Liverpool HEP Computing

SITE INFRASTRUCTURE

CURRENT PROBLEMS
AND
PROPOSED SOLUTIONS

Cooling

- Air condition and water cooling units that need regular manual checking of status.
  - Time is wasted performing manual checks. No metrics to locate hot spots or degrading units. No alarms on failure.
- Enclosed rack units that need regular manual checking of temperature, unclean shutdowns.
  - Time is wasted performing manual checks. Equipment at risk of damage from overheating or sudden power loss. Current rack sensors proprietary and essentially unsupported.
- Non-optimal rack layout (no hot/cold aisles, bad use of space).
  - Hot exhaust air is mixing with cold intake air, lots of hot spots requiring fans, cooling units far from where cooling is needed, hard to add equipment without increasing problems.
- Aging, overloaded equipment and badly situated
  - High failure rates, reduced capacity. High load also reduces life span. Hard to get heat load to all units: some idle whilst others overloaded.

Cooling - Problems

The information is there, but there is no way of getting it or generating alerts from it other than by manually going to the cluster room to check.

Some units run hot (25-30C) while others are sitting idle.

Continual round of hardware problems.

Cooling - Problems

In-house temperature monitoring of water-cooled racks. Again needs manual checks. Running on aging hardware, no formal support, no clean shutdowns.

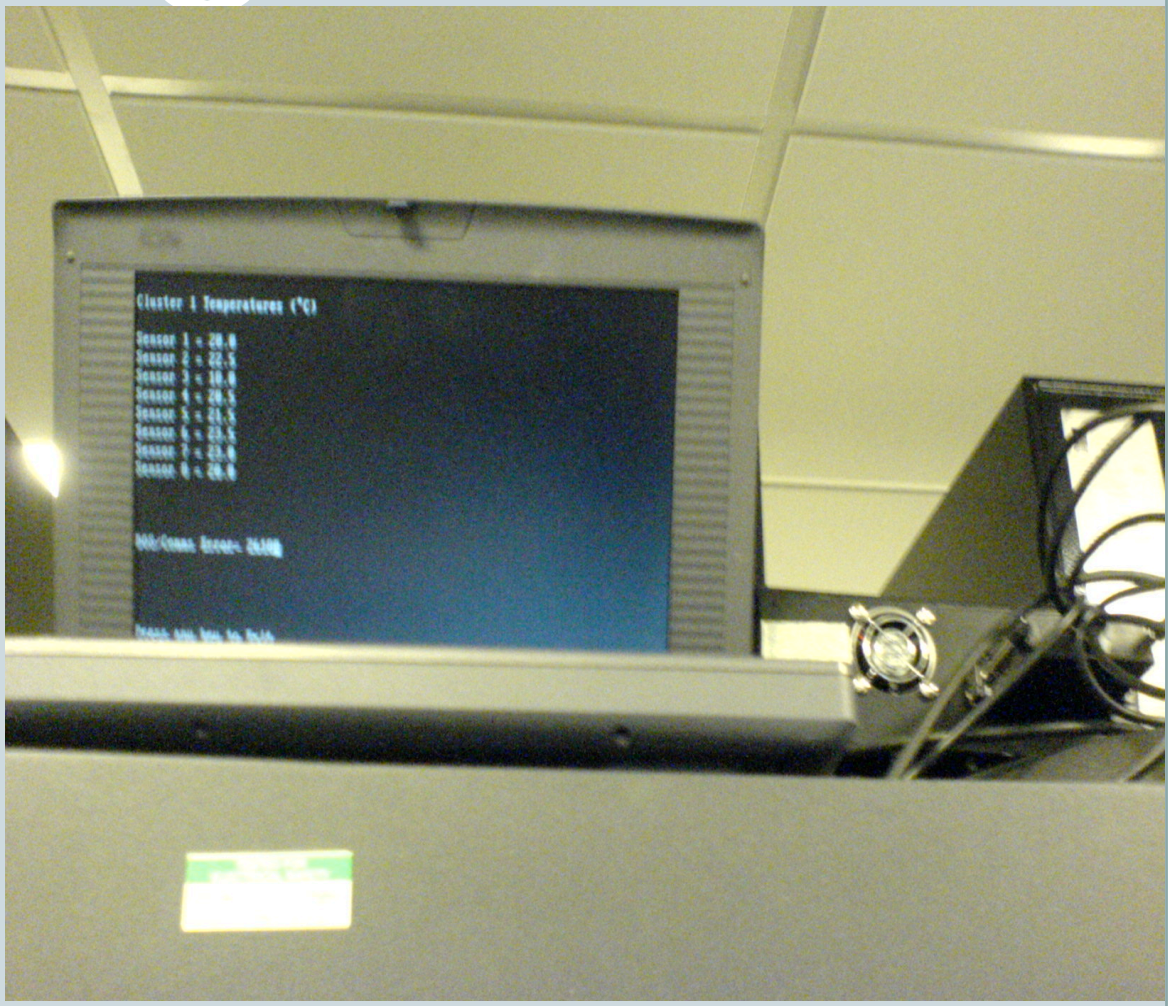

Cooling – Makeshift solutions

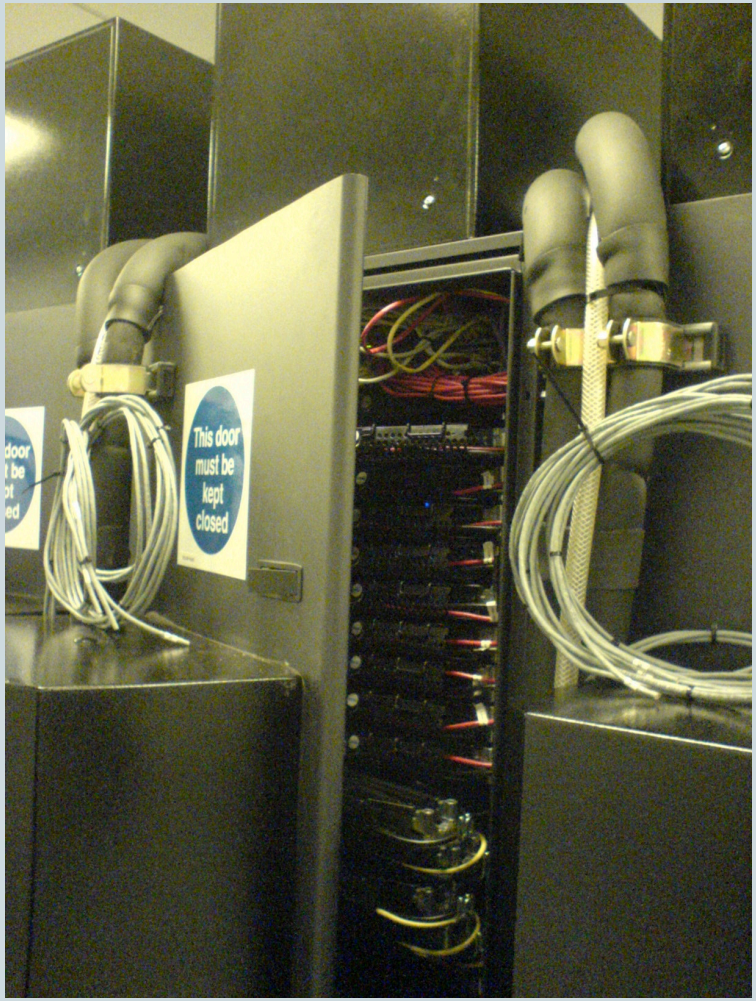

Cluster Room Layout

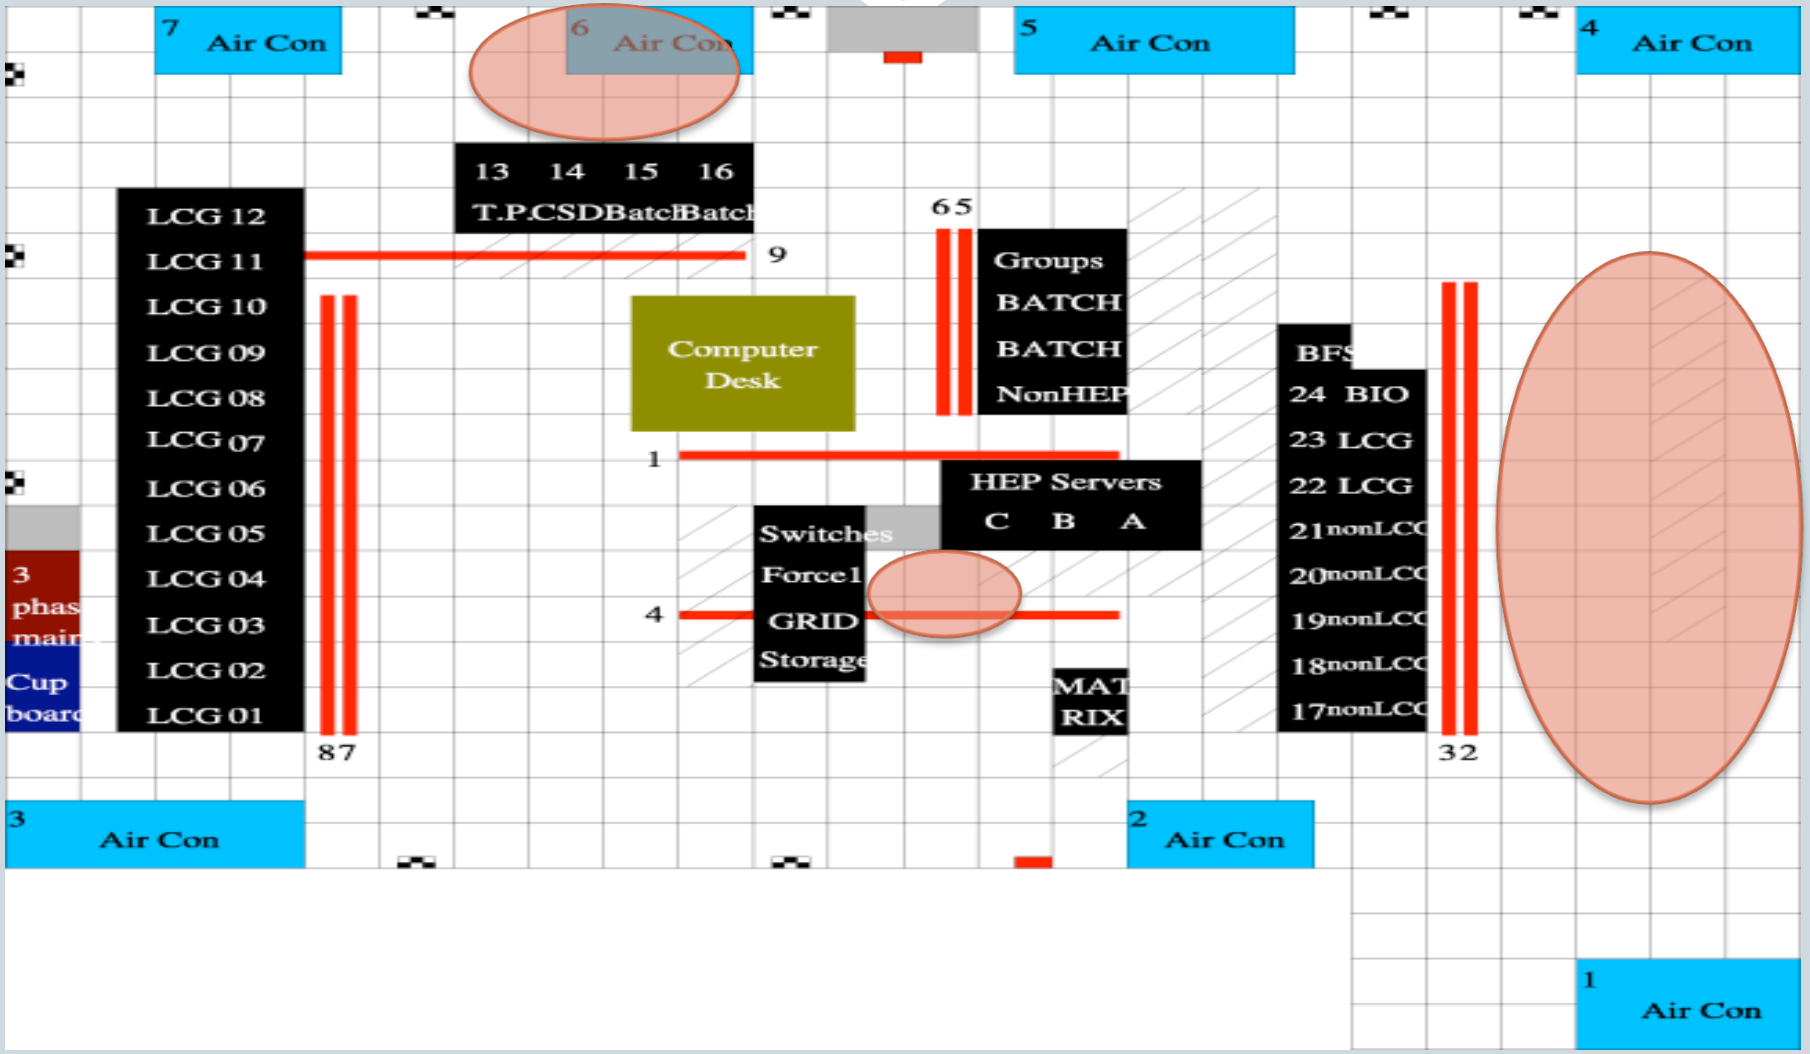

Power

- All Worker Nodes and many servers have no UPS backup of any kind
  - Potential damage from unclean power supply, in event of any power loss all running jobs are lost, unclean OS and hardware state on nodes.
- Large number of Worker Nodes
  - More hdds, fans, switches, motherboards per job, more power consumption, more space and cooling required.
- No remote electrical power monitoring/control
  - Unsure of peak and average loads or distribution, unable to hard reboot any crashed systems remotely.
- Unsupported fuse board equipment
  - Expect high failure rates, long lead times on repairs.

Maintenance

- All Worker Nodes and many servers have no KVM support.
  - Wastes time and unsafe using ‘trolleys’ with extensions, damages chassis connectors, many systems now don’t work through front connectors, cannot easily use keyboard and USB stick at same time.
- Large number of Worker Nodes.
  - Increased probability of system or component failure, higher turnover of components, increased load and latency on monitoring systems.
- Unreliable network booting of Worker Nodes.
  - Cannot use automatic installation methods seamlessly, wastes lots of time.
- Cluster room steps
  - Difficult to move heavy and bulky equipment in or out of cluster room.

Maintenance - Problems

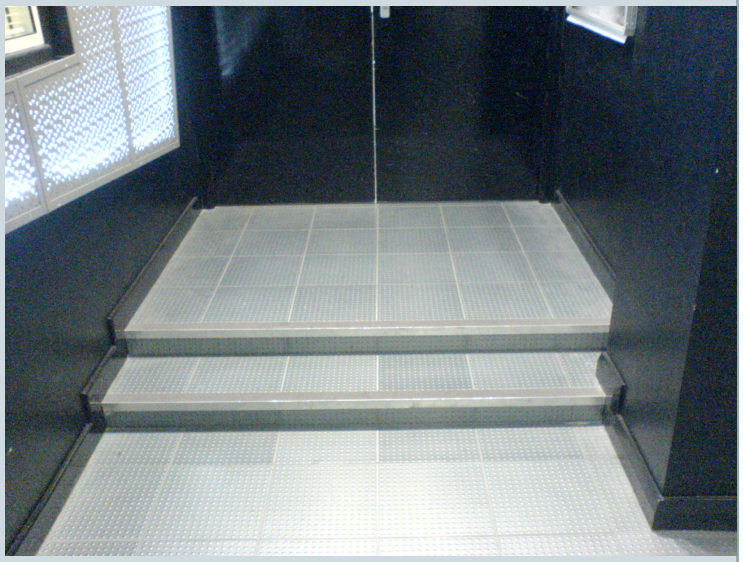

Maintenance – Manpower

- Perennial problem of not enough staff
- Low priority tasks build up
  - To Do lists always get longer, pile of broken nodes gets higher
- Projects take longer to complete
- No time for staff development, research
- Currently down to *two* staff members to cover grid middleware and hardware, local batch farms, helpdesk, cleanroom upgrades, security, supplies, hardware commission and system planning
  - New grid admin by September? Needs training and acclimatisation. Maybe taking up some slack by Christmas.

Network

- Current cluster backbone acceptable through high density Force10 switch
  - But no 10Gbit capability, requires extensive wiring.
- University uplink limited to 1Gbit
  - See limit of ~300Mbit during LCG tests, recent multi-site tests show Liverpool is severely limited somewhere.
  - http://pprc.qmul.ac.uk/~lloyd/gridpp/nettest.html
- Building network very old, badly wired, undocumented and unmanaged
  - Wastes time tracing problems, very little monitoring capability, shoddy wiring is unreliable and limits bandwidth
  - Office connections at high risk and many multiple points of failure.

Network - problems

These undocumented, badly made, easily dislodged wires keep most of the HEP offices and clean rooms connected to the HEP servers and internet.

We don't know exactly where the other ends are or how they get there.

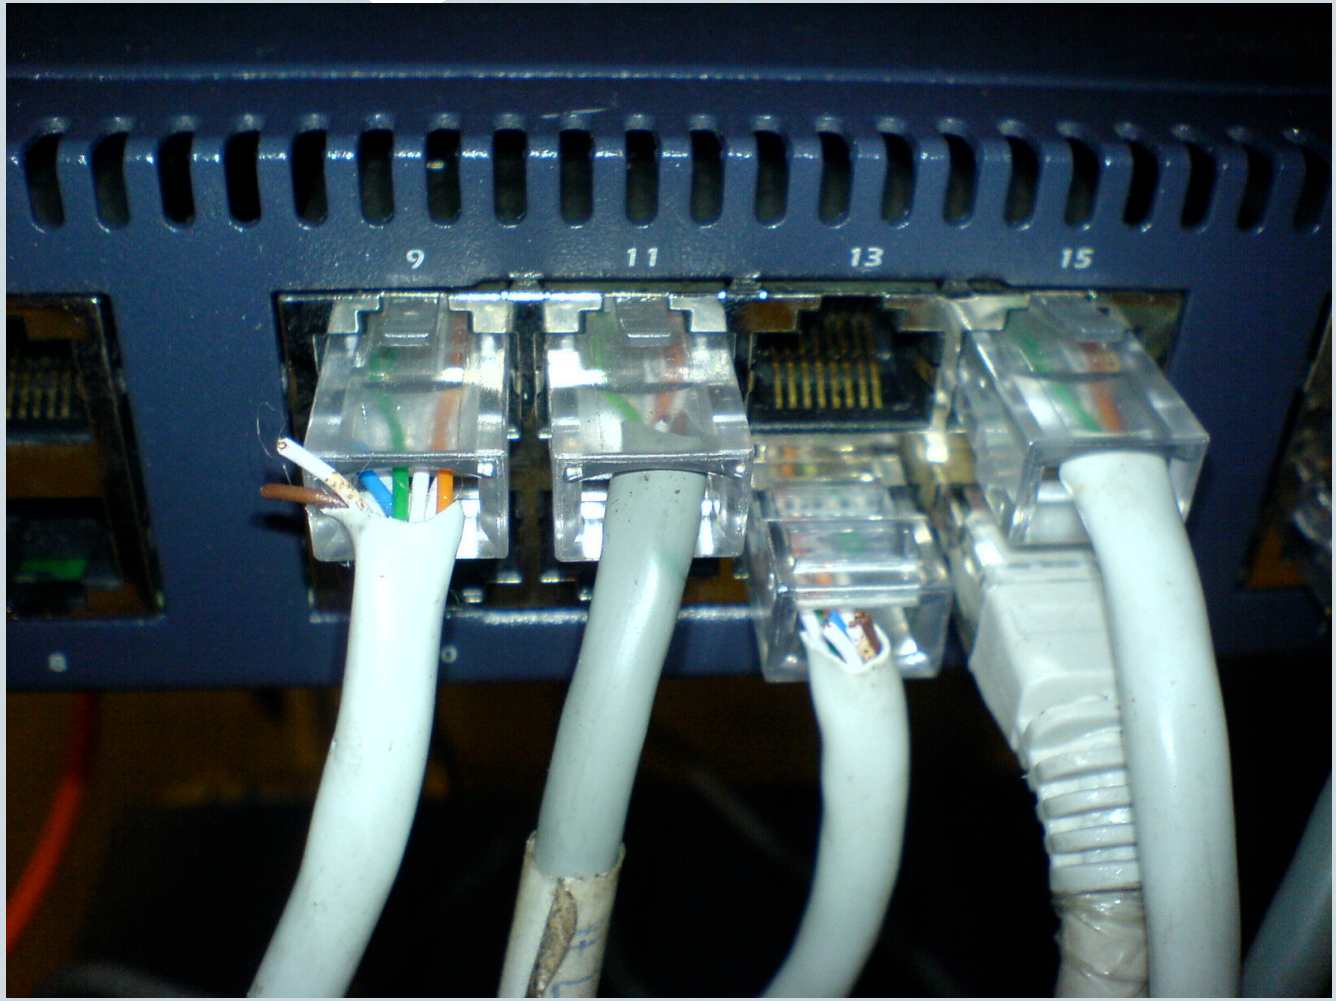

Network - success

High availability, high performance gigabit network backbone through Force10 switch.

But required extensive recabling, still more cabling to be added.

For cross-campus links and high performance storage, fibre networking would be a good but expensive addition (aggregate links aren't a perfect answer).

Rack and server VLANs working well, reduced broadcast traffic, more manageable routing, isolates problems and good for testing.

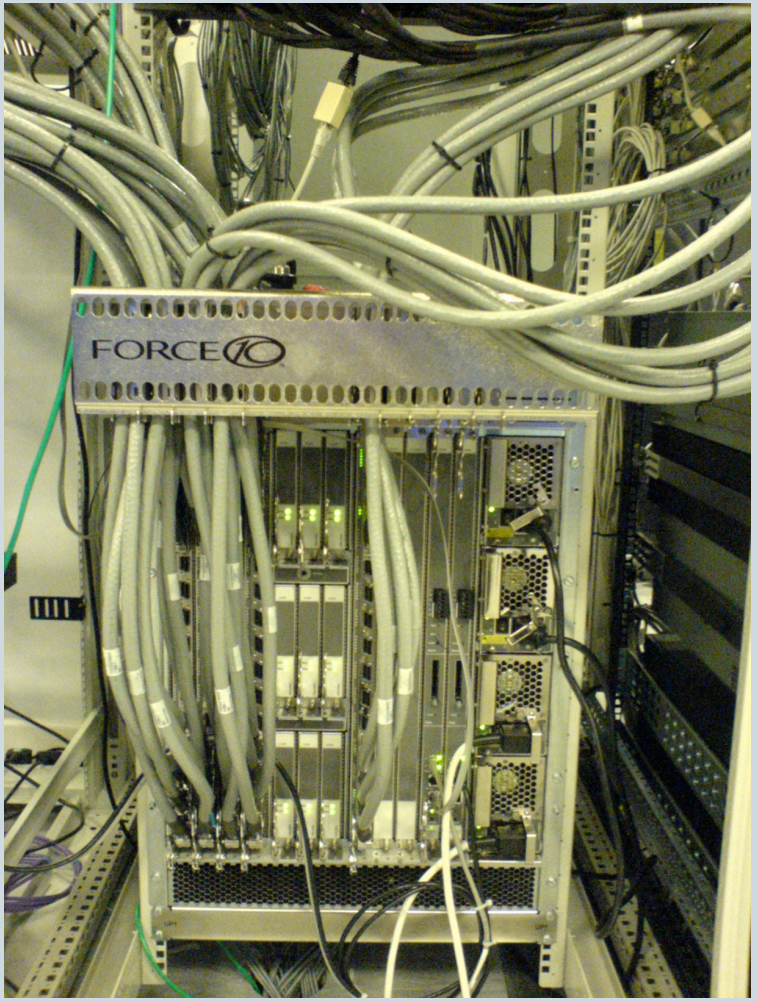

Conclusions - Solutions

- Integrated systems
  - Racks, cooling, UPS, KVM, environmental monitoring as a total, supported package. Ideally self-contained to prevent disruption while replacing or moving old equipment.
- Monitored systems
  - Any metric must be remotely accessible directly or through data logger: Temperature, humidity, PDU, UPS, chiller status.
- Remote systems
  - Remove the need to physically intervene for regular maintenance, better working environment, off-site and out-of-hours support, faster response.
- High density systems
  - Multiple cores per node, virtualised where necessary. Low energy CPUs: more bang per buck per Watt.

Solutions - Examples

- KVM over IP
  - Remote access via web browser to KVM
  - Can be KVM or server module
  - Security implications
  - Ultimate in server management
    - Control all machines from one desktop!
    - Control right down to BIOS level
    - With virtual CDs can even install remotely

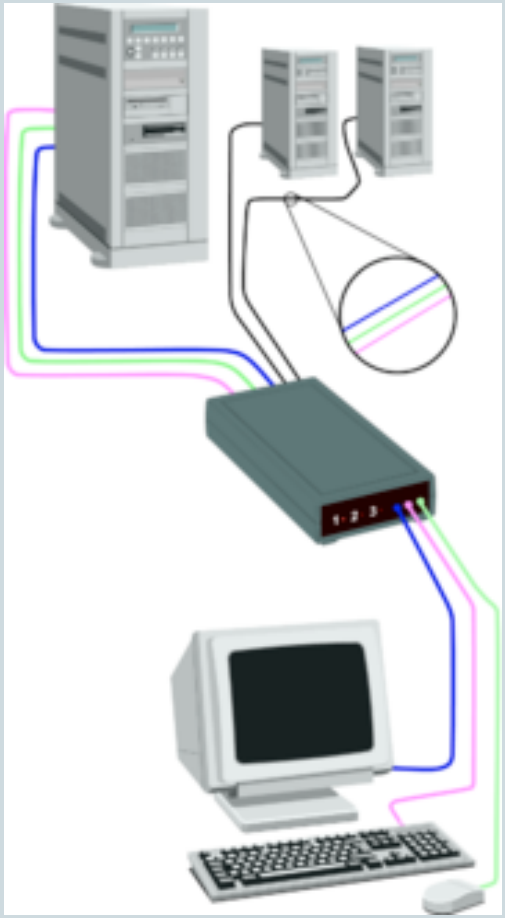

Solutions - Examples

- APC InfraStruXure
  - Self contained, high efficiency cooling
  - Integrated UPS
  - Monitorable and remote controlled PDUs
  - Modular, can easily add extra capacity
  - Doesn't require raised floors
  - Flexible in location
  - Bit pricey!

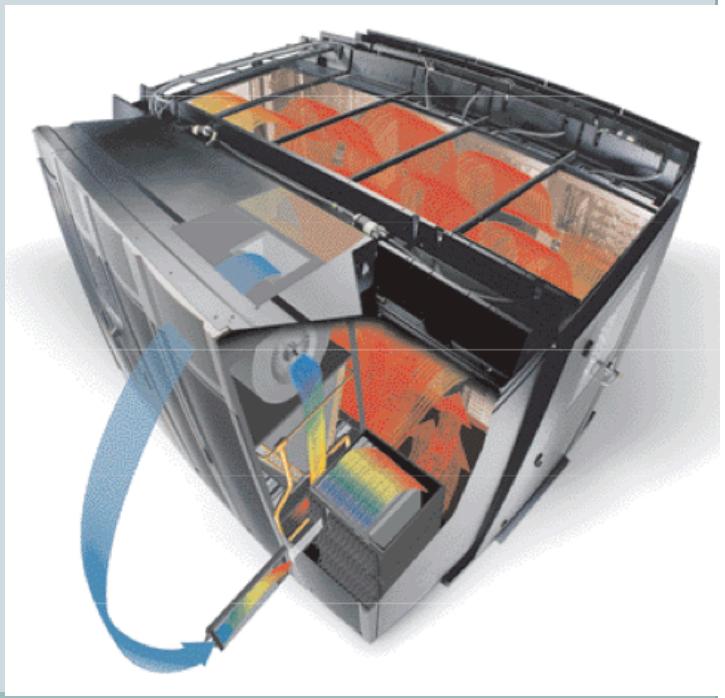

Solutions - Examples

- Building network rewire
  - Gigabit-rated cables
  - Managed switches
  - Fully documented
  - Easy access for maintenance
  - Redundant trunk cables
- Go from this:

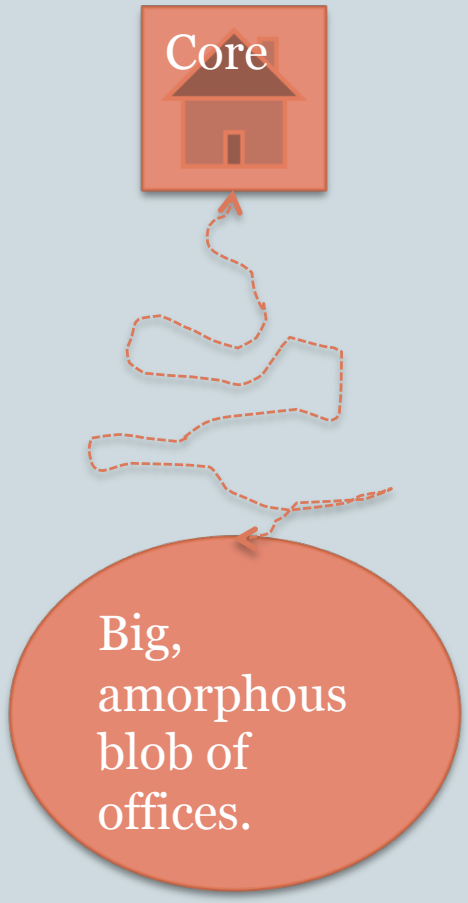

Solutions - Examples

- To this: