zero-U) pretty extensively.

APC, page 3

- ✱ Cacti APC files: <http://forums.cacti.net/about7612.html>
- ✱ APC SNMP MIB: <http://tinyurl.com/3tu9mv>

Conserver

- ✱ Multi-user application that helps manage serial consoles.
- ✱ Awesome tool, must have for integrating your various OOB console tools into one common interface.
- ✱ Does logging, authentication, multi-user console access for watching what others are doing (such as mentor/mentee work)

Conserver, page 2

- * Scales reasonably well with some tweaks.
- * Default build will work up to a few hundred consoles being controlled.
- * Have to tweak number of consoles controlled per daemon if you scale up higher.
- * We run about 1500-2000 connections per conserver master.

Conserver, page 3

- ✱ Multi-master configuration so you can make all masters aware of each other.
- ✱ Conserver client will redirect to right master automatically if all are aware of each other.
- ✱ Collectively have tens of thousands of consoles managed worldwide with one interface to them all
- ✱ Makes cross-site staff training easy! One method to rule them all.

Conserver, Final

✱ <http://www.conserver.com/>

Questions?