# Building A Computer Centre

# HEPiX Large Cluster SIG

October 21st 2002

Tony.Cass@CERN.ch

# Why do you need a Computer Centre?

- Because you have some combination of CPU servers, Disk Storage and Tape Storage.
- All of this equipment
  - consumes electrical power
    - » which may need to be protected against failure
  - turns electrical power to heat at 100% efficiency,
  - needs to be put somewhere, and
  - needs to be protected against fire.
- Your Computer Centre or Computer Room must
  - have an adequate power supply
  - have adequate cooling,
  - be large enough, and
  - have smoke detection and, possibly, fire suppression.

# Electrical Power

# How much power do you need?

- **Tape drives and robotics are low power**. O(100W) per device and you don't have many of them.
- **Disks are also low power**. O(10W) per device.
  - Main concern is inrush current as drive spins up. Don't start them all at once.
- **CPUs are the problem**. **Lots of them** and **power efficiency is not improving** as raw performance improves.
  - A study of processor specifications suggests constant power consumption of **1W/SpecInt95**.
  - **Power factor** of the power supply is **important**. Off the shelf supplies have power factor of **~0.7**.
    - » Power supplies with power factors **better than 0.8** are available, but these are generally **more expensive**.

# Switched Mode Power Supplies

- PC power supplies only draw current when the voltage exceeds a certain threshold:

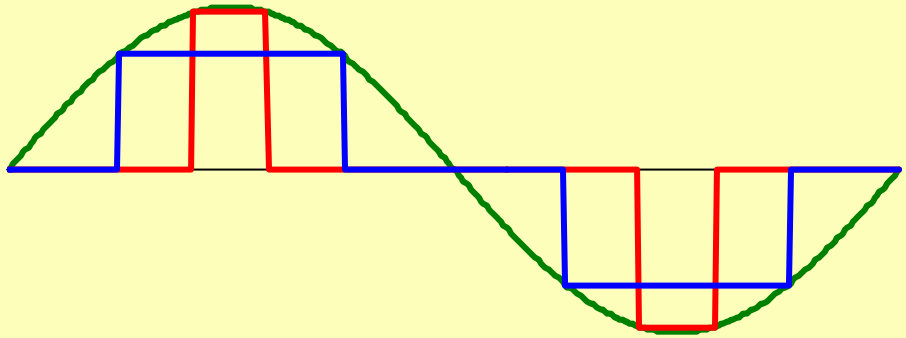

- Result? Lots of **harmonics**, especially **3rd**. These lead to **high currents in the neutral** conductor.
  - Electrical distribution can need **neutral conductor** with **twice the cross section** of the phase conductors.
    - » **All the way** from room distribution, through PDUs to the transformers.
    - » c.f. no neutral needed for ideal 3-phase. US building code used to allow half cross section neutral conductors. No longer!
    - » **Potential fire risk**.

# Power Protection—High Level

- Usual question: "What happens if your primary supply fails?"
- Better question: "**How long does it take you to perform a controlled shutdown?**"
  - My guess? **Over half an hour**. Probably **2 hours**.
- **High power, long autonomy** battery backed (static) **UPS systems are expensive**. **Need diesel** backed (rotary) systems **for decent autonomy**. These are also expensive!
- If you can, **install a 5 minute static UPS to smooth power and cover microcuts**.
  - **expect to lose physics computing due to power cuts**.
  - **consider a 2 hour static UPS for critical systems** (networks, databases, …) if you have no diesels.

# Power Protection—Local

- Many "multi sockets" come with **differential** (or earth leakage) **circuit breakers**.
  - These trip on small difference (usually 30mA) between the phase and neutral currents.
- This is fine for general home use, but **bad for connecting many PCs**.
  - Power supply design gives natural difference of 3-5mA. If you have lots of PCs behind one differential breaker this may trip if there are small disturbances.
- This is **our theory**. We will try omitting these breakers for new installations.
  - But need to provide dedicated sockets for, e.g. consoles.
  - **Ask me for an update next year**…

# HVAC

# HVAC Options

- **Water cooling**
  - **Efficient**, **but no pipes in equipment today**. Does any vendor want to be the first to add these?
    - » IBM Ice Cube would have internal water cooling, but even here, closed system is mentioned as option.
    - » HP talk of using "ink jet" heads to spray water on cpu. This cools the chip faster, but the heat still ends up in the machine room.
  - Can put water cooled pipes in racks (common practice in online areas for experiments) but rack layout has to be fixed.
- **Air cooling**
  - Air conditioning units are connected to the chilled water supply and circulate chilled air in the machine room.
  - **Flexible**, but…

# Air Cooling

- … Heat capacity of air is limited. Extracting **lots of heat** means **lots of air changes** per hour.
- Acceptable level depends on room usage and occupancy.
  - **above 60-70 AC/h** the environment becomes **uncomfortable** for general work.
    - » **Good! Make SysAdmins work from their offices!**
  - **120AC/h acceptable for technical areas**
- If possible, **avoid closed systems**. Injection of external fresh air allows "free cooling" in winter.
- Cold air injection usually under false floor
  - need to ensure air flow through equipment racks.
  - Roof level injection possible, but requires high ceiling.

# Air Cooling Limits at CERN

| | *Clearance* | |
|---|---|---|
| *Air Flow* | 2.5-3m | >6m |
| **Comfortable** | 600W/m$^2$ | 1200W/m$^2$ |
| **High** | 1200W/m$^2$ | 1800W/m$^2$ |

For guidance only: your mileage may vary…
Don't forget to account for overheads (e.g. solar heating)! These depend on location, but we have 3-400W/m$^2$ for the main machine room in summer.

# CERN Planning

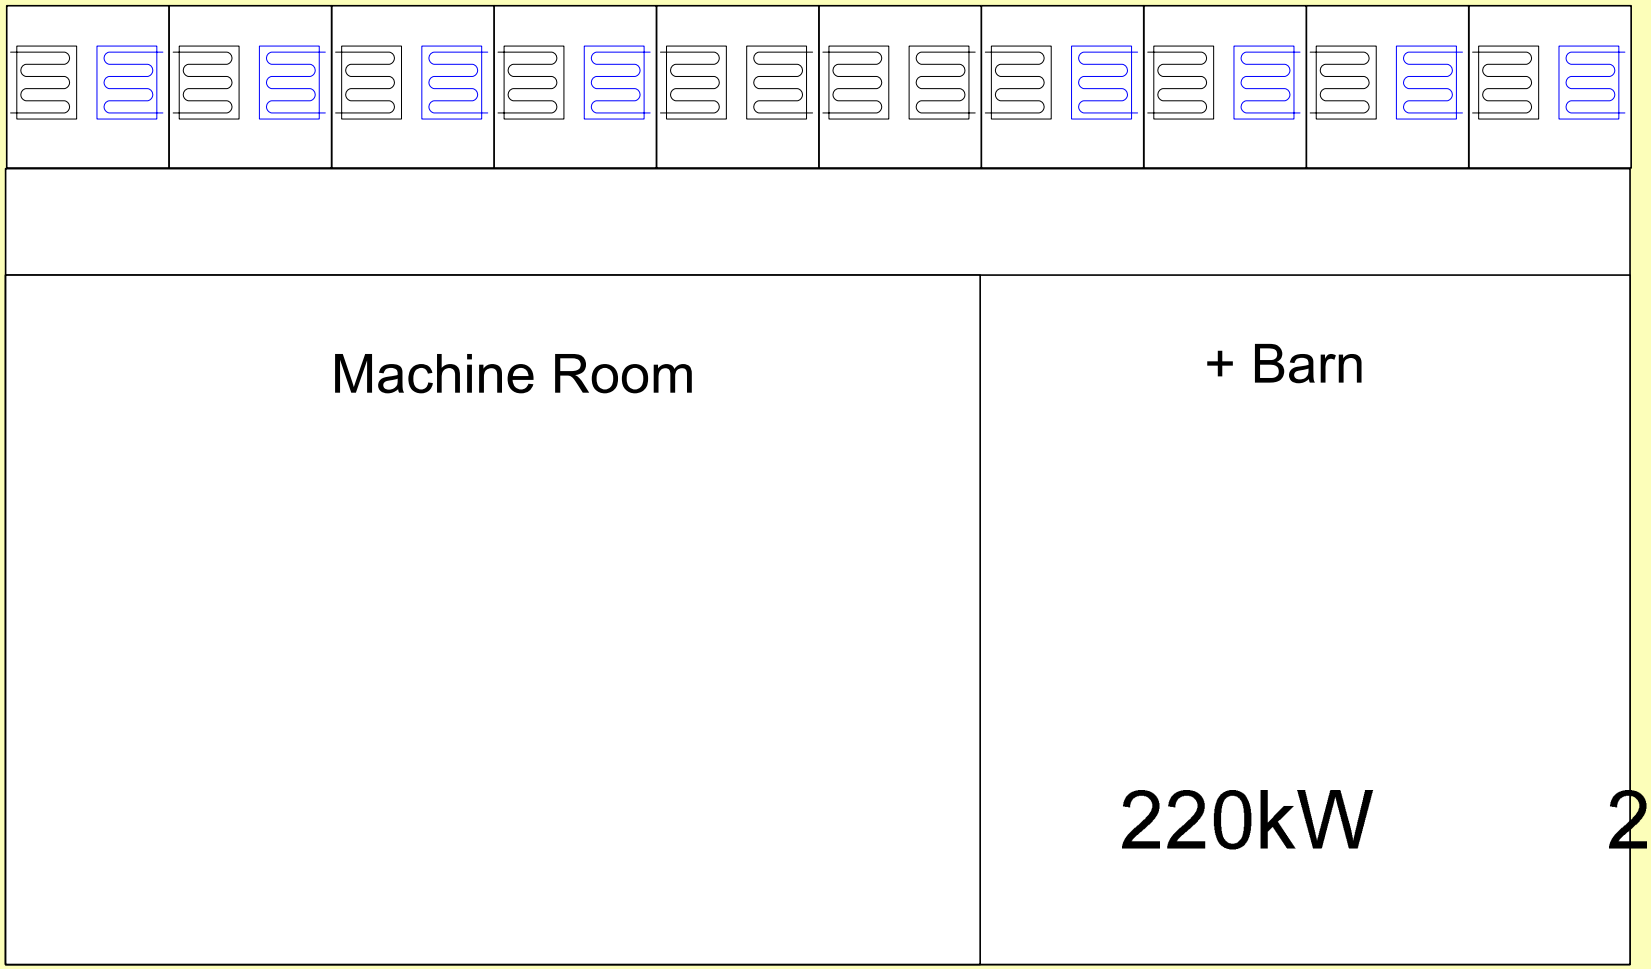

# Space

# Space

- **Box size may not be the determining factor**!
  - 40 1U boxes in 19" rack: 4kW in 1m$^2$ (including clearance).
- So, good if you can **mix CPU and disk servers**.
  - Better still if you have tape robots!
- Otherwise, space is relatively easy. Just **remember** to add enough **clearance**.

# Fire Precautions

# Smoke Detection

- **Laser based** smoke detection systems are **extremely sensitive**.
  - Use these to **raise early alert**, but
  - **need less sensitive detectors as "second knock"** to trigger any fire suppression or prevention.
- **Localisation is a problem**.
  - The combination of sensitive detectors and rapid air flow means all detectors will trigger in time—usually sooner, not later.
  - **Locality of smoke detection competes against efficient cooling.**
    - » You can divide a big room into smaller ones, but it will be harder to cool.
    - » The same is true for closed racks. And, anyway, smoke will escape into the room eventually.

# Fire Suppression or Fire Prevention?

- **Fire Suppression if you want to keep services running** at all costs (e.g. banks, phone companies, ISPs).
  - Usual system is to **flood machine room with an inert gas** (Nitrogen/Argon mixture now that Halons are banned).
    - » Gas is expensive and bottles are bulky.
    - » Machine room needs to be hermetic!
  - Some modern water based systems can be used on running equipment and are good for scrubbing smoke particles from the air.
- **Fire Prevention** is the alternative **if you want to avoid fire at all costs**.
  - **Cut electrical power as soon as smoke is detected**.
    - » Risk of progression to fire cut by 20x-100x for smouldering PCBs
- **Fire Prevention chosen at CERN** for our new machine room in the former tape vault. Ask our opinions next year…
  - We're still not sure what to do for our main machine room because of the localisation problem.