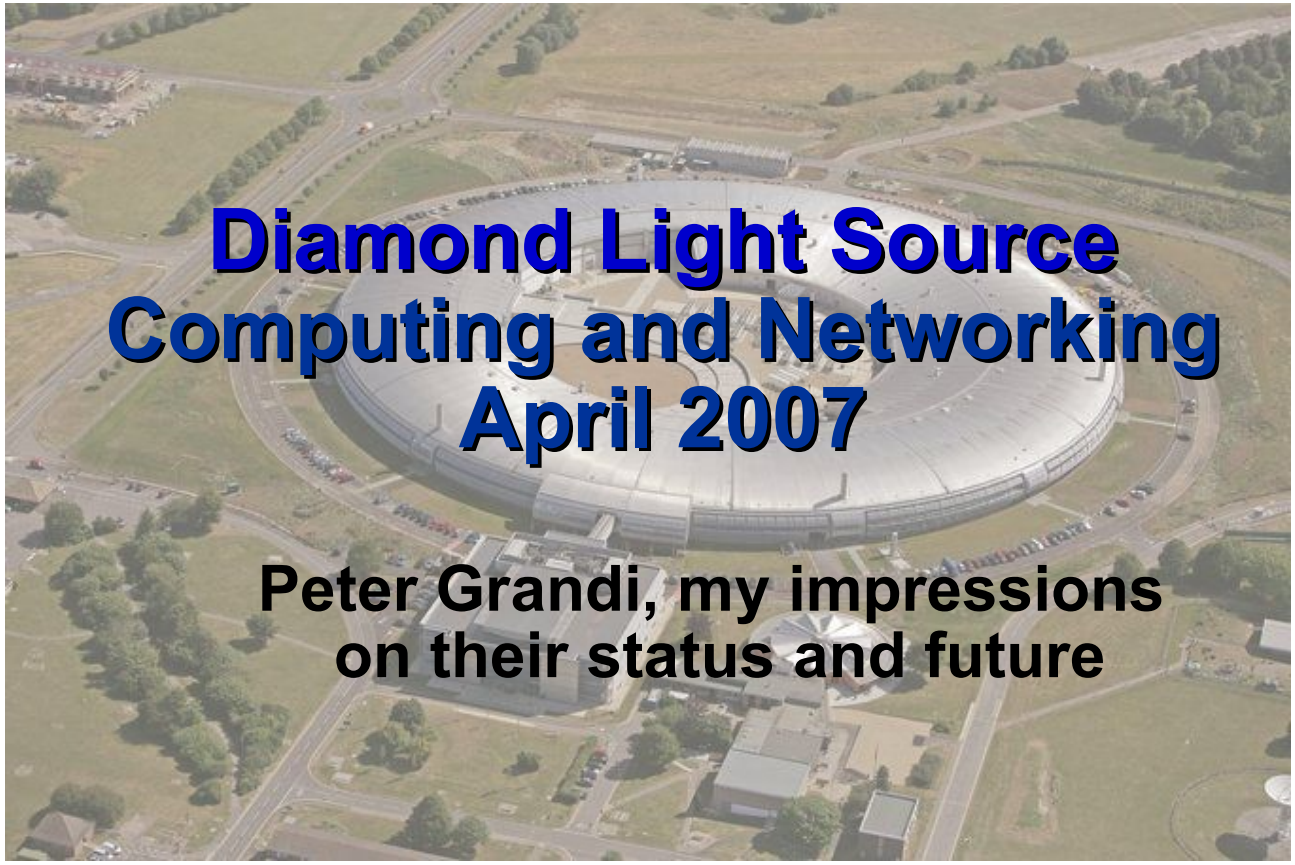

Diamond Light Source Computing and Networking April 2007

Peter Grandi, my impressions on their status and future

diamond

Context

- **Diamond Light Source themes:**
  - **Production: support researchers.**
  - **Research: push the boundaries.**
  - **Very new, building cointinues for many years.**
- **A factory and a laboratory for the beam, but for computing and networking too.**

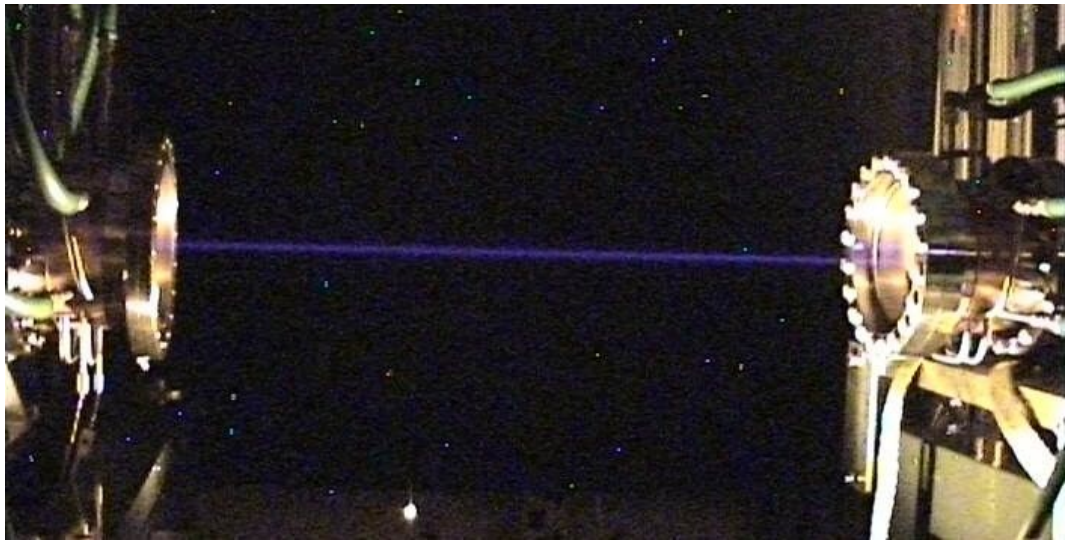

diamond

Computing and networking

- Three separate organisations:
  - *Technical*: light source control and tuning.
  - *Science*: beamline control, tuning, DAQ.
  - *Business*: office and administrative applications.
- Three separate infrastructures, but:
  - Business IT also does cross divisional MS-Windows support.
  - Science IT also does office GNU/Linux second line support.

diamond

Technical goals

- **Without networking and computing the Machine continues to work for a while, if no other accidents.**
- **Major requirements**
  - **high light source availability;**
  - **high event frequencies.**
- **Non-requirements:**
  - **growth in time;**
  - **computing power;**
  - **bandwidth.**
- **Context:**
  - **One organisation, stable technology.**
  - **24x7 operation during beamtime.**
  - **Over a 560m ring.**

diamond

Technical architecture

- **Guiding ideas:**
  - **time to diagnose and repair instead of redundancy;**
  - **stability more important than performance;**
  - **use of established standards;**
  - **long term outlook.**
- **Guiding principles:**
  - **simplicity of design, simplicity of implementation;**
  - **advanced but consolidated technology;**
  - **reuse of existing commodity components.**

diamond

Technical solutions

- Central computer room plus 24 local ones
- Networking:
  - OM3 (multimode 50/125) fibre infrastructure, LC sockets;
  - CAT6 LSZH copper local wiring;
  - 1000BASE-SX/-T Ethernet;
  - 1 level tree topology for both fibre and Ethernet;
  - two separate parallel networks for critical and less critical functions;
  - quality and simple 3com switches with no configuration.
- Computing:
  - x86 GNU/Linux based servers. PowerPC VxWorks based controllers, ARM GNU/Linux based monitors;
  - C EPICS distributed control software;
  - Python GUI.
- 1 FTE for both networking and computing.

diamond

Science goals 1

- **Without network and computing beamlines just don't work.**
- **Major requirements**
  - **minimize all-beamlines downtime;**
  - **handle operations and construction at the same time;**
  - **extremely diverse science setups;**
  - **very high, ever increasing bandwidth and computing power needs;**
  - **low cost, especially in manpower.**
- **Non requirements:**
  - **always-on connectivity outside the beamline;**
  - **full avoidance of downtime for single beamlines (rare, up to 2 hours downtimes per single beamline seem acceptable);**
  - **totally uniform solutions.**

diamond

Science goals 2

- **Context:**
  - **multiple organisations, unstable technology;**
  - **some experiments destroy their samples;**
  - **24x7 operation during beamtime, 9-to-5 monitoring outside beamtime;**
  - **over 60m long, very packed, beamlines sited in a ring of 250m in diameter**
  - **3,000-5,000 user community doing 3,000 to 6,000 hours of experiment a year per beamline (35%-70% of the year).**
- **Data acquisition demands:**
  - **single beamlines can be a challenge, but the big deal is that there are many;**
  - **3 MX beamlines at 100MB/s 24x7.**
  - **some beamlines at 100MB-400MB/s for 1-2 days;**
  - **some beamlines at 20-40MB/s for several days;**
  - **aggregate demand not far off LHC range (500GB-1TB/s).**

diamond

Science architecture

- Large changes planned: from an office-style, 9-to-5 mindset to a 24x7, production one.
- Guiding ideas:
  - time to diagnose and repair but also some redundancy;
  - performance and flexibility more important than stability;
  - use of established standards;
  - medium term outlook.
- Guiding principles:
  - simplicity of design;
  - whatever-it-takes implementation, aiming for simplicity;
  - advanced technology, even if not yet common;
  - use of very recent components.

diamond

Science solutions 1

- **Per-beamline computer and network resources plus shared computer and network room.**
- **Networking:**
  - **OS1 (singlemode 9/125) fibre infrastructure, LC connectors;**
  - **CAT6 copper local wiring;**
  - **10GBASE-LR and 1000BASE-LX/-T Ethernet;**
  - **1 level tree topology for both fibre and Ethernet;**
  - **advanced Nortel routing switches, stacked for both capacity (servers) and hot sparing (beamlines) with minimal configuration.**

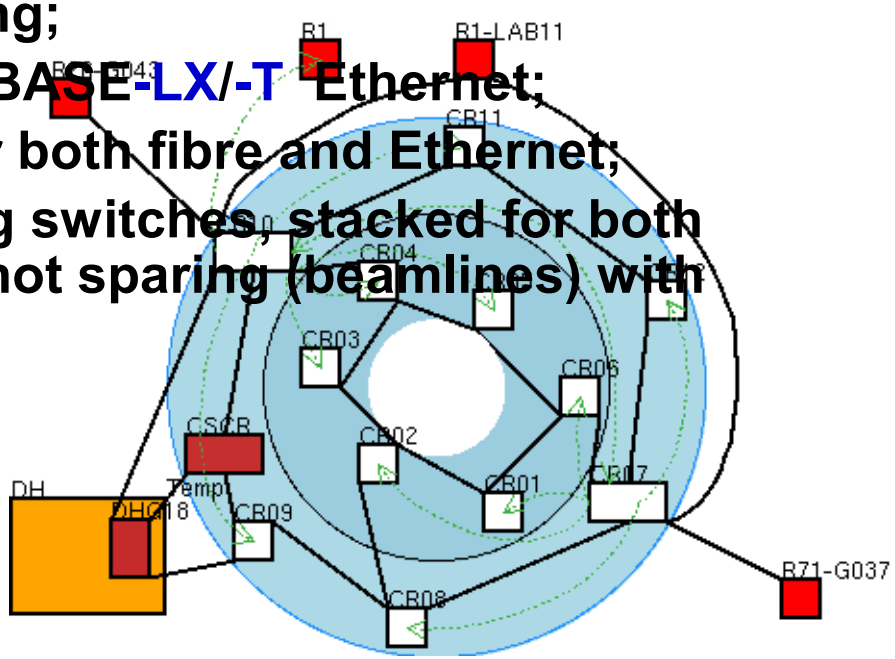

diamond

Science solutions 2

- **Computing:**
  - single sign on using Active Directory and Vintela;
  - x86 GNU/Linux based servers (Dell) and clusters (IBM) with 1gb/s Ethernet (some Infiniband);
  - x86 GNU/Linux and MS Windows based (Dell) workstations with 1gb/s Ethernet;
  - beamline local storage servers (5GB-20GB);
  - shared storage servers (20GB-180GB);
  - streaming of data to vast (petabyte) on-campus archive;
  - C EPICS for some low level control functions;
  - Java GDA with Jython scripting for detector control and GUI;
  - data delivery for now on DVD or removable hard disk, in the near future local processing or fetched over the network from on-campus archive;
  - 3 FTEs.

diamond

Science lessons

- **Physical infrastructure:**
  - Cabling a lot more work than desired.
  - Racks a lot more work than desired.
  - A plant manager probably a good idea (not me!).
  - Singlemode fibre only sensible choice.
- **Ethernet:**
  - Single global VLANs seductive, too much trouble.
  - -LX and -LR cost a bit more, but site is big and new.
- **Computing**
  - Don't know yet. So far we get 25MB/s write rates with RH EL4, ext3 and EMC Clariion RAID3. Not very good news. We need at least simultaneous 100MB/s write and read.
  - Power and cooling rather more trouble than desired.
- **Overall: science users like to spend money on physics; redundancy overrated.**

diamond

Science future

- **Issues:**
  - ever more beamlines on the same infrastructure;
  - detectors improve constantly: talk of 700MB/s detectors.
  - conflict between growing data requirements and use of RH EL4 (no XFS!) and RAID3 or RAID5 and NFS.
- **Upgrades:**
  - ensure support for jumbo frames;
  - what to do for large storage (fsck times!);
  - 10gb/s server connections (here 10GBASE-T looks interesting) will be needed, even if 1gb/s bonding could be used here and there.
- **I am also looking speculatively at:**
  - Coda for network file system;
  - cluster file systems like Lustre or GPFS.

diamond

Business goals 1

- Without business networking and computing the light source continues to work, but people are unhappy.
- Context:
  - cross divisional support, but uniform requirements;
  - 9-to-5 operation;
  - over both a building and office and laboratories around a ring 250m in diameter.

diamond

Business goals 2

- **Major requirements**
  - minimize 9-to-5 downtime;
  - safe record-keeping and backups;
  - compatibility with campus standards;
  - great flexibility.
- **Non requirements:**
  - fully in-house support;
  - ever increasing requirements (even if a doubling of the user population is happening);
  - high data rates or computing power.

diamond

Business solutions 1

- **Networking:**
  - **OS1 (singlemode 9/125) fibre infrastructure, FC connectors;**
  - **CAT6 copper local wiring;**
  - **10GBASE-LR and 1000BASE-LX/-T Ethernet;**
  - **2 level tree topology for fibre;**
  - **single global LAN without routing but multiple VLANs;**
  - **advanced Nortel routing switches, stacked for capacity, with extensive configuration;**
  - **network integrated with the campus network and managed by the campus network support organisation as a black-box.**
  - **Who knows how many FTEs? Perhaps 0.5, perhaps 2.**

diamond

Business solutions 2

- **Computing:**
  - **single sign on using Active Directory and Vintela;**
  - **x86 MS-Windows based servers with 1gb/s Ethernet and Fiber Channel;**
  - **MS Windows, x86 GNU/Linux and Sun based workstations with 1gb/s Ethernet;**
  - **MS-Office and MS-Exchange applications;**
  - **extensive web based applications, both internal and external;**
  - **mostly Dell systems.**
  - **Around 10 FTEs, mostly help desk support.**

diamond