Karlsruhe Institute of Technology

# Benchmarking and Power Measurements at GridKa

Manfred Alef

# Power Measurements at GridKa

- Power measuring during benchmark runs
- Power analyzer (Infratek 101A), connected to a dedicated controller system, >1 readings per s
- Comparing W and $A_{max}$(*) during HS06 benchmark runs
- Used for TCO estimations in procurements ...
- ... and for rack design, ...

* How many servers can I connect to a PDU protected by a 32 A breaker?

# Power Measurements at GridKa

- The power analyzer (Infratek 101A) has C14 inlet, allowing 10 A max
- Increasing number of servers coming with C20 inlets, power consumption exceeding 10 A
- Dedicated controller system: Windows 95, not connected to LAN, data transfers using USB device

- Looking for replacement

# Alternatives

- The good news (?):
  - ☺ Modern servers are equipped with a baseboard management controller (BMC), many of them providing power reading ...

```
ipmitool sdr | grep -i watts
```

... and inlet air temperature (also affecting power consumption!)

```
ipmitool sdr | grep -i degrees
```

  - ☺ May also provide access via web interface

# Alternatives

- The bad news:
  - ☹ No common keywords for power reading ...
    - ☹ "PS1 Power In" + "PS2 Power In"
    - ☹ "SYS_POWER" (PSU output!)
    - ☹ "PSU1PIN" + "PSU2PIN" > "SYS_POWER"
    - ☹ ...
  - ☹ ... or ambient temperature
  - ☹ No min, max, and accumulation (kWh)
    (Min and max possibly available via web interface)

# Alternatives

- The bad news:
  - ☹ Interpretation strongly system dependent:
    - ☹ PSU-input, and/or PSU-output
    - ☹ Redundant PSU: total power vs. power per PSU module
    - ☹ How many nodes (e.g. 2U4N) are connected to the PSU?
  - ☹ Customized, host model dependent service script required when using power reading in benchmark tool

# Alternatives

- The bad news:
  - ☹ Frequently running 'ipmitool sdr' command on system under test, alongside the benchmark, will possibly affect accuracy of benchmark results
    - ✷ However, can query the BMC remote

# Alternatives

- Are measuring PDUs an option?
  - ☹ Individual SNMP GET command line syntax, or
  - ☹ Individual web access

# Alternatives

- Open questions:
  - ☹ Accuracy of such "cheap" power measuring devices?
  - ☹ Number of readings per time unit?

# Power Measurement in SPEC Benchmarks

- Some SPEC® benchmarks, e.g. SPEC CPU® 2007, incorporating power measurement
- Setup:
  - Dedicated controller system, running SPEC PTDaemon® tool, connected to an accepted power analyzer* and to an accepted ambient temperature sensor*

    * https://www.spec.org/power/docs/SPECpower-Device_List.html
  - System under test running the benchmark, talking to PTDaemon on controller system

# Conclusions

- Coupling of benchmarking and power reading, for TCO estimations in procurements, rack layout, ...
- Using a high-quality power analyzer so far, C14 (10 A) inlet, connected to dedicated controller system
- Looking for alternatives because of increasing number of WN models with C20 inlets
- Many servers coming with BMC providing power reading, but no standards, lower accuracy
- Using dedicated controller system for measurements while running benchmark on system under test

# Questions, Comments