# Benchmarking, Power Consumption and Thermal/Power Measurements in the CERN Computer Center

# Some Background – The Computer Center

- The CERN Computer Center was built in the early 1970s
  - Designed for the mainframes of those days
    - Esp. in terms of the cooling / air conditioning system
  - Recently upgraded to better accommodate rack-mounted systems
- Certain limitations remain
  - Most importantly, only 2.5MW of heat can be cooled
  - ... which means only 2.5MW of power can be provided

# Some Background – The Computers

- Computing at CERN is done using standard rack-mounted servers
  - We decided to use “1U” rackmounted machines as compute nodes
  - Pretty much all systems have an Intel processor
    - ~700 systems with two “Woodcrest” processors (Xeon 51xx)
    - A new delivery of chassis with two motherboards (== two independent systems) and two “Clovertown” processors (Xeon 53xx)
    - The tendering process does not specify Intel ...

# Basics of Benchmarking

- The goals of benchmarks are to provide
  - Comparable results – the benchmark results of different computers are directly comparable
  - A “simple” number or set of numbers as the direct measure of the compute power of a computer for a specific use-case
- CERN IT uses the SPECInt2000 benchmark for comparing offers during the tendering process for new computers
  - The configuration is customized, so the results reflect the general use-case at CERN
- The experiments use, for example, special versions of their online software for taking decision on purchases for the online compute farms

# Basics of Power Consumption

- Over the years the power consumption of the computers has become a serious concern
- Several sub-systems in a computer have a significant power consumption
  - CPU: Intel “Harpertown” Xeon 54xx: 80W - 150W (TDP – Thermal Design Power)
  - Memory
    - FB-DIMMs: ~10W per 1GB module; ~15W per 2GB module
    - DDR2: ~5W per GB
  - Hard disks: depending on model ~10-20W under load
  - The rest – chipset, fans, etc.: ~80-100W
- Other variables like “power factor” and the efficiency of the power supply are also important factors

# Benchmarking and Power Consumption

- Benchmarks and power consumption are correlated
  - The “raw” benchmark results for compute power and the “raw” value of power consumption by themselves have only a limited significance.
  - The metric that provides a reasonable basis for decision making is compute power over electrical power (SPECInt/Watt)
  - A computer system that provides the most compute power per unit electrical power is the most efficient...
- Factors that drive our decision making process are
  - SPECInt/Watt for the overall most efficient system
  - and also power supply quality (power factor and efficiency)
  - Both factors translate also into the financially most effective solution when considering the lifetime costs (“TCO”)

# Power Measurements

- Extensive power measurements during the tendering process and also in openlab in order to better understand the behaviour of different types of systems
  - For example measurements by a summer student using “standard” batch nodes

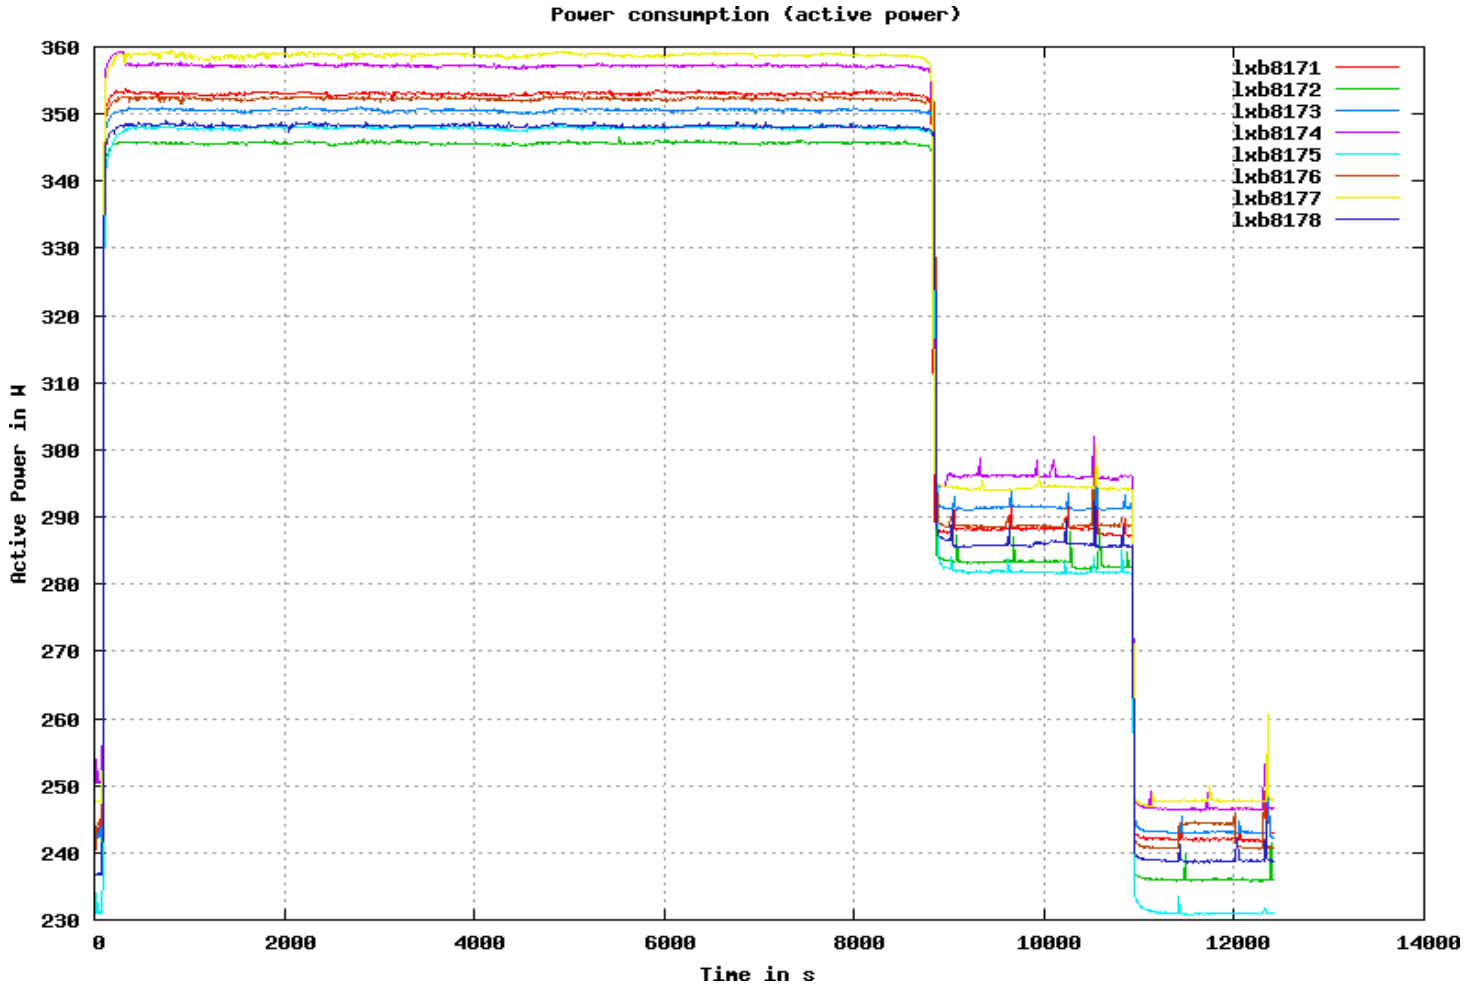

# Power measurements

- Under the load conditions the power consumption should have been constant ... but for some machines it wasn't ...
  - It turned out that in this case the power factor of the power supply “jumped” for an unknown reason...

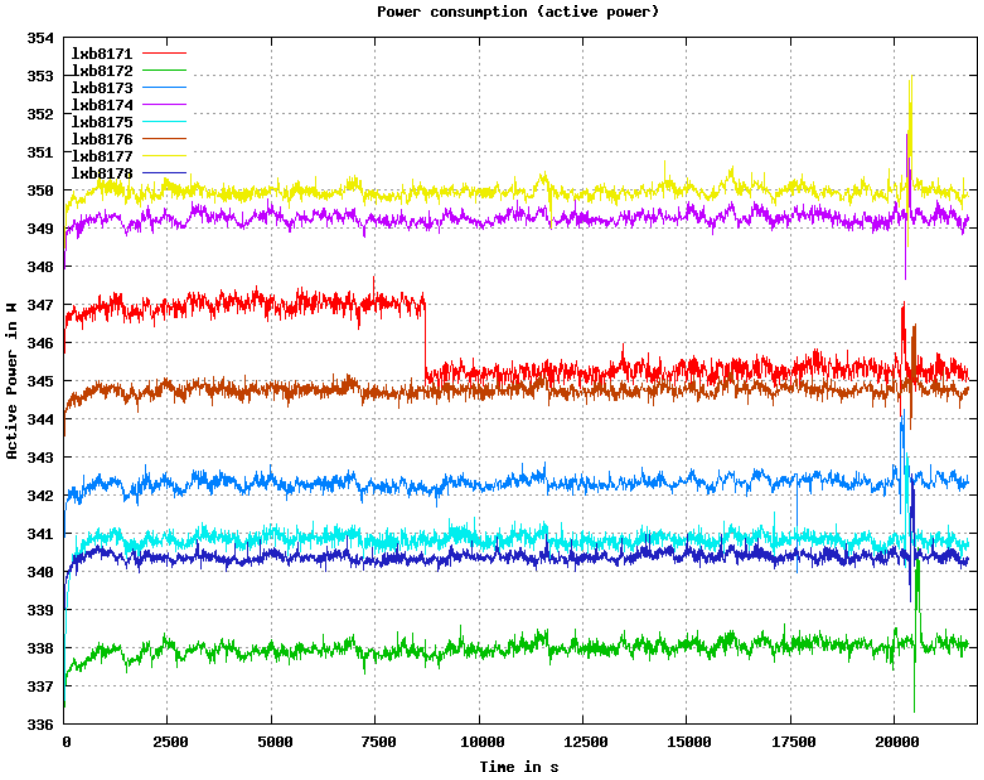


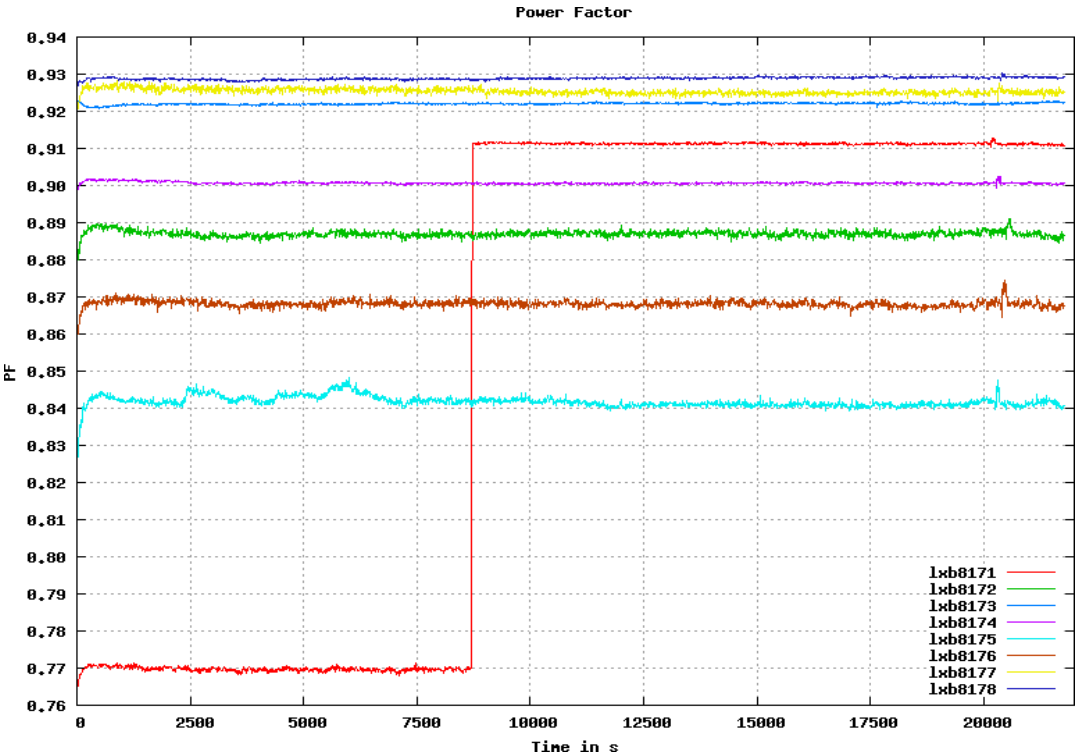

# ... the end result of the efforts ...

- After being able to specify power related measurements and requirements in our purchasing procedures we were able to purchase machines that were very power efficient in every aspect possible!!!
- This does, of course, not mean that they cost less!!!
- It means that over their lifetime they will provide much better efficiency per unit electrical power and per CHF (electricity costs money...)

# A quick outlook

- Power consumption/power efficiency is the major driving force in the IT industry and will continue to grow in importance
- The SPEC consortium is preparing a benchmark targeted at power consumption/efficiency: SPECpower
  - And CERN is a beta-tester ...