GreenSONAR

A multi-domain energy profiling system based on perfSONAR

Lutz Engels Todor Yakimov

University of Amsterdam
System and Network Engineering

February 8, 2013

Presentation overview

- Green IT
- Energy Profiling
- Proposed solution
- Conclusion
- Questions

The need for GreenIT

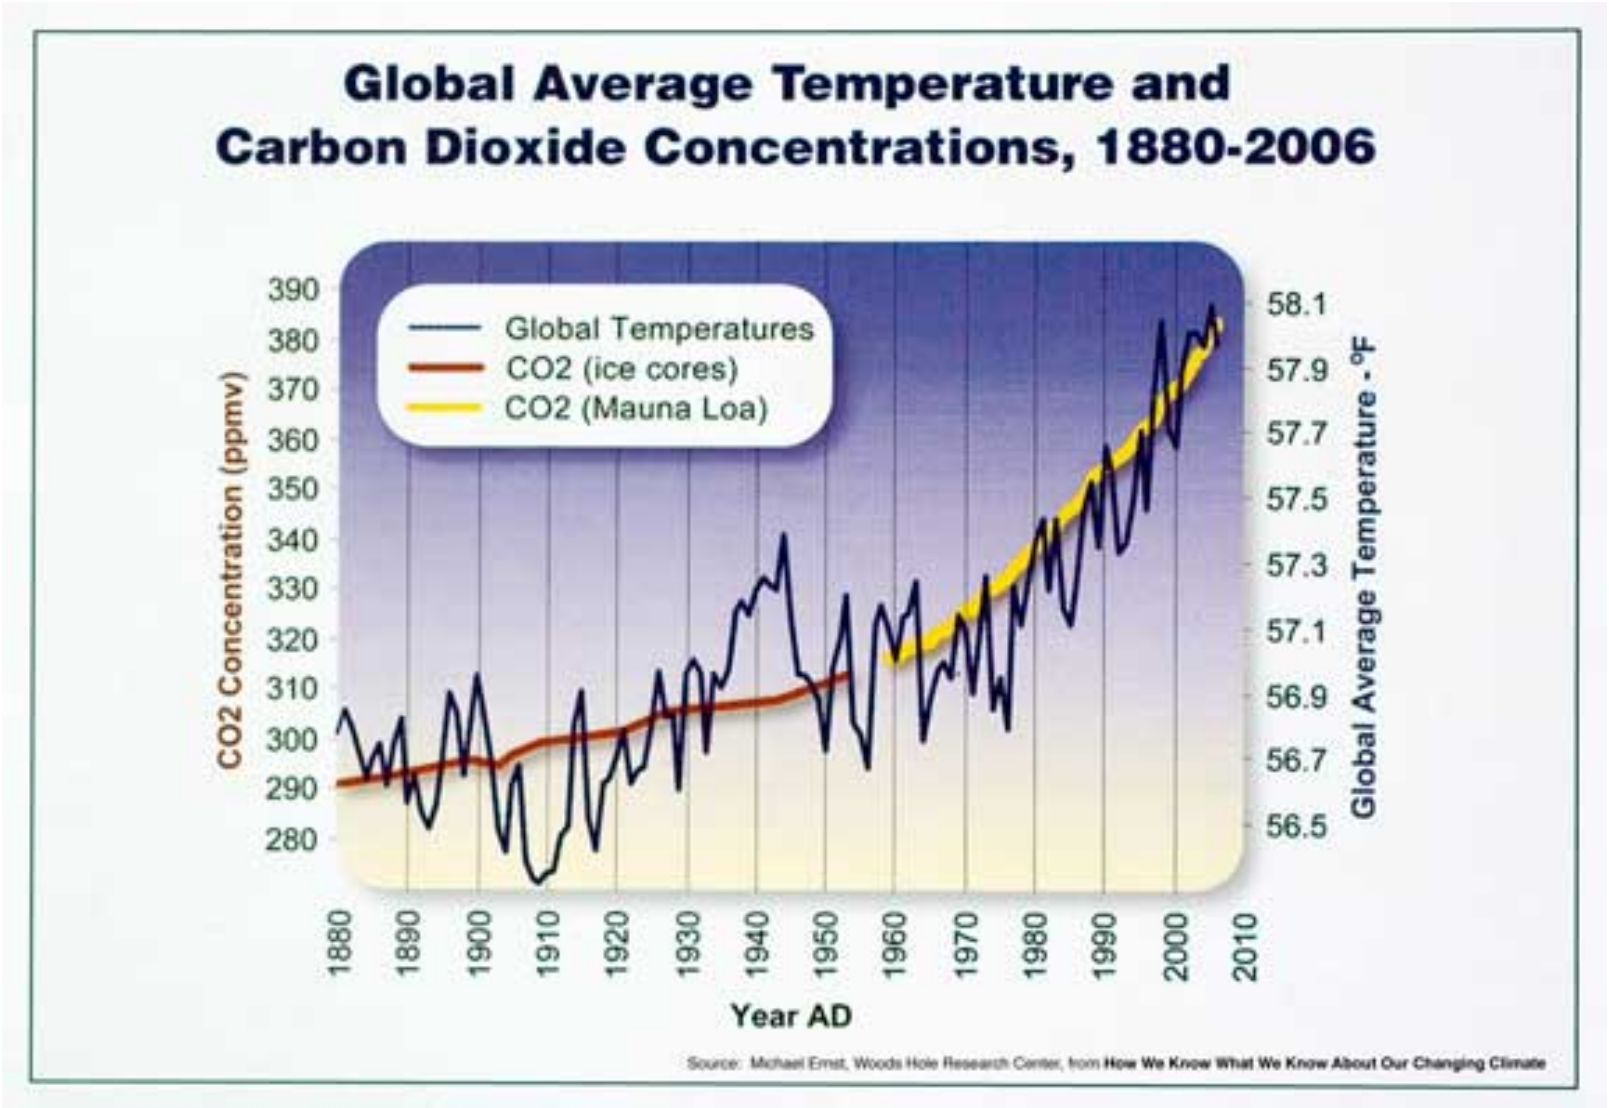

The need for GreenIT

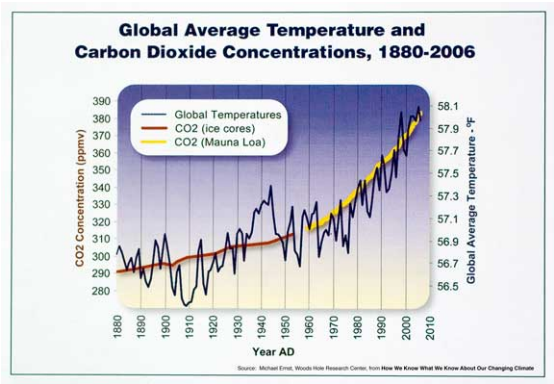

- ICT $\sim$2% of global energy consumption/CO_2-emission
- CO_2-emission: likely to overtake airline industry in near future

How does Green IT help?

- Implementing virtualization.
- Equipment consolidation.
- Upgrading to newer, more energy-aware equipment
- Employing management systems to increase efficiency and availability.

How can we further extend Green IT in Networking?

- 802.3az Energy-Efficient Ethernet
- Our focus:
  - find metrics as base for profiling
  - software to distribute the metrics/profiles
  - those lead to the following research questions

Research Question(s)

"What **metrics** need to be considered in order to build **energy profiles** of networking devices and how can such data be published by using **distributed multi-domain monitoring systems**."

Research Question(s)

"Is **perfSONAR-PS** a suitable architecture to achieve energy profiling of computational devices, and what are the **necessary steps** to be undertaken to evolve perfSONAR-PS in a system we can call '**GreenSONAR**'?"

Metric: absolute energy efficiency

- metric by *Parker et al.* (2011) to make efficiency comparable

$$dB\varepsilon = 10 log_{10} \left(\frac{Power/BitRate}{kTln2} \right)$$

- where:
  $dB\varepsilon$: absolute energy efficiency
  k: Boltzmann constant ($1.381 * 10^{-23} J/K$)
  T: temperature in Kelvin
  $kTln2$: absolute minimum energy per bit dissipated

Metrics

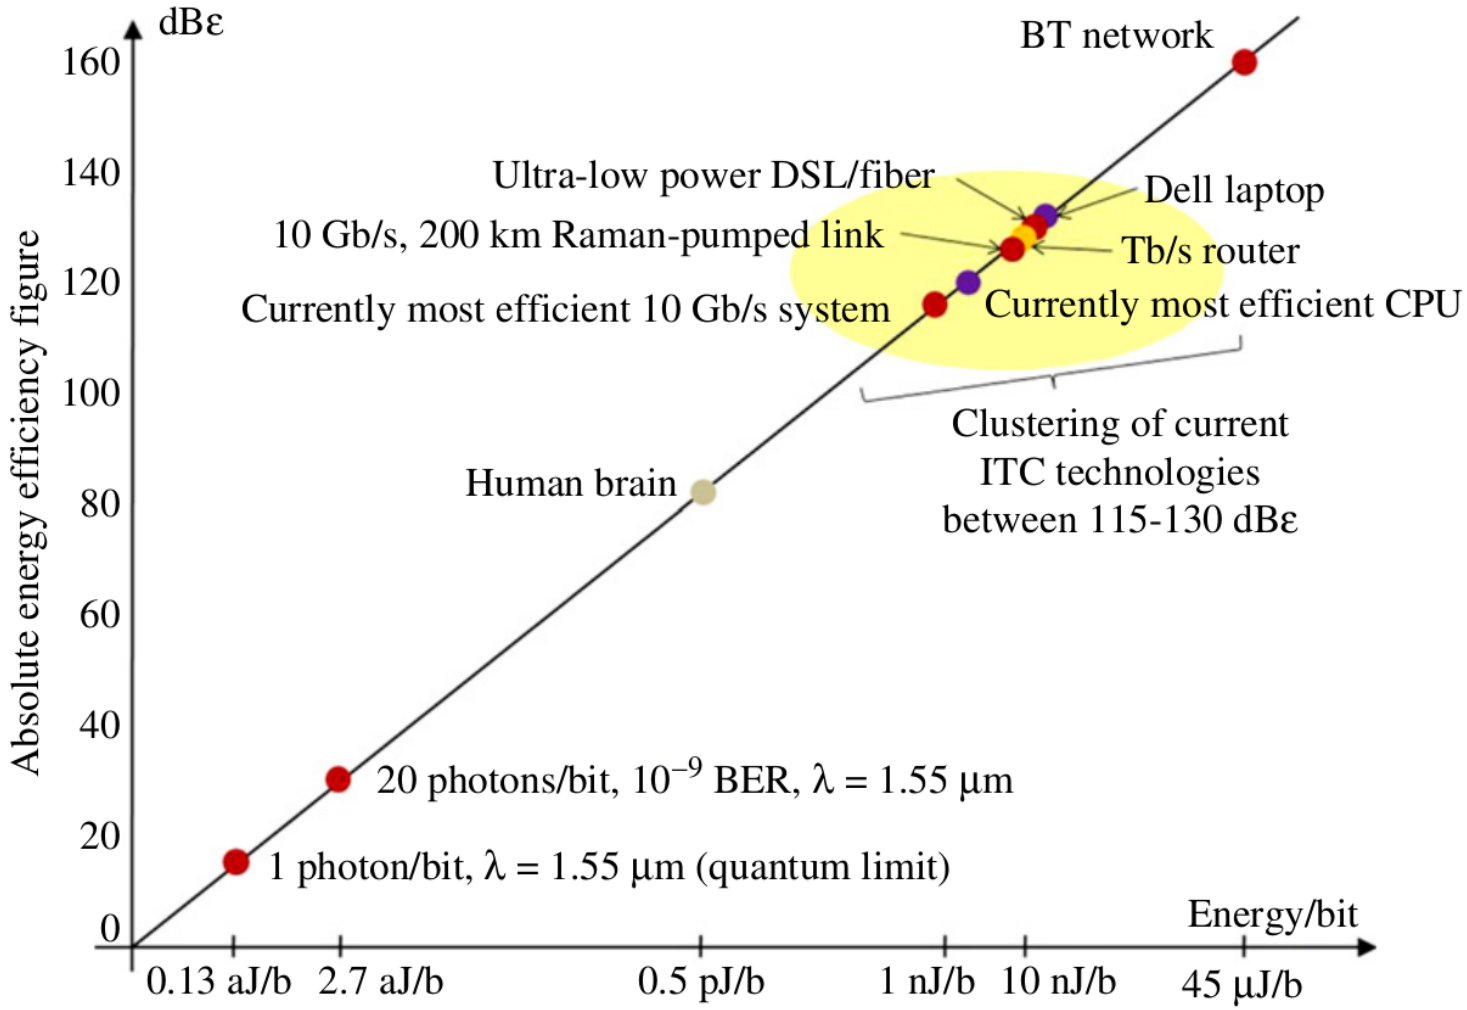

source: Parker et al., *Roadmapping ICT: An Absolute Energy Efficiency Metric* (2011)

derived: absolute current energy efficiency per port

Switch
100 Mb/s
123 W
24 ports

Utilisation:
Port 1: 10%
Port 2: 1%
Port 3: 80%
...

$$dB\varepsilon_{cpp} = 10log_{10}\left(\frac{\frac{P_{total}}{N_{ports}} / Util_p * Speed_{max}}{kTln2}\right)$$

Metrics

Switch
100 Mb/s
123 W
24 ports

Utilisation:
Port 1: 10%
Port 2: 1%
Port 3: 80%
...

$$dB\varepsilon_{cpp} = 10log_{10}\left(\frac{\frac{P_{total}}{N_{ports}} / Util_p * Speed_{max}}{kTln2}\right)$$

Metrics

Switch
100 Mb/s
123 W
24 ports

Utilisation:
Port 1: 10%
Port 2: 1%
Port 3: 80%
...

$$dB\varepsilon_{cpp} = 10log_{10}\left(\frac{\frac{P_{total}}{N_{ports}} / Util_p * Speed_{max}}{kTln2}\right)$$

Metrics

Switch
100 Mb/s
123 W
24 ports

Utilisation:
Port 1: 10%
Port 2: 1%
Port 3: 80%
...

absolute min. energy/bit dissipated
with k = 1.381 * 10^{-23} J/K
and T = 300 K

$$dB\varepsilon_{cpp} = 10log_{10}\left(\frac{\frac{P_{total}}{N_{ports}} / Util_p * Speed_{max}}{kTln2}\right)$$

Metrics

Switch
100 Mb/s
123 W
24 ports

Utilisation:
Port 1: 10%
Port 2: 1%
Port 3: 80%
...

absolute min. energy/bit dissipated
with k = 1.381 * 10^{-23} J/K
and T = 300 K

$$dB\varepsilon_{cpp} = 10log_{10}\left(\frac{\frac{P_{total}}{N_{ports}} / Util_p * Speed_{max}}{kTln2}\right)$$

Metrics

Decreasing $dB\varepsilon_{cpp}$ (=higher efficiency) on higher utilization

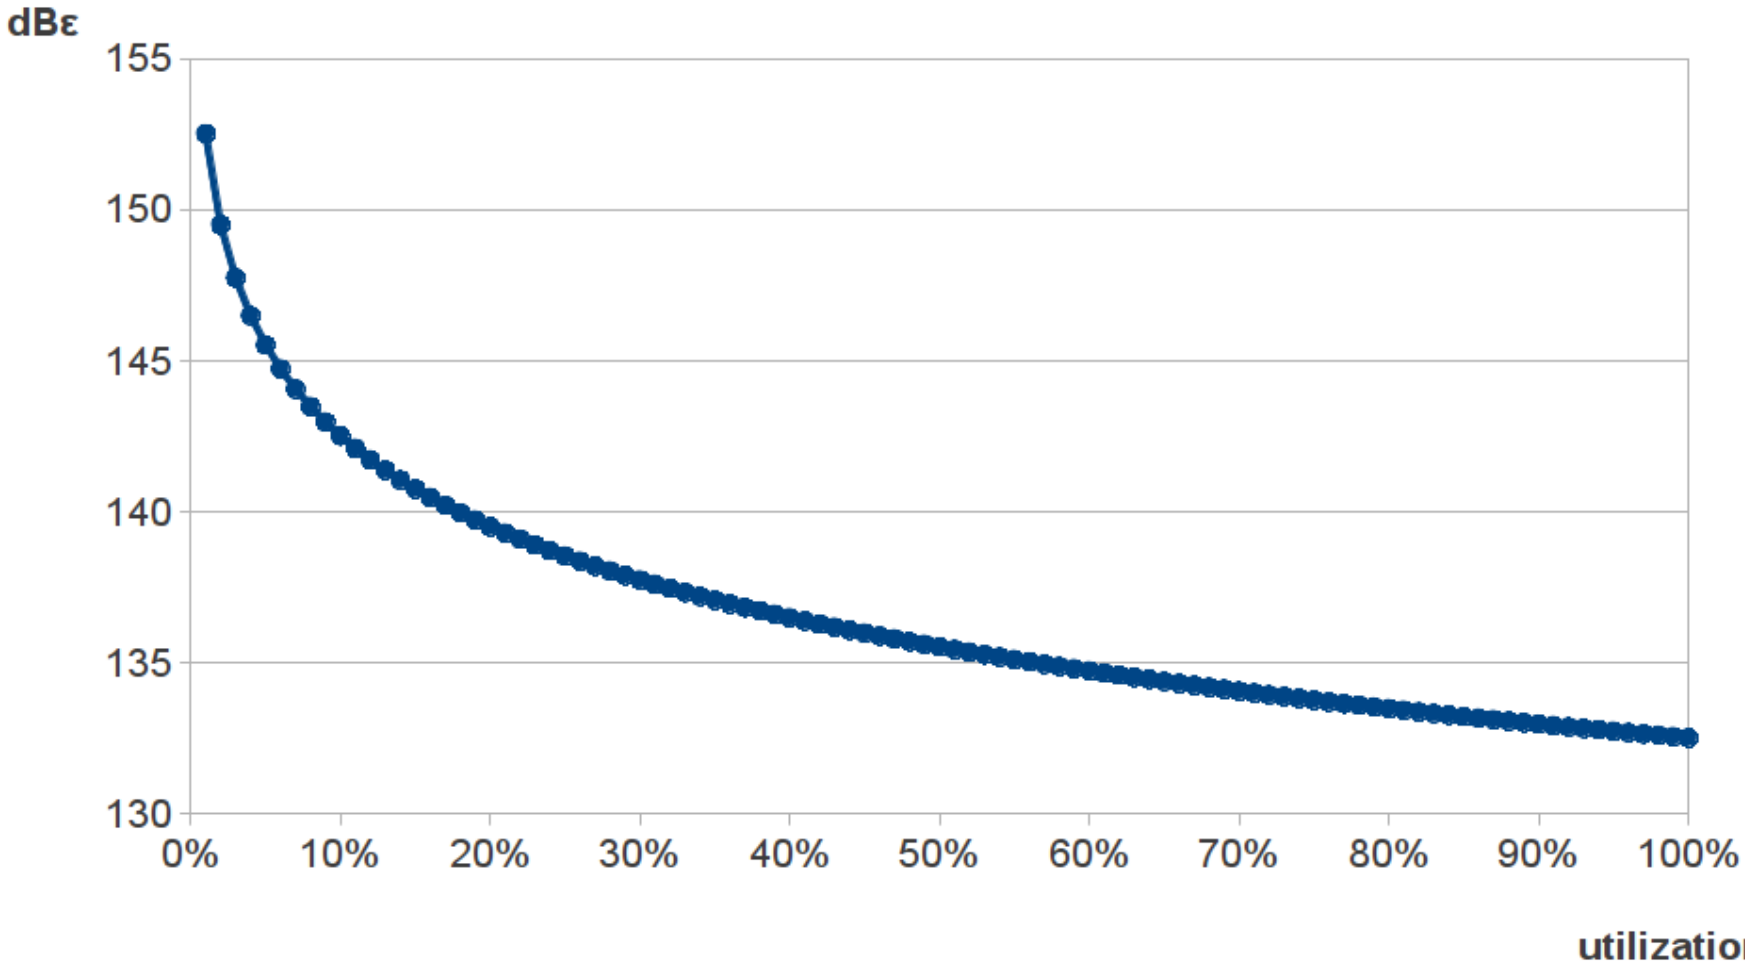

Gathering data needed for the metric

- Power Distribution Units provide energy data
- port utilization can be retrieved by SNMP
- derived metric can be used as 'path cost' for path decisioning

Data dissemination: perfSONAR

- A conjoint effort between the EU-funded GN2 JRA1 project, Internet2 and ESnet
- Facilitates multi-domain monitoring by removing administrative limitations
- Implements a network monitoring Service-Oriented Architecture (NMSOA)

Data dissemination: The perfSONAR framework

A three-tier architecture:

- (L1) Measurement Point Layer (MP) - Domain-specific interface wrappers around existing measurement tools
- (L2) Service Layer - Inter-domain exchange of measurement data and management information
  - Measurement Archives (MA)
  - Lookup Service (LS) - MA discovery
  - Authentication Service (AS)
- (L3) User Interface Layer - Reporting and visualization tools

Data dissemination: The perfSONAR framework

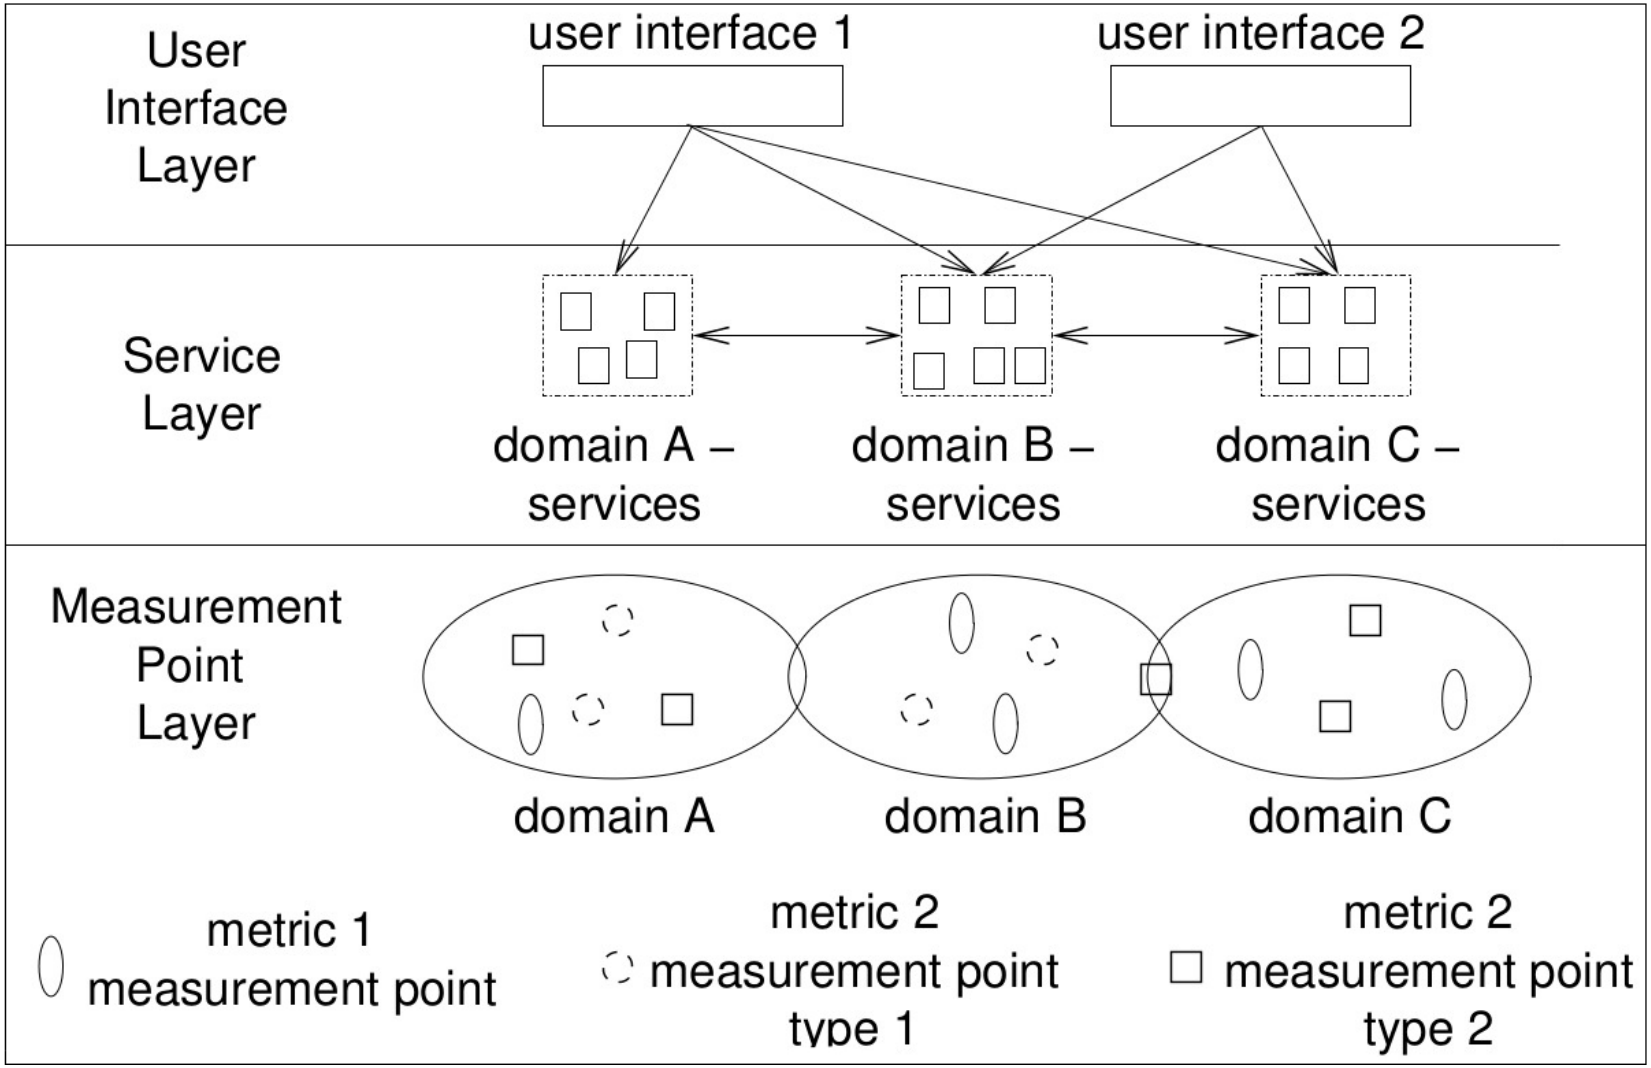

Data dissemination: perfSONAR releases analysis

Comparison of releases

Feature	perfSONAR MDM/PS	perfSONAR NC
Services	MP, MA, LS, AS, TS	MP
Protocol	NMWG	NETCONF
Data models	none	Generic model for each service defined using YANG
Query Mechanisms	MA/MP dependent	Xpath
Scalability	~1000 domains	>200K domains
Performance	Slow (seconds)	Fast (milliseconds)
Code base	>6000, MA specific ~3000	<1000, MA specific <200
Programming Language	Java and Perl	PHP

PS/MDM protocol

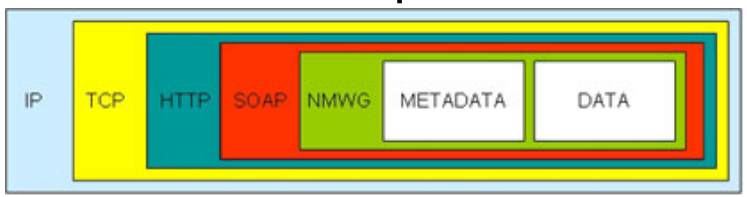

Data dissemination: perfSONAR analysis steps

perfSONAR PS/MDM/NC documentation

- Installation
- Configuration
- Code base
  - MPs/MAs
  - Protocol management information

Test Setup

- Two perfSONAR instances
- Measurement Points that wrap around RRD archives for reading
  - PDU data
  - Network utilization (SNMP) data

Data dissemination: perfSONAR analysis steps

Results PS/MDM

- No generic APIs for defining new MPs/MAs;
- No documentation kept on the code base
- Reverse-engineering the code base was not achievable within timeframe

Results NC

- Easily extensible MP facilities
- Basic implementation for trasnfering utilisation data, gathered by SNMP

General Results

- Fast query responses
- However, a lot of XML overhead

Proposed Traffic flow model

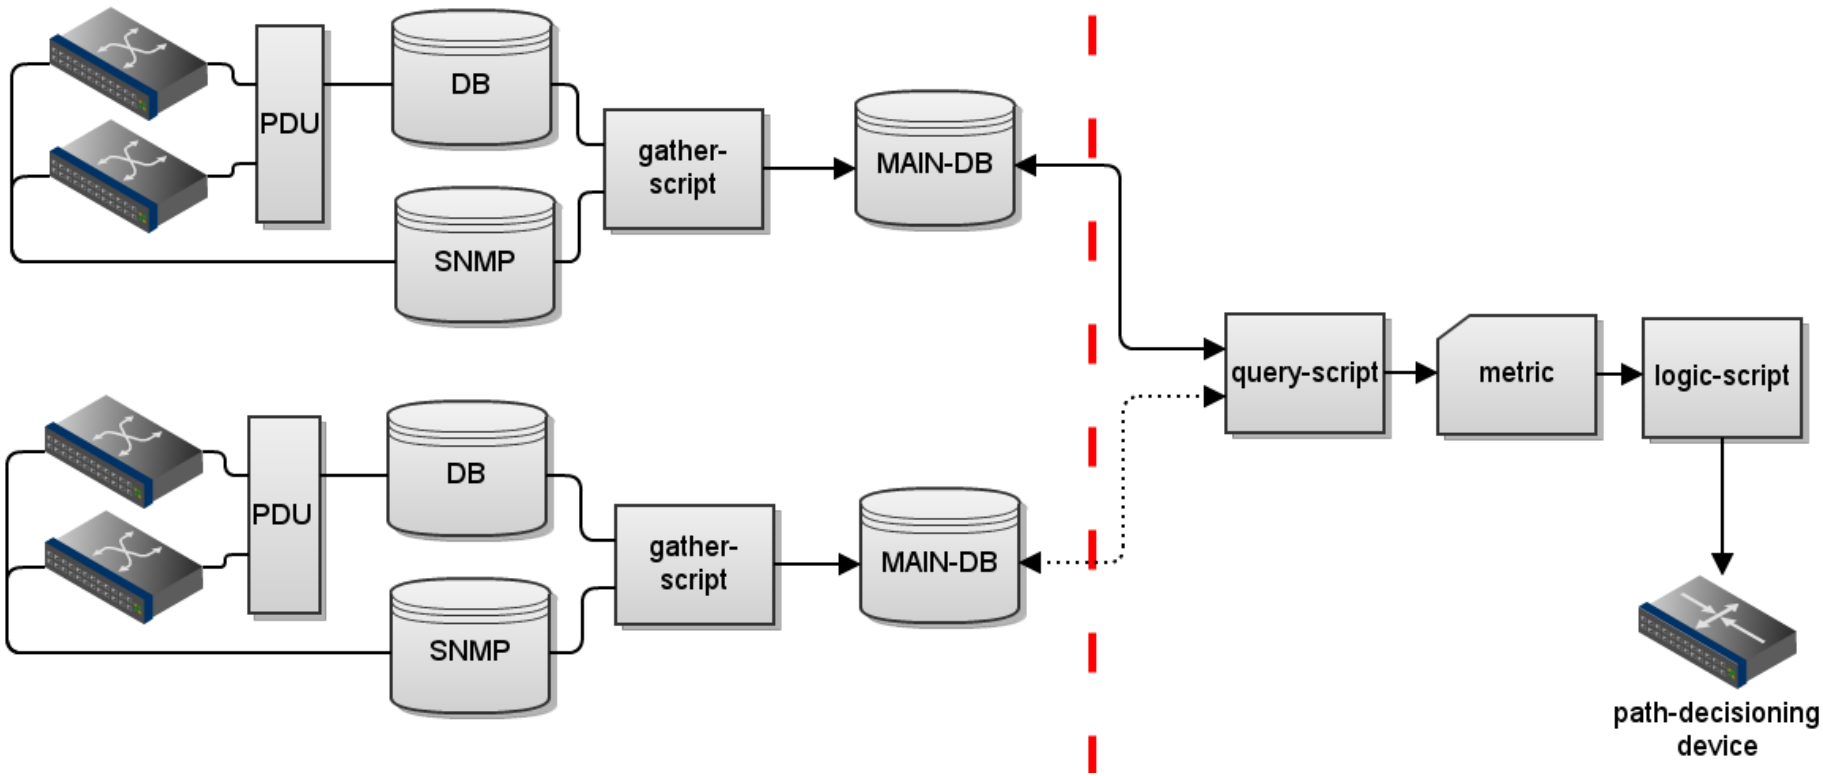

Conclusions

PerfSONAR applicability

- Hard to implement new functionality due to lack of documentation
  - Solution: cooperation with current developer team needed
- A large code base that does not follow a data model
  - Solution: cooperation needed
- Large XML overhead when requesting small data sets
  - No straightforward solution

Summary

"What metrics need to be considered in order to build energy profiles of networking devices and how can such data be published by using distributed multi-domain monitoring systems."

- $dB\varepsilon_{cpp}$

"Is perfSONAR-PS a suitable architecture to achieve energy profiling of computational devices, and what are the necessary steps to be undertaken to evolve perfSONAR-PS in a system we can call 'GreenSONAR'?"

- **Partially, because...**

Questinons

References

- perfSONAR framework:
  `http://www.perfsonar.net/`
- Related projects:
  `http://ext.delaat.net/smartgreen/index.html`