Macro and Micro-level Energy Savings Through COTS Technologies

Markus Leberecht
Technical Alliance Manager
16 September 2010

Total vs. Local Energy Optimization

Power Per Unit

< 50 Watts

~50-200 Watts

Kilowatts

Megawatts

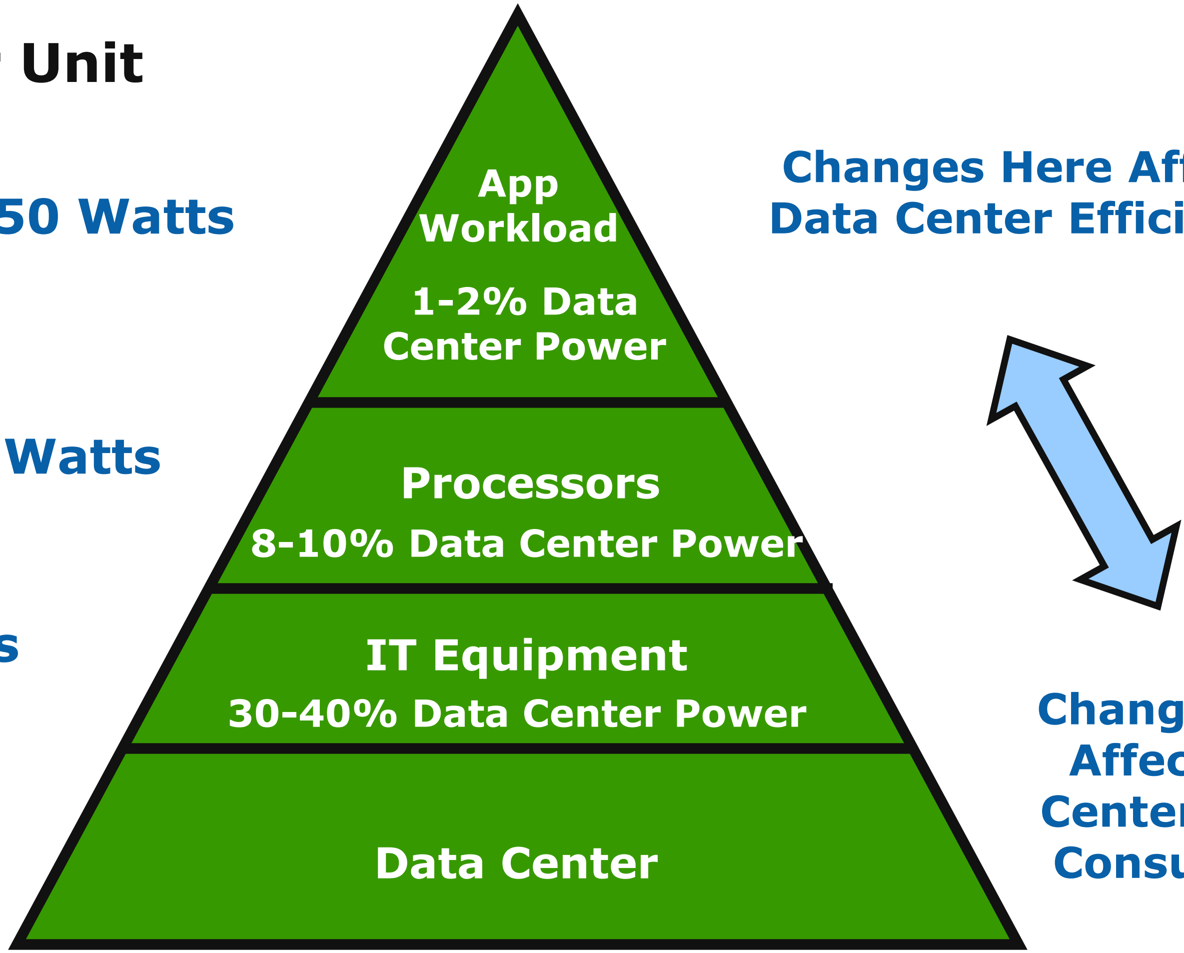

Changes Here Affect Data Center Efficiency

Changes Here Affect Data Center Power Consumption

Total vs. Local Energy Optimization

Power Per Unit

< 50 Watts

~50-200 Watts

Kilowatts

Megawatts

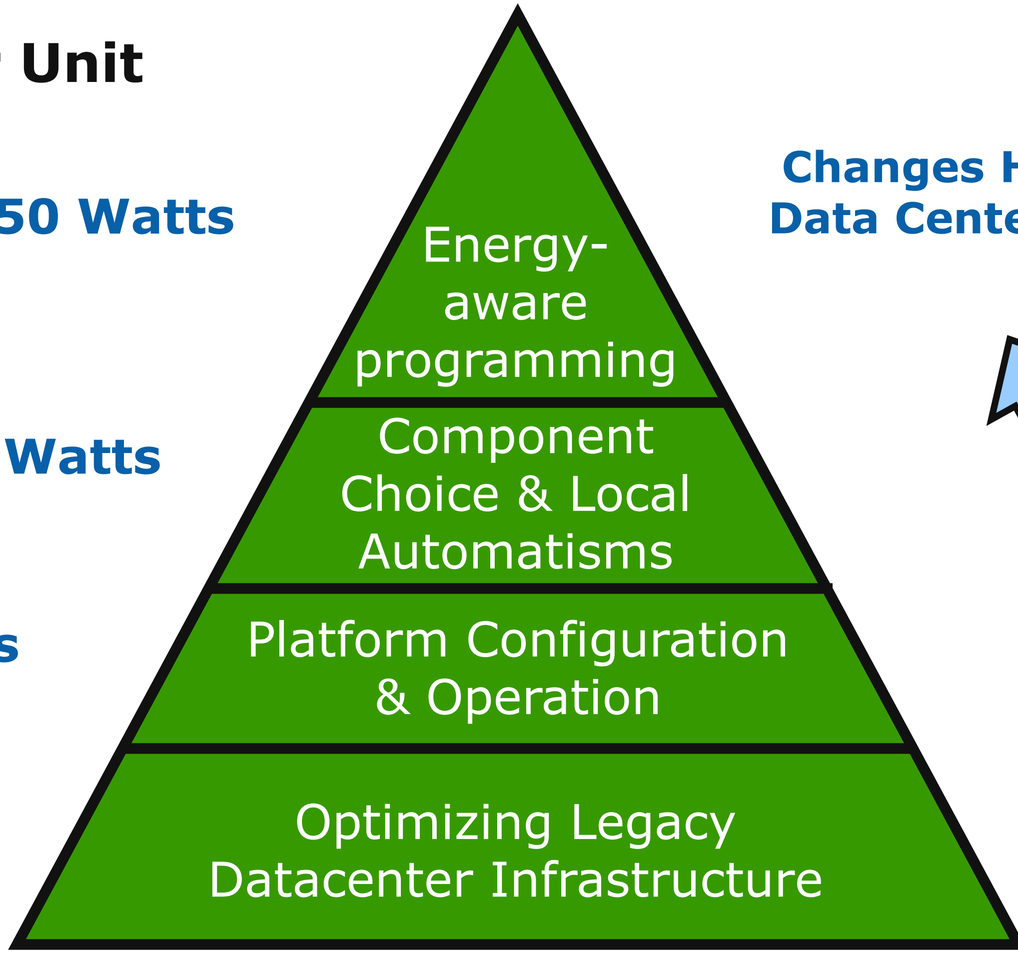

Changes Here Affect Data Center Efficiency

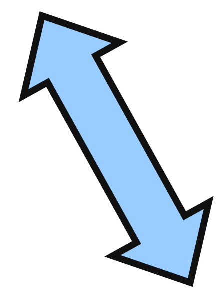

Changes Here Affect Data Center Power Consumption

Total vs. Local Energy Optimization

Power Per Unit

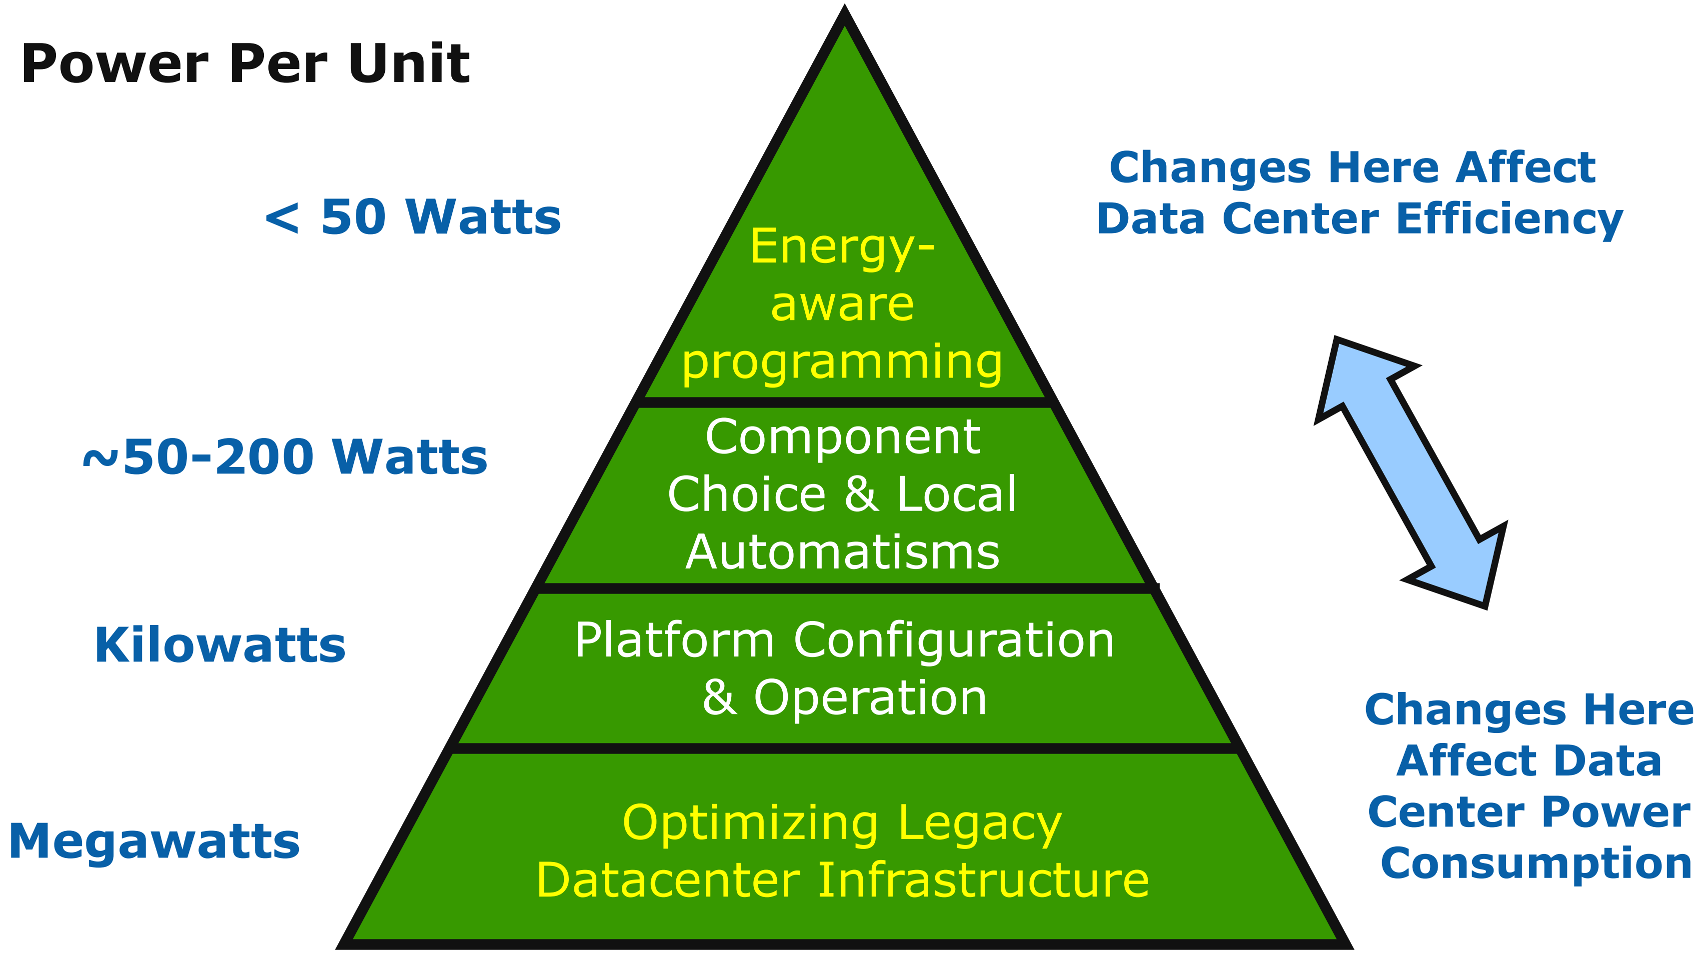

End-to-End Datacenter Efficiency: www.datacenter2020.com

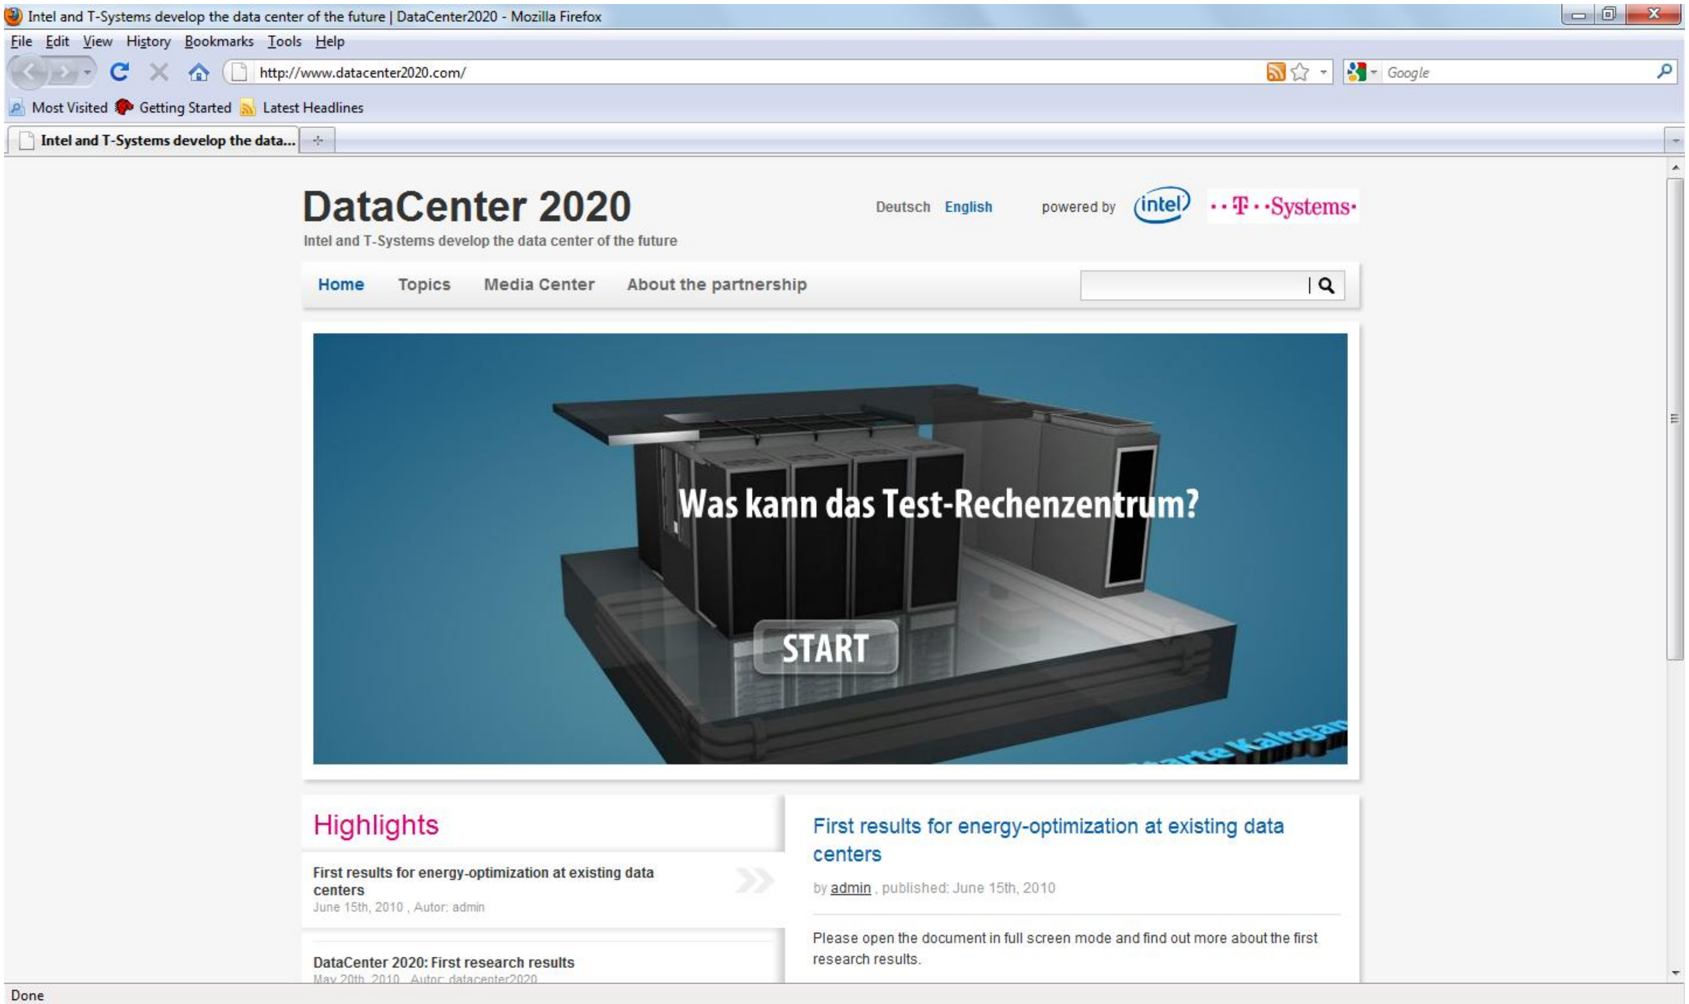

Optimizing Legacy Data Center Infrastructure

Average Data Center Power Allocation

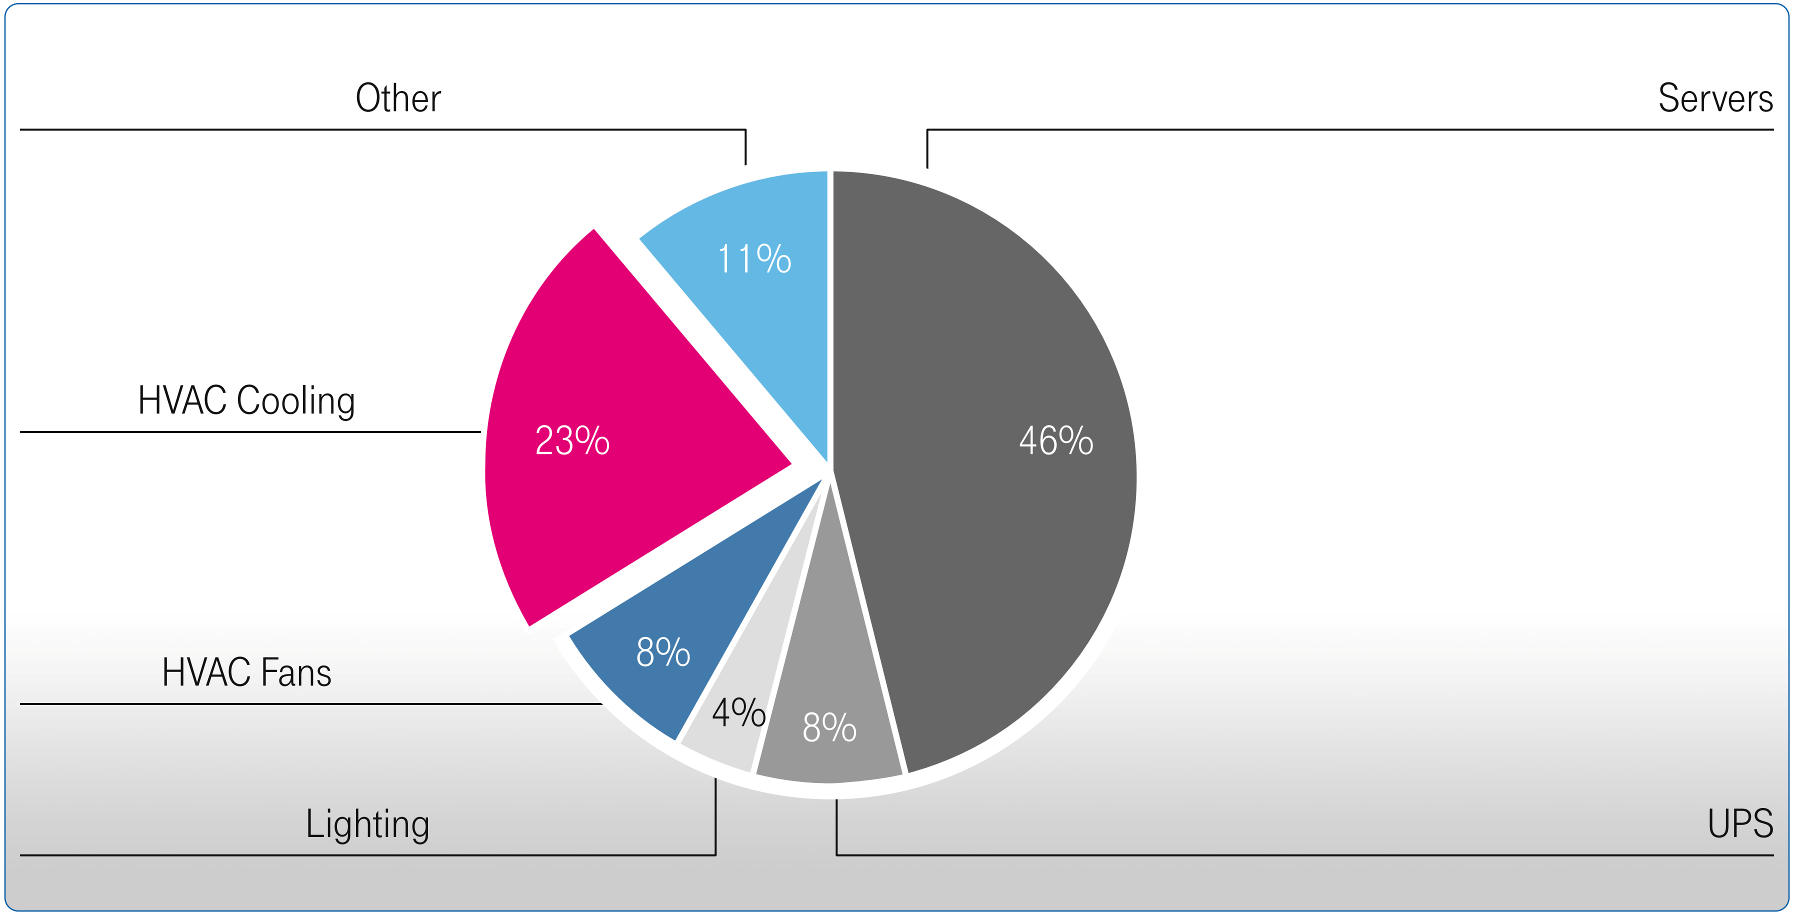

Optimizing Legacy Data Center Infrastructure

Test Lab Fast Facts.

- Cold water supply
- CRAC EC
- CRAC FT
- CRAC Humidity
- Liquid cooled Rack
- Cold aisle containment
- Hot aisle containment
- Building mgmt. system
- Raised floor
- RFID
- Permanent inertization

- 180 real Servers
- DC-Supply48V
- AC-UPS
- Adjustable ceiling
- Coated walls
- Smoke generator

- 8 Racks

Optimizing Legacy Data Center Infrastructure

First Phase Equipment Utilized.

Cold water supply

CRAC EC

CRAC FT

Cold aisle containment

Building mgmt. system

Raised floor

Permanent inertization

180 real Servers

AC-UPS

Adjustable ceiling

Coated walls

8 Racks

Optimizing Legacy Data Center Infrastructure

Simulating a Typical Enterprise DC with Indirect Free Cooling.

- No Enclosure
- FT-CRAC
- 8° C water supply
- 100% fan speed
- Max. 100% Leakage
- No blind blanks
- 100% CPU-Load
- perforated tiles (38% opening)
- Cold water valve setting: Auto
- Room height 2.7 m

- **Assumed IT power delivery limit of 5 kW/rack.**

Infrastructure Optimization Effects.

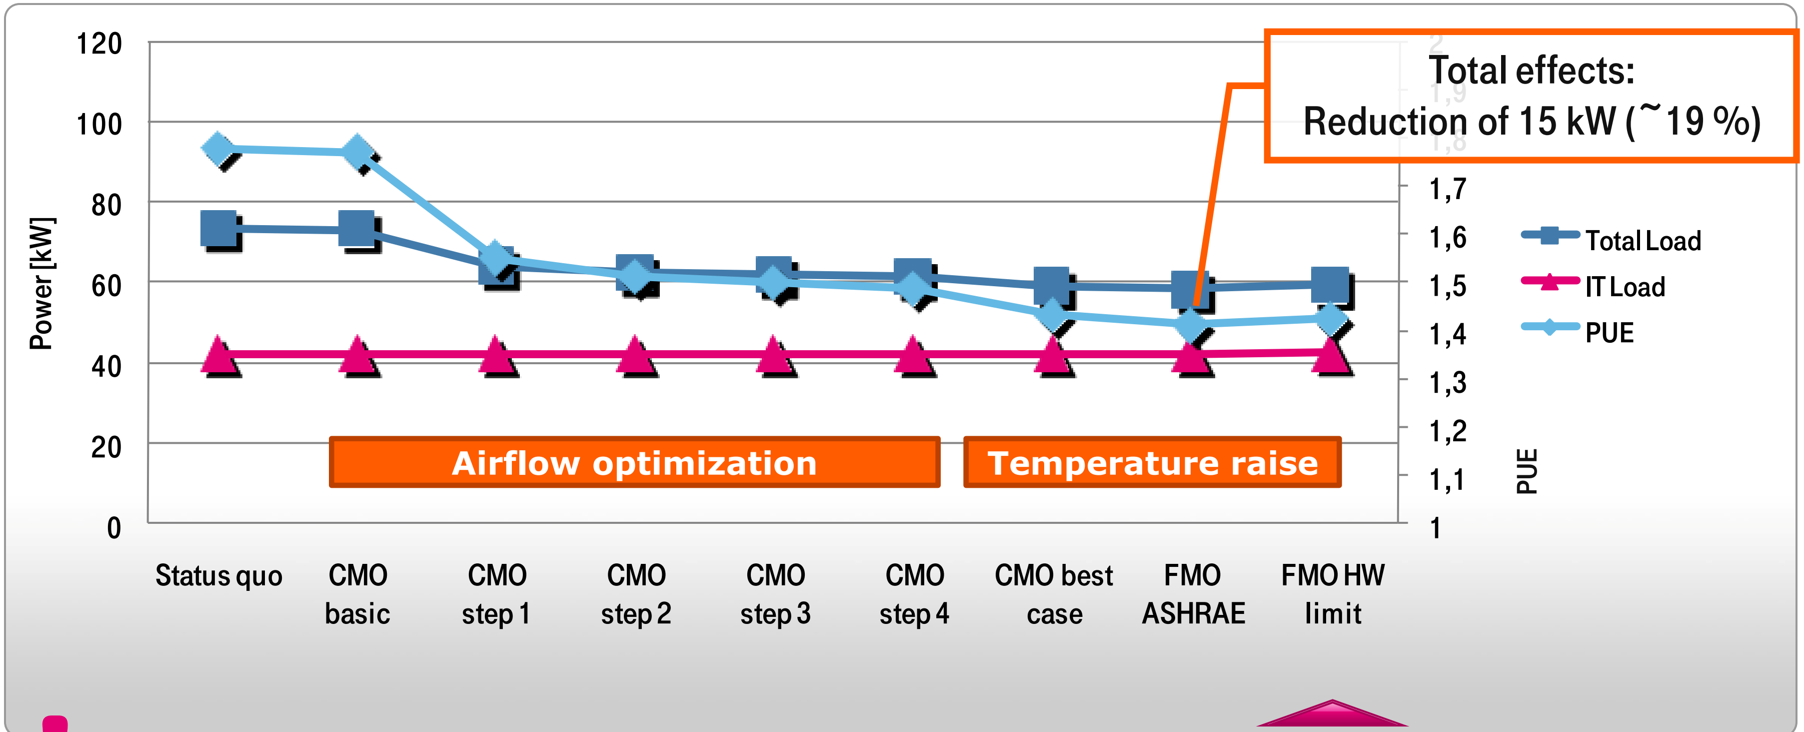

- Start with reducing leakages, separating hot and cold air, reducing fan speed
- Further gain by increasing water supply and room temperature towards ASHRAE recommendations

Optimizing Legacy Data Center Infrastructure

Summary of Infrastructure Changes.

Status quo	CMO basic	CMO step 1	CMO step 2	CMO step 3	CMO step 4	CMO best case	FMO ASHRAE	FMO HW Limit
No Enclosure	No Enclosure	No Enclosure	Doors	Doors	Cold Aisle Enclosure	Cold Aisle Enclosure	Cold Aisle Enclosure	Cold Aisle Enclosure
FT-CRAC	FT-CRAC	FT-CRAC	FT-CRAC	FT-CRAC	FT-CRAC	EC-CRAC	EC-CRAC	EC-CRAC
8° C water supply	8° C water supply	8° C water supply	8° C water supply	8° C water supply	8° C water supply	14° C water supply	24° C water supply	34° C water supply
100% fan speed	100% fan speed	optimize fan speed	optimize fan speed	optimize fan speed	optimize fan speed	optimize fan speed	optimize fan speed	optimize fan speed
100% Leakage	20% Leakage	20% Leakage	20% Leakage	20% Leakage	10% Leakage	10% Leakage	10% Leakage	10% Leakage
No blind blanks	blind blanks	blind blanks	blind blanks	blind blanks	blind blanks	blind blanks	blind blanks	blind blanks
100% CPU-Load	100% CPU-Load	100% CPU-Load	100% CPU-Load	100% CPU-Load	100% CPU-Load	100% CPU-Load	100% CPU-Load	100% CPU-Load
perforated tiles	perforated tiles	perforated tiles	perforated tiles	grating tiles	grating tiles	grating tiles	grating tiles	grating tiles
Cold water valve Auto	Cold water valve Auto	Cold water valve Auto	Cold water valve Auto	Cold water valve Auto	Cold water valve Auto	Cold water valve 75%	Cold water valve 75%	Cold water valve 75%
Room height 2,7 m	Room height 2,7 m	Room height 2,7 m	Room height 2,7 m	Room height 2,7 m	Room height 2,7 m	Room height 2,7 m	Room height 2,7 m	Room height 2,7 m

First optimization step — Second optimization step — Third optimization step

Optimizing Legacy Data Center Infrastructure

Raising Data Center Capacity.

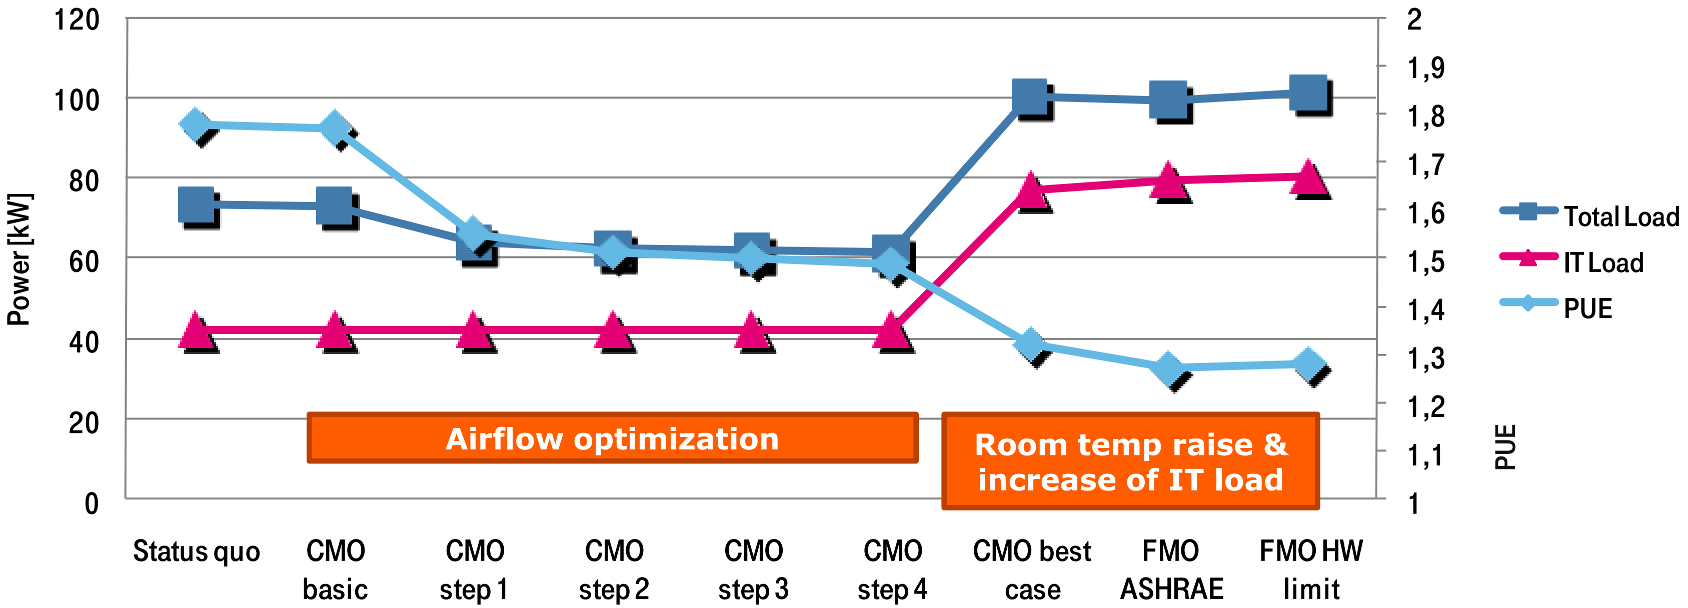

- Doubling the IT load from approx. 5 kW/rack to approx. 10 kW/rack further improves efficiency.

Optimizing Legacy Data Center Infrastructure

"HPC-like" Thermal Densities

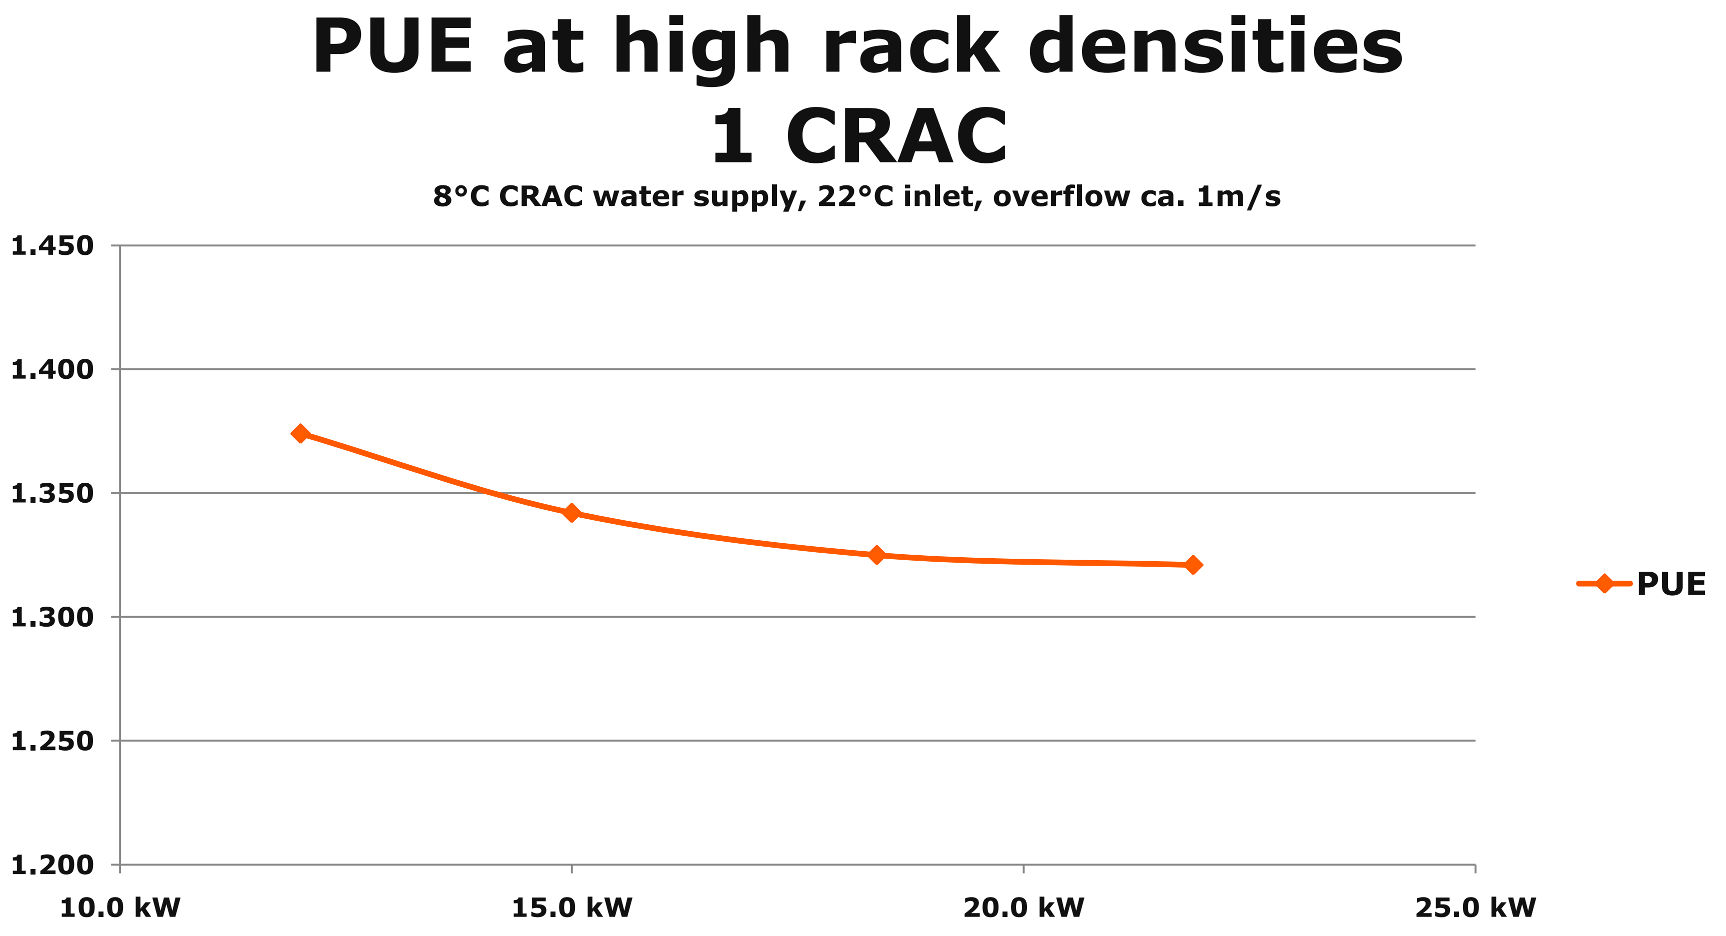

Optimizing Legacy Data Center Infrastructure

Increasing the Cooling Capacity

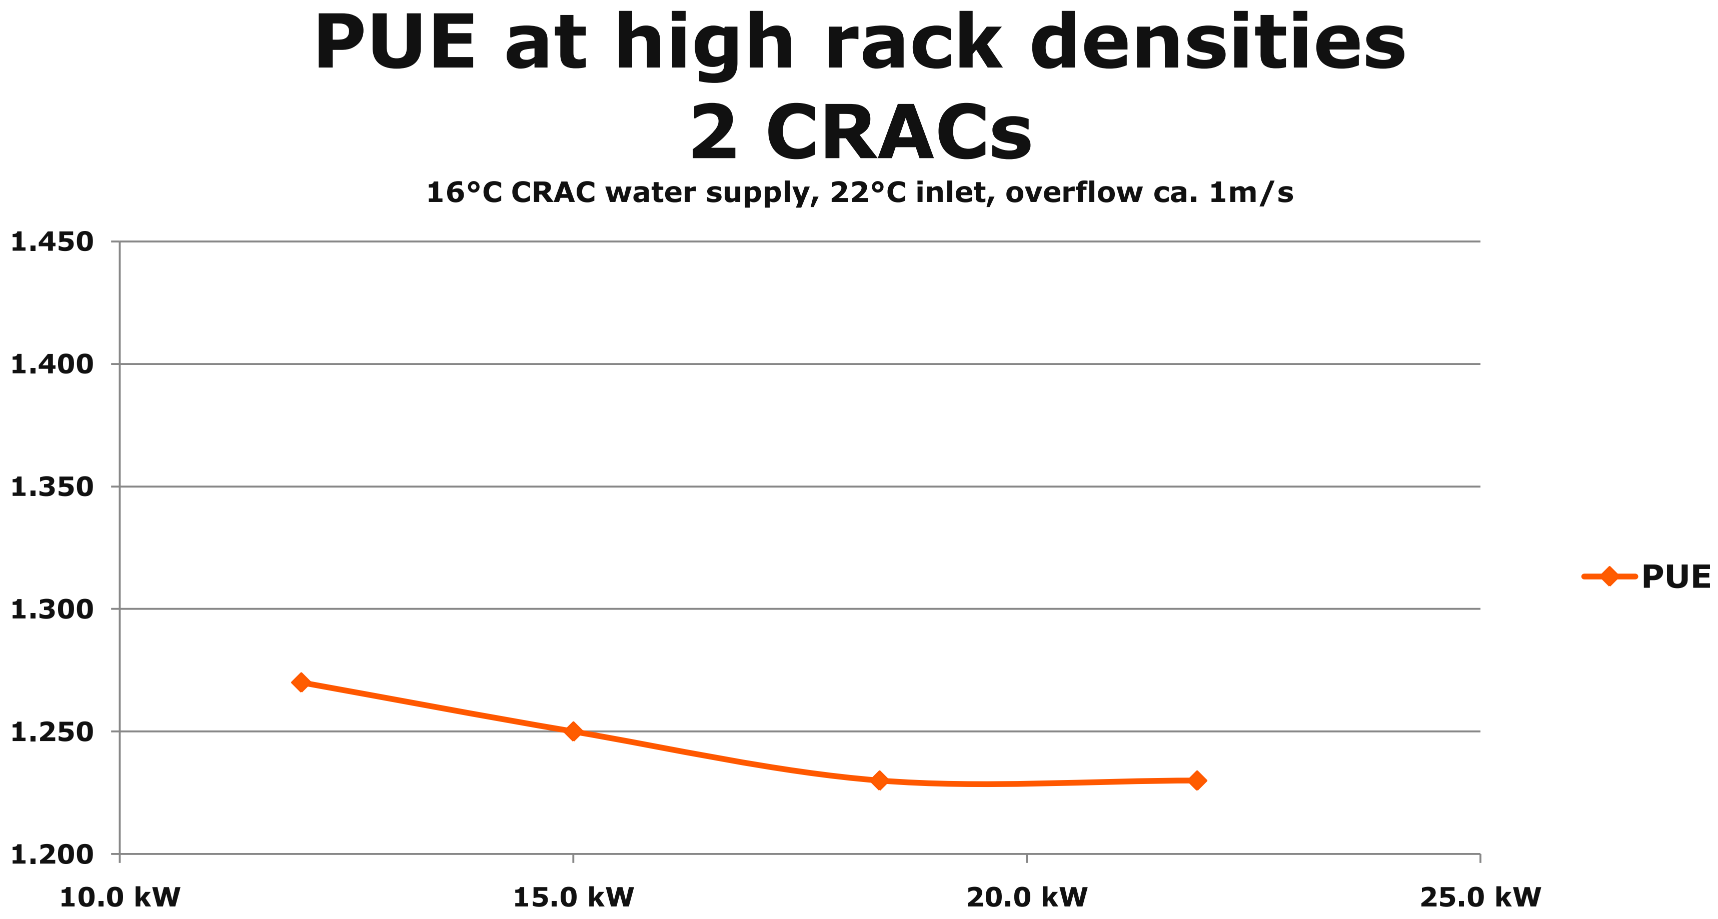

Optimizing Legacy Data Center Infrastructure
Summary & Next Steps

- Findings
  - **High thermal densities possible with conventional data center cooling infrastructures**
    - Clean airflow separation permits lower fan speeds
    - Raised CRAC water supply temp requires less compressor cooling
  - **Dense packaging beneficial for efficiency**
    - Approach your cooling capacity limit!

- Next Steps
  - HPC building practices (cabling & hot spots, CFD modeling)
  - End-to-end power management
    - Node / rack / row / DC power policies
    - Capping & Smart ride-through capabilities
  - **Targeting energy-proportional computing**

Measuring Efficiency

- Efficiency = Useful Work / Resource Consumed
  - Increasingly applied to Energy Efficiency:
    Energy Efficiency = Useful Work / Energy Consumed
- Seems simple, except:
  - No agreement on what constitutes “useful work”
  - Often difficult to measure energy consumption

There is no universal quanta of work

Common Metric Example: CPU Utilization

- CPU utilization often used as a proxy for the work done by a server
- Things that can make utilization go down:
  - More efficient software algorithms
  - Faster/more memory or disk
- Things that can make utilization go up:
  - Low/no-value background applications
  - Intrinsic overhead (e.g. on context switches)
- CPU utilization metric can penalize efficiency!

Energy-aware Programming

The Intel® Energy Checker SDK

- Intel® Energy Checker SDK is a series of routines that can be integrated into any application to write out counters in a standard manner
  - No external libraries or run-time software must be installed with the application; this is standalone
  - Becomes part of the application, not the system

- These counters can be easily read and aggregated to report the productivity of a system
- These counters can be used during benchmarking
- They are also lightweight enough to be used in production systems with negligible impact on performance

Energy-aware Programming

Intel® Energy Checker (EC) SDK

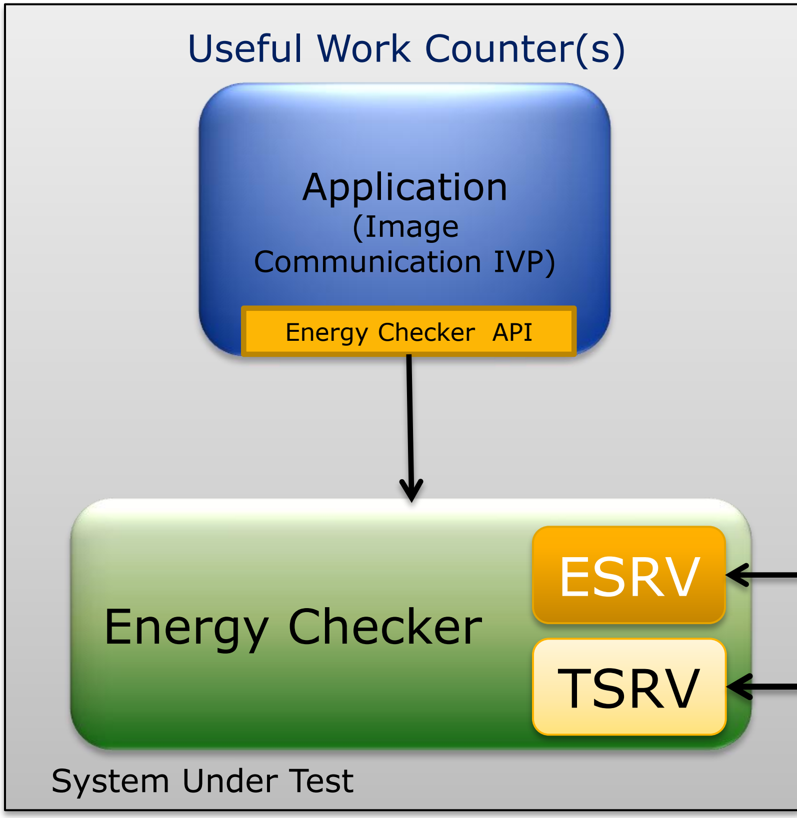

Measure energy consumed for a workload

- ➢Define “Work” to be measured
- ➢Instrument code
- ➢Collect data and compute Energy Efficiency
- ➢Analyze
    - •Sys1tem productivity
    - •Application’s energy profile

Write energy-aware SW with minimal effort, focusing on relevant energy heuristics

*Not included in EC SDK
•Use Recommended models or follow User Guide to enable another model

Energy-aware Programming

Major Pieces of the Intel® EC SDK

- Windows*, Linux*, Solaris* 10, and MacOS* X support
- Core API (C/C++, C#, and Java* interfaces)
- PL Scripting Tools
- ESRV/TSRV energy/temp monitoring tools
- Windows interoperability tools
- GUI Monitor and CSV Logger
- Additional sample code
- Installation code
- Documentation (3 manuals)

- See http://whatif.intel.com for details.

Energy-aware Programming
Example of Energy Checker Data Collected[1]

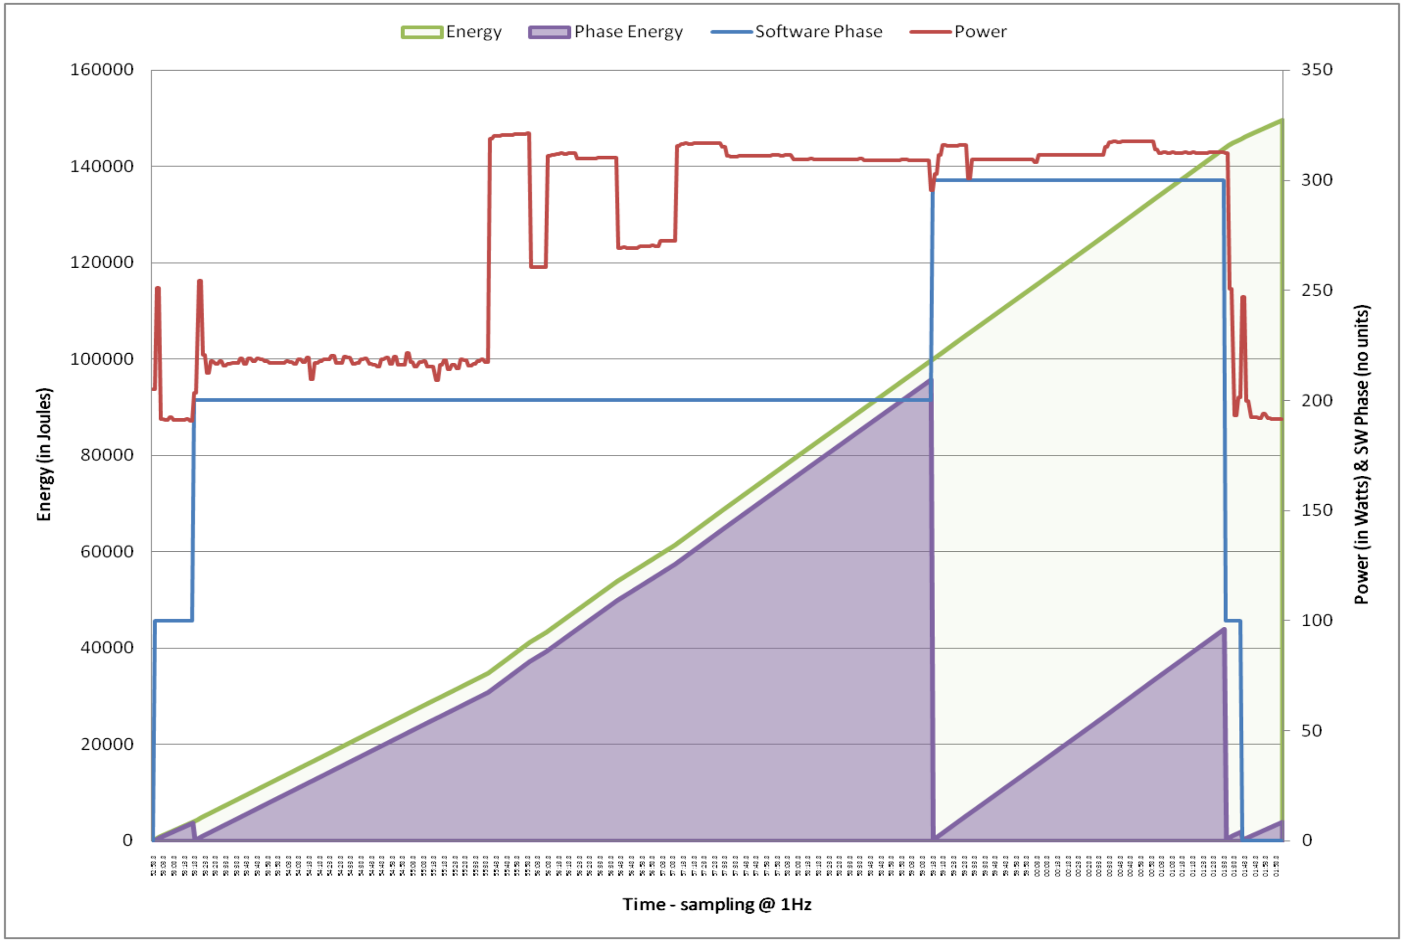

[1] Data was collected using pl_csv_logger (a tool shipped with the Intel® EC SDK). Instrumented from outside (scripting tools) and VBA [no source code access].

Summary

- Build your energy savings pyramid from bottom to top: optimize from macro to micro scale!
  - Each homework done in the lower layer maximizes leverage in the upper layer.
- **Macro-scale:** typical data center infrastructure offers many improvement opportunities for rather simple best practices.
- **Micro-scale:** energy-aware programming needs to be based on consistent measurement of executed work and Watt-hours consumed.