A study on energy consumption of cluster nodes toward Green Clouds

Vesselin Hadjitodorov
Master education System and Network Engineering

February 02, 2011

Energy consumed by ICT

- Amount of energy consumed
  - Doubles every 6 years
  - Consumed energy costs more than equipment
- Additional infrastructure costs:
  - Cooling, UPS, Power lines
- Solutions:
  - Outsourcing data centers
  - Improving energy efficiency

Research questions

What is the influence of system components in cluster nodes on the total energy consumption?

- better understanding of the energy consumption
- finding the best way to optimize the system

How does one measure the energy consumption of cluster nodes, during the execution of various testing scenarios?

What improvements can be made without sacrificing performance?

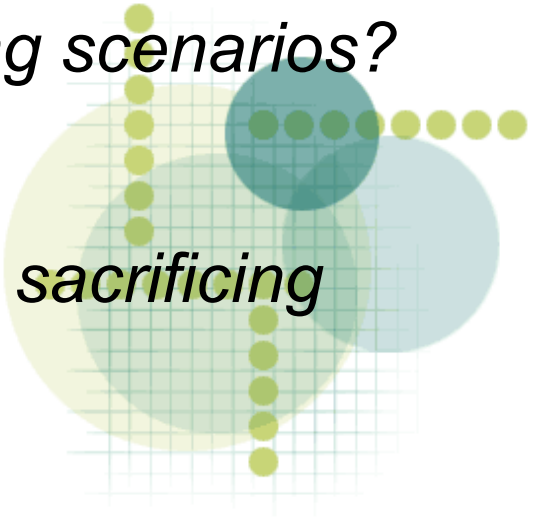

GreenClouds

- Project by VU and UvA, funded by NWO
- Uses DAS-4 cluster infrastructure
- The goal is to find methods to reduce energy and pollution footprint of HPC
- Hardware and software based solutions for increasing efficiency:
    - Accelerators
    - Virtualization
    - Task scheduling
    - Hybrid networks

DAS-4 cluster nodes

- Test bed for scientists
- Hardware of each tested twin node:
    - PSU: 2 x 1.4 kW Supermicro (shared)
    - CPU: 2 x Intel Xeon E5620 @ 2.4 GHz
    - Memory: 6 x 4GB DDR3 @ 1333 MHz
    - Storage: 1 x 1TB SATA II Western Digital RE3
- Currently running CentOS

DAS-4 twin node

Performed tests and measurements

- Optimization requires power consumption measurements

- Measured components:
  - Hard drive
  - Memory
  - Power supplies
  - CPUs

Measurement setup

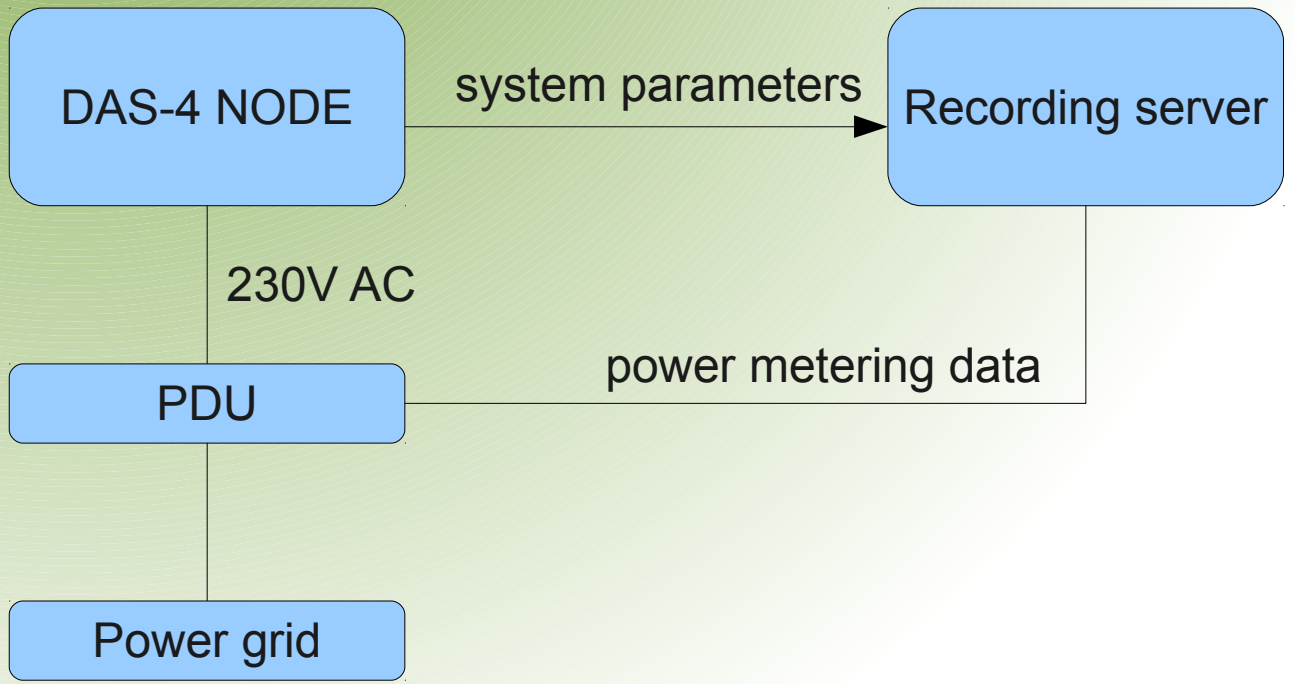

- Data can be stored in CSV, XML or database
- Currently LAMP server + adaptive DB used
- Data is sent using GET requests to server

Measurement instruments

- Data obtained using Schleifenbauer PDU:
  - Voltage
  - Current
  - Power factor

- Perl API used
- Requires separate machine
- Measurement density is too low

image source: http://www.schleifenbauer.eu

Data collection from nodes

- Number of cores
- Frequency of cores
- Kernel and user load

- Sysstat and /proc/cpuinfo used for obtaining data

Hard drive

- Minor difference between idle and active
  - Values are in the error range of the current setup
- According to specifications:
  - Idle consumption – 7.8W
  - Read/write consumption – 8.4W

Memory measurements

- Hard to measure separately from CPUs
- Dependent on number of memory modules

According to Kingston datasheets:

6 x 4 GB = 23.8W

Measurement confirm result

3 x 8 GB = 13.1W – better ?

PSU measurements

- Turned off nodes consume around 35W each
- Load and power factor are not distributed equally
- Energy efficiency is dependent on the load of PSUs

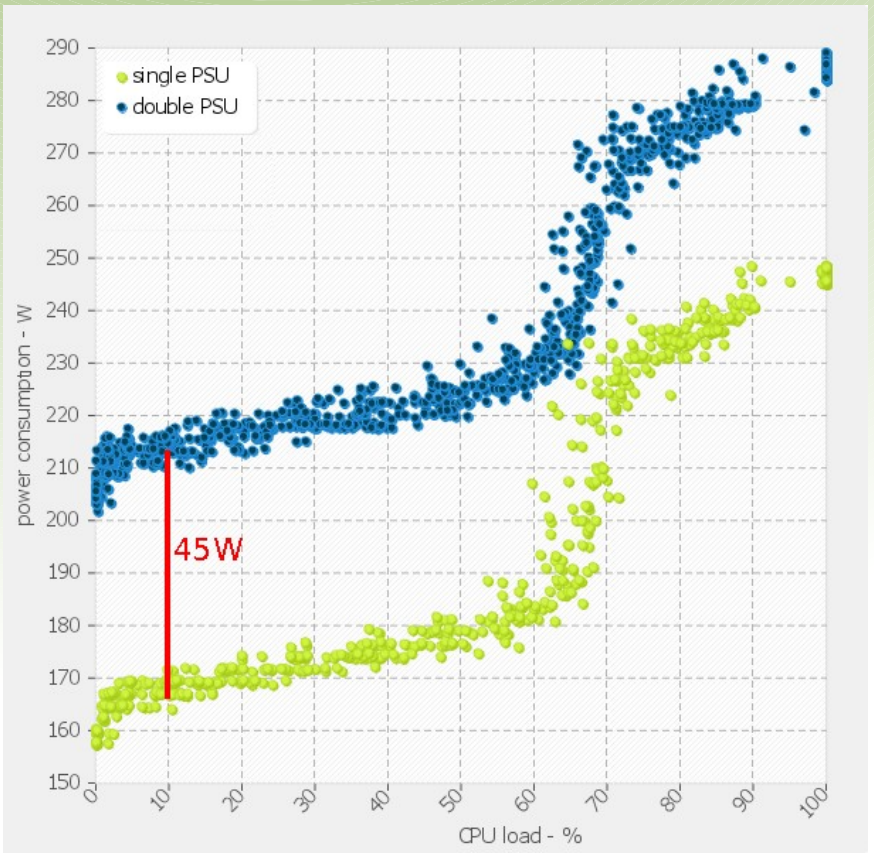

one node

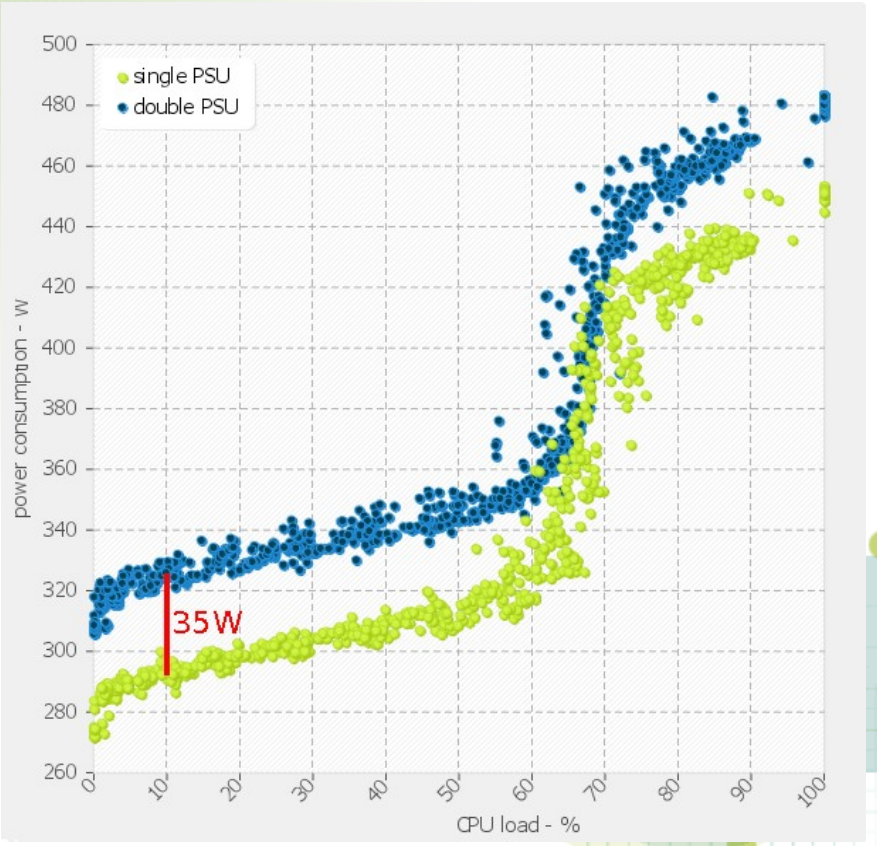

two nodes

CPU measurements

Load generating:
- cpuburn
- gradually from 0 to 100%

- CPU load is not equivalent to real performance

- 100% loaded CPUs requires 90W additional power
- Scaling is efficient under low load
- Power consumption is dependent on frequency
- Frequency is scaled at around 70% load of core

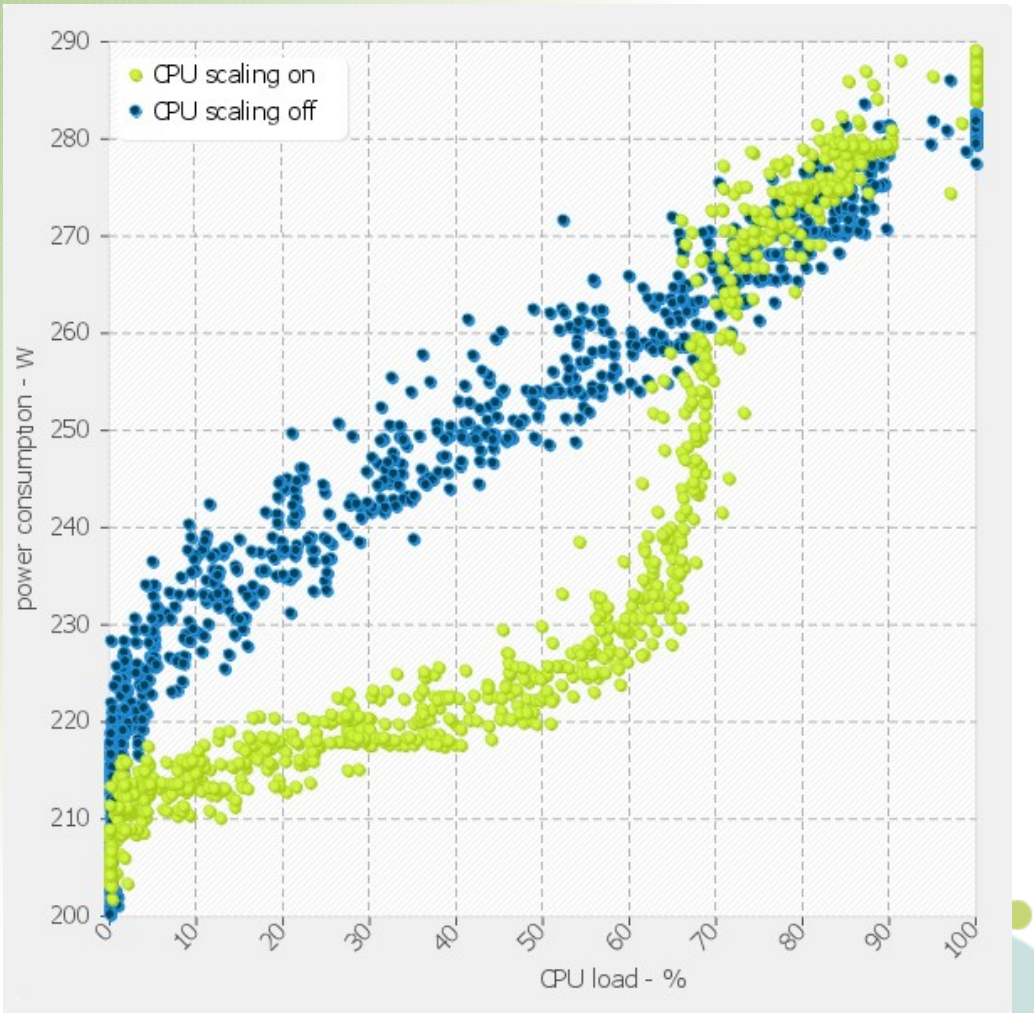

scaling vs. non scaling CPUs

How measurements can be improved

- Higher measurement density and precision
    - New API or firmware
    - Acquiring power analysis equipment
- Measuring on low voltage power rails
    - PSU efficiency can be monitored
    - Individual components can be distinguished
- All in one system monitoring tool
    - Detailed system information with higher refresh rate

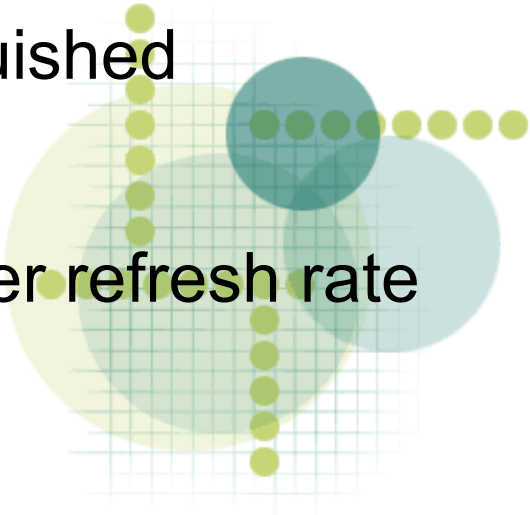

Conclusion

What improvements can be made on the nodes without sacrificing performance?

- Replacing PSUs or using a single one when possible
- Replacing memory modules

Future work:

- Using SSD instead of HDD
  - SSDs have 80% lower consumption for the same capacity
- Using accelerators such as GPUs and FPGAs
  - They have higher performance per watt for certain tasks

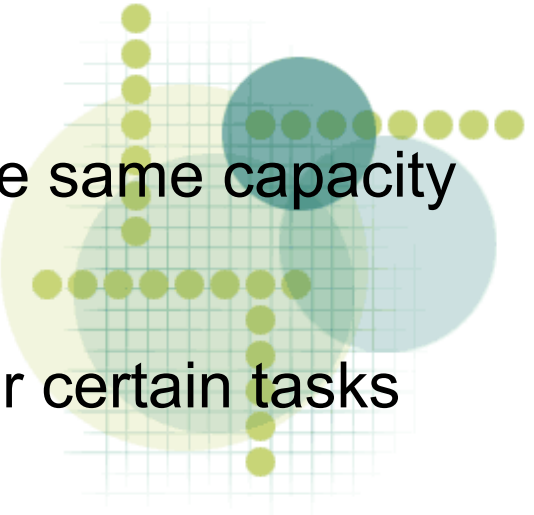

Questions ?