Future Console Servers

devproj project #31

Overview

- Requirements / motivation
- Current approach
- Possible future options
  - KVM over IP
  - IPMI
  - Serial concentrator cards
  - Commodity solutions
- Summary / recommendations

Requirements

- Ideally, want a simple and inexpensive solution that allows us remotely to:
  - look at serial console output (even for dead/locked/unresponsive machines)
  - examine and set BIOS/bootprom values
  - do machine installations
  - power cycle machines
- Ideally, one such solution per 'bank' of racks; capable of handling up to ~80 machines

Motivation

- New building
- Some of the hardware in use in the current system is no longer obtainable
- The current approach requires a lot of serial cabling – can we rationalise this?
- Are there better/cheaper solutions?

Current Approach

- Six console servers fitted with 32-way serial cards, using conserver (`www.conserver.com`), configured via LCFG
  - ```
    ssh consoles
    console karajan
    Enter idurkacz@exeter.inf.ed.ac.uk's password
    ...
    karajan.inf.ed.ac.uk login:
    ```
- Each console server looks after its own client machines, but the entire system is federated
- *But:* Cyclades serial cards we use are no longer obtainable!
    ```

Option 1: KVM over IP

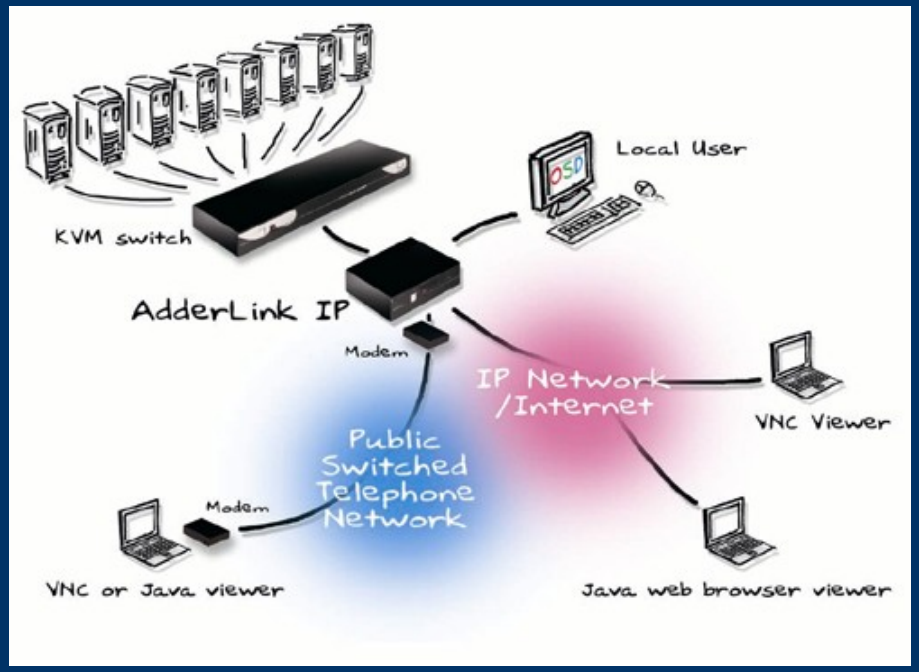

- AdderLink IP:
  - Multiple clients via KVM switch
  - Up to four user connections
  - Needs only one IP address per unit
- But:
  - No buffering
  - Messy cabling
  - Expensive

KVM over IP (cont)

- Lantronix Spider:
  - One client per unit
  - Multiple user connections
  - Neat cabling
- But:
  - Needs one extra IP address per target machine
  - Still no buffering
  - Still expensive

Option 2: IPMI

- IPMI – 'Intelligent Platform Management Interface'
- Defines a 'standardised message-based interface to intelligent platform management hardware.'
- Provides a standard interface to:
  - Environmental sensors (temperature, voltage)
  - Power control
  - Event logs
- Implemented by the Baseboard Management Controller (BMC)

IPMI (cont)

- v1.0 (1999)
  - Access to the BMC via system buses only
- v1.5 (2002)
  - Access to the BMC via serial port and LAN – the BMC has its own IP address
- v2.0 (2004)
  - Support for encrypted network traffic
  - Serial-over-Lan (SOL) – redirection of monitored system's serial port over the BMC's LAN connection

IPMI v1.5

- Supported by 8th generation Dell servers – e.g. PowerEdge 850, 860 and SC1425
- SOL *is* implemented by some vendors, but is *not* standardised. Requires auxiliary s/w on the client. Not usable in practice
- Still useful though! – e.g. remote power off/on, get sensor data, etc.:
- `ipmitool -I lan -H splitbmc -U root power off`
  `ipmitool [...snip...] sensor get Temp`

IPMI v2.0

- Supported by 9th generation Dell servers e.g. PowerEdge 1950, 2950; SC1435
- SOL *is* standardised:
- `ipmitool -I lanplus -H pastabmc -U root sol activate`
- Note that:
  - Still need to use `conserver` to get buffering of console output
  - Need an additional IP address per machine (but this could be on a management VLAN)

Option 3: Serial concentrator cards

- The current approach
- Requires:
  - Multi-port serial cards and drivers
  - One LCFG-configured server per 32 (or so) target machines
- Cyclades serial cards are no longer obtainable, but Perle cards are

Option 4: Commodity solutions

- Drop-in 'boxed' solutions that are functionally similar to the current `conserver`-based set-up:
  - Connected to target machines via serial ports and cables (generally: 16, 32 or 48 ports)
  - Remote access to target consoles via ssh
  - Provide buffering
  - Various authentication possibilities (local DB, RADIUS, Kerberos, etc.)
- Vendors: Avocent, Lantronix, ...

Commodity solutions (cont)

ssh access to consoles

```
ssh kbslc
[kbslc1]> connect direct deviceport 1
Fedora Core release 5 (Bordeaux)
Kernel 2.6.18-1.2257_FC5_dice_1.2 on an i686

selidor.inf.ed.ac.uk login:
```

– or –

```
ssh kbslc -p 3001
Fedora Core release 5 (Bordeaux)
Kernel 2.6.18-1.2257_FC5_dice_1.2 on an i686

selidor.inf.ed.ac.uk login:
```

Commodity solutions (cont)

- *In principle*, provide a low maintenance and cheap solution, but they:
  - Need some thought regarding integration (e.g. can we make it easy to locate the console of interest?)
  - May require everybody to change habits
  - Don't necessarily scale up neatly (i.e. it may not be possible to federate multiple boxes)
  - Don't address cabling concerns

Recommendations

- *In the future:*
  - IPMI SOL (v2.0 or above) + `conserver`
- *Now: either*
  - Continue the current scheme, *or*
  - Use commodity boxes

Open questions

- How many of our machines need remote console access?
- Of these, how many support IPMI v2.0?
- How independent of the rest of our infrastructure (DHCP, authentication etc.) should any solution be?
- How successful in practice are the commodity boxes?

What are (some) others doing?

- *Ed Uni IS:* “FMD UNIX are the only team to have remote console/lights out management ... Primarily all our kit is SUN ... we use serial management – primarily CISCO 3745 with 128 serial lines.”
- *A ShefLUG contact:* “We use a dual solution of IPMI for fan/temp monitoring and power control, plus KVMoIP+virtual media for remote access. It's highly reliable and wonderful to work with, but not cheap!”
- *From the beowulf.org list:* “... get boxes which support IPMI-2.0 ...”