Electronic Devices & Networks Annex EDNA

The Idle Coefficients

KPIs to assess energy wasted in servers and data centres

SEPTEMBER 2021

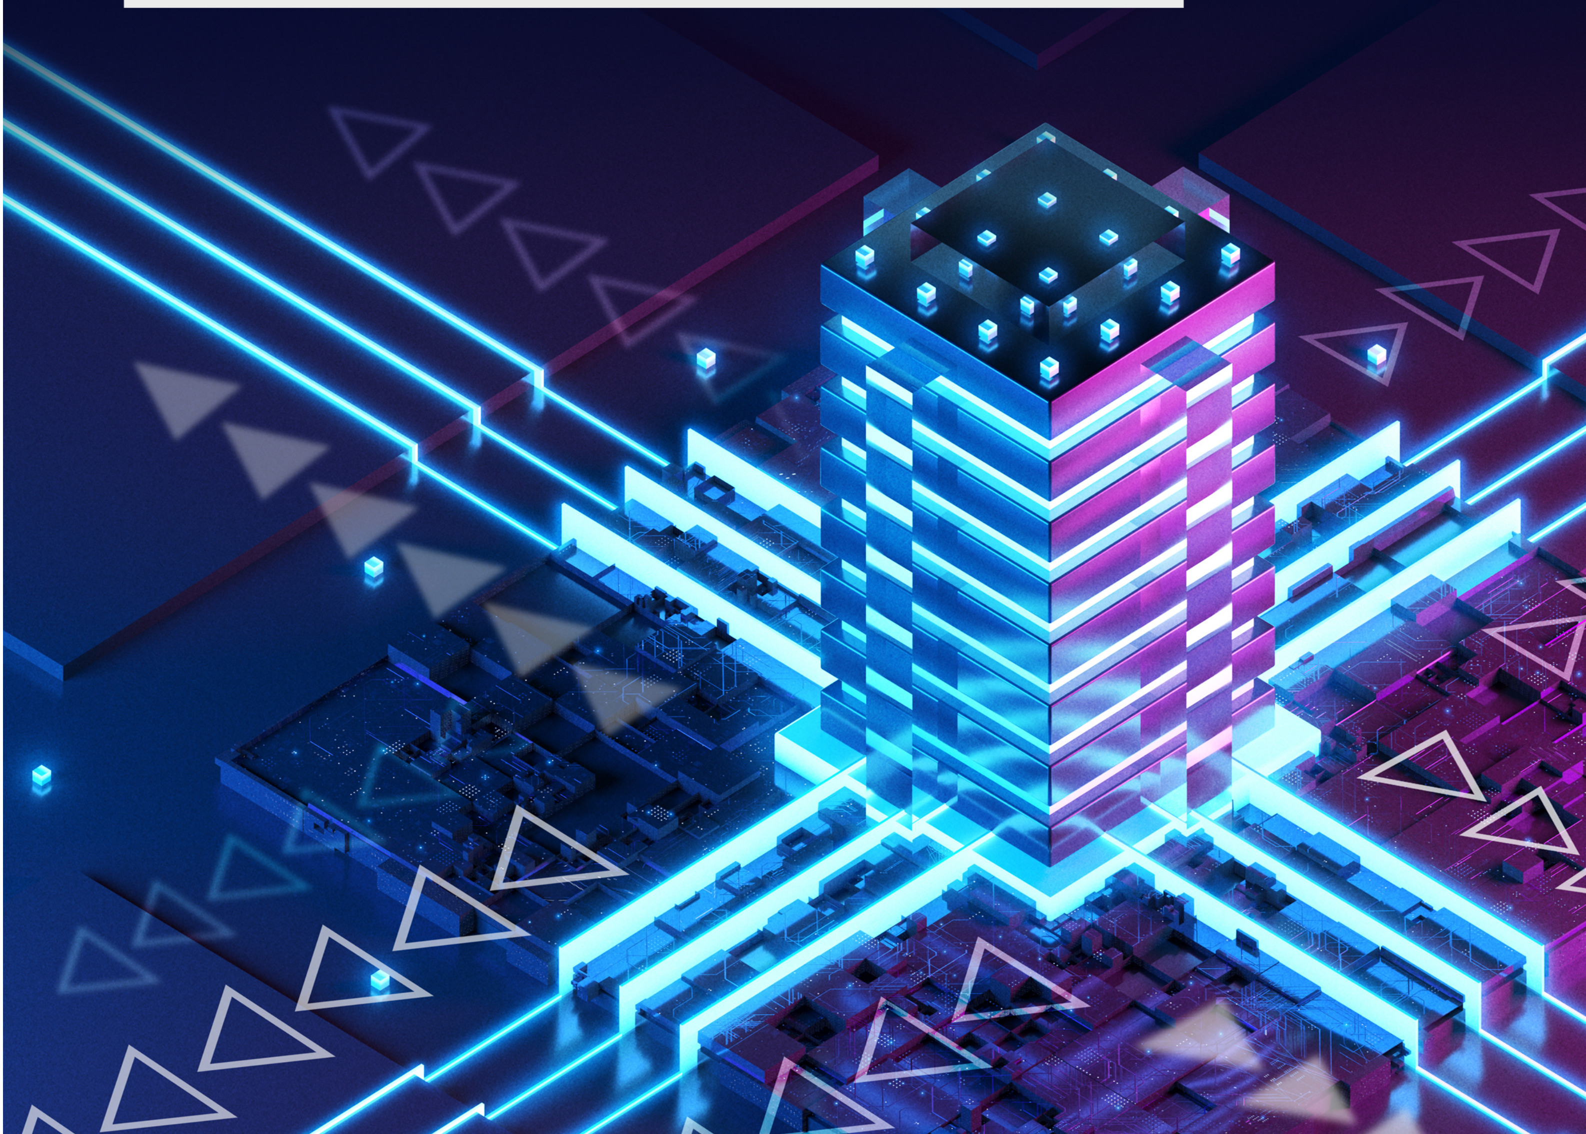

Technology Collaboration Programme
by iea

The Technology Collaboration Programme on Energy Efficient End-Use Equipment (4E TCP), has been supporting governments to co-ordinate effective energy efficiency policies since 2008.

Fifteen countries have joined together under the 4E TCP platform to exchange technical and policy information focused on increasing the production and trade in efficient end-use equipment. However, the 4E TCP is more than a forum for sharing information: it pools resources and expertise on a wide a range of projects designed to meet the policy needs of participating governments. Members of 4E find this an efficient use of scarce funds, which results in outcomes that are far more comprehensive and authoritative than can be achieved by individual jurisdictions.

The 4E TCP is established under the auspices of the International Energy Agency (IEA) as a functionally and legally autonomous body.

Current members of 4E TCP are: Australia, Austria, Canada, China, Denmark, the European Commission, France, Japan, Korea, Netherlands, New Zealand, Switzerland, Sweden, UK and USA.

Further information on the 4E TCP is available from: **www.iea-4e.org**

The EDNA Annex (Electronic Devices and Networks Annex) of the 4E TCP is focussed on a horizontal subset of energy using equipment and systems - those which are able to be connected via a communications network. The objective of EDNA is to provide technical analysis and policy guidance to members and other governments aimed at improving the energy efficiency of connected devices and the systems in which they operate.

EDNA is focussed on the energy consumption of network connected devices, on the increased energy consumption that results from devices becoming network connected, and on system energy efficiency: the optimal operation of systems of devices to save energy (aka intelligent efficiency) including providing other energy benefits such as demand response.

Further information on EDNA is available at: **iea-4e.org/edna**

This report was commissioned by the EDNA Annex of the 4E TCP. It was authored by Dr. D.H. Harryvan of Certios. The views, conclusions and recommendations are solely those of the authors and do not state or reflect those of EDNA, the 4E TCP or its member countries.

Views, findings and publications of EDNA and the 4E TCP do not necessarily represent the views or policies of the IEA Secretariat or its individual member countries.

The IEA, 4E, EDNA and the authors make no conclusions, endorsements or disendorsements in relation to any organisations or brands mentioned in this report.

THE IDLE COEFFICIENTS

KPIs to assess energy wasted in servers and data centres

Dr. D.H. Harryvan (Certios)

For

IEA 4E Technology Collaboration Programme
Electronic Devices and Networks Annex (EDNA)

VERSION | STATUS	1.0 | Final version
DATE	September 27th, 2021
DISSEMINATION LEVEL	Open
AUTHORS (PARTNER)	D.H. Harryvan (Certios B.V.)

EXECUTIVE SUMMARY

Addressing a request from EDNA, this document presents information on novel data centre KPIs, the Server Idle coefficient (SIC) and the Data Centre Idle Coefficient (DCIC). The goal of this document is to inform EDNA members about the existence of these metrics, how these are calculated and what the impact of the use of these metrics could be, for monitoring and improving data centre energy effectiveness, potentially lowering the total energy consumption of a data centre. This total energy consumption can be expressed as:

$$\boldsymbol{E(DC) = \ E(IT)x\ PUE}$$

where

- *E(*DC) is the total data centre energy consumption (annual) in kWh;
- *E(*IT) is the total IT equipment energy consumption (annual) in kWh;
- PUE is the power usage effectiveness (annual).

Because the improvements made on the PUE over the past years are such that further major improvements are no longer economical, a logical next step is targeting E(IT).

The complete ICT hardware architecture inside a data centre is comprised of servers, storage and networking equipment and the resulting IT equipment energy consumption is the basis for all energy usage in a data centre.

$$\boldsymbol{E(IT) = \sum Server\ Energy + \sum Storage\ Energy + \sum Networking\ Energy}$$

Where

- *E*(IT) is the total IT equipment energy consumption in kWh;
- $\sum \boldsymbol{Server\ Energy}$ is the total energy consumption of all the servers in kWh.
- $\sum \boldsymbol{Storage\ Energy}$ is the total energy consumption of all the storage equipment in kWh.
- $\sum Networking\ \boldsymbol{Energy}$ is the total energy consumption of all the networking equipment in kWh.

The reason for targeting energy use in servers is that, as was published in the ICT impact study prepared for the European Commission (VHK and Viegand Maagoe, 2020) servers are responsible for over 80% of the total IT energy use in a data centre. Since PUE is by definition greater then 1, reducing energy use in servers is a very effective way of reducing the total energy use by data centres. The KPIs proposed in this document quantify the energy waste in servers and the reductions in this part of the energy use do not hamper the productivity of an IT infrastructure, regardless of what this productivity entails or how it is expressed. Reducing the energy waste will result in reduction of the total IT energy assuming that the workload is unchanged. When workload increases, the energy waste reduction will slow the increase in energy demand.

The SIC / DCIC metrics take a different approach than most other data centre KPIs in that the proposed metrics are **in**effectiveness metrics. Determining **in**effectiveness rather than (in)efficiency is grounded in the fact that there is no known generic metric indicating the amount of work that is done by a data centre. Therefore it is impossible to define an efficiency metric that would by definition be of the format; Unit of Work per Unit of Energy. For servers however, there is a single identifiable process that is common to all servers, regardless of make, model and even architecture, that indicates that the server has no useful workload to run. This process is known as ‘Idle’. It is possible to calculate the energy used for running these idle cycles and express this as a percentage of the total energy use of the server:

$$\boldsymbol{SIC} = \frac{\boldsymbol{Server\ Energy\ (Idle)}}{\boldsymbol{Server\ Energy\ (total)}} \text{ x } 100\%$$

Expanding this for the total data centre to:

$$\boldsymbol{DCIC} = \frac{\sum Server\ Energy\ (Idle)}{\sum Server\ Energy\ (total)} \text{ x } 100\%$$

Obtaining the data needed for the SIC calculation requires the cooperation of the owner of the server, who is often not the same as the owner of the data centre. However, the data needed is simple, limited in volume and does not contain or refer to application data. The pilot described in this document proves the feasibility of collecting the needed data and illustrates the usefulness of calculating the SIC from this data.

The pilot data showed values for the SIC that ranged from 50% (best case) to over 90% (worst case). These values indicate a huge potential for savings, that can be targeted using well known practices, used in modern IT environments. The SIC also hints to a possible means of evaluating IT hardware efficiency through the equipment's electrical power draw in relation to workload (or lack thereof). As such the SIC represents a potential useful metric whose usefulness will increase with increased adoption and experience in interpreting results from data centres across the globe.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

The energy use of data centres is a growing concern for policy makers. This energy use is significant and growing, despite the improvements in energy efficiency[1] of the ICT equipment, servers, storage devices and networking equipment, and despite the continuing improvements in the efficiency of data centre facility infrastructure (mainly cooling). An estimation of data centre energy use, including a breakdown in these elements is shown in table 1. This information was published in an ICT Impact study done for the European Commission by VHK and Viegand Maagøe in 2020 (VHK and Viegand Maagoe, 2020).

ICT Electricity Use EU27 *(in TWh/year)*	**2010**	**2015**	**2020**	**2025**
Servers	18.66	18.66	22.05	27.24
Storage	1.80	1.80	4.35	4.45
Networks	0.53	0.53	0.74	1.06
Cooling etc.	23.74	23.74	12.40	10.07
Total Data Centres	**44.73**	**44.73**	**39.54**	**42.82**

Table 1 : Estimated electricity use of data centres, including breakdown (VHK and Viegand Maagoe, 2020)

Policy making for data centres has proven itself to be difficult, mainly for two reasons;

1) A data centre is a system
2) It is difficult to impossible to quantify the function of a data centre.

Despite these difficulties, a large number of data centre key performance indicators (KPIs) and data centre standards have been developed in an attempt to both quantify data centre efficiency as well as improve upon it. However, the lack of policy puts most of these efforts into the domain of 'voluntary actions'.[2]

Perhaps the best known and internationally recognized KPI for data centres bypasses the aforementioned difficulties by solely quantifying the effectiveness of the data centre facility infrastructure:

- PUE Power Usage Effectiveness, ISO 30134-2

The PUE was originally developed by the Green Grid, however, it has been transferred to the ISO organization. It has been published under ISO 30134-2 (ISO, 2016) and is available from https://www.iso.org/.
The exact same information is also contained in the European norm, EN 50600-4-2 (NEN, 2017).

$$PUE = \frac{E(DC)}{E(IT)}$$ **Eq 1: PUE**

where
- *E*(DC) is the total data centre energy consumption (annual) in kWh;
- *E*(IT) is the total IT equipment energy consumption (annual) in kWh;
- PUE is the power usage effectiveness (annual).
-

Continuous attention of the industry, customers and governments for this metric has resulted in a marked improvement of the facility infrastructure of data centres. The data in table 1 shows that averaged PUE in 2010 was 2.1 and has dropped to 1.45 in 2020. It is expected to improve to 1.3 in 2025. The ideal value of a PUE is 1, meaning that there is absolutely no other energy use than that of the ICT equipment. However, since the PUE calculated here represents an average over all of Europe, including areas with a warmer climate, further improvement below 1,3 will be difficult. As a consequence, the growth in E(IT) -

[1] While there is no known generic metric indicating the amount of work that is done by a complete data centre, the efficiency of a single piece of it equipment can be determined through the use of benchmark workloads. These benchmarks show for instance that newer generations of servers use less energy per computation than older generations. (Standard Performance Evaluation Corporation, 2021)

[2] An exception to voluntary use of PUE is made in the Amsterdam region. In Amsterdam a building permit is only supplied to a data centre if this data centre can prove that a PUE <1,2 will be achieved. Data centres are also obliged to report the annual PUE to the appropriate authorities (regionale uitvoeringsdienst noord holland noord, 2021).

mostly caused by the increase in energy use by servers - will result in an increase in total data centre energy use over the coming years (see table 1).

A closer examination of the first difficulty mentioned reveals a secondary observation:

> All energy usage in the data centre is linked to the primary function of the data centre, the running of IT equipment. Since the data centre is a system, reduction of the energy usage by the IT equipment will lead to a larger reduction of the total energy used by a data centre.

Bearing in mind that the PUE is by definition greater than 1, the observation is substantiated when rewriting the PUE to show the direct relation of the IT energy usage with the total data centre energy usage:

$$\boldsymbol{E(DC) = E(IT).PUE}$$ **Eq 2: Total data centre energy**

where
- *E(*DC) is the total data centre energy consumption (annual) in kWh;
- *E(*IT) is the IT equipment energy consumption (annual) in kWh;
- PUE is the power usage effectiveness (annual).

Since PUE optimizations from facility improvement are no longer expected, data centre energy usage can only be curbed by optimizing the E(IT). Therefore, there exists a need for a KPI that can spark a drive for IT energy savings. This paper discusses two candidates for such KPIs, the Server Idle Coefficient (SIC), and the cumulative version, the Data Centre Idle Coefficient (DCIC).

2 THE SERVER IDLE COEFFICIENT, SIC

As mentioned in the introduction, there is need for increasing the effectiveness of the ICT infrastructure. Fortunately, the technological means for accomplishing this goal are already in place.

2.1 ADVANCED CONFIGURATION AND POWER INTERFACE

Any server in operation today can be operated in a **dynamic power state**. Simply put, a server has the ability to match its electrical power usage to its ICT workload in some degree. The control of the dynamic range is either in the hardware itself (through BIOS settings) or in the base Operating System (OS)[3] running on the hardware.

It is important to note that the ability itself and the degree of the dynamic response is determined by the system administrator through power management settings. The control mechanisms OS and BIOS are presented to the system administrators simultaneously, the correct settings for OS controlled power management would be a BIOS setting 'OS-controlled', followed with an appropriate setting within the OS. If any BIOS setting other than OS-controlled is chosen, the firmware will be in control of the power management.

Power management allows the computer hardware components to be put in a variety of low power modes when the demand on the component is low. These low power modes are described by so called Advanced Configuration and Power Interface (ACPI) states. The ACPI provides an open standard that operating systems can use to discover and configure computer hardware components, to perform power management by (for example) putting unused components to sleep, and to perform status monitoring. First released in 1996, ACPI brings the power management under the control of the Operating System, as opposed to the previous BIOS-centric system that relied on platform-specific firmware to determine power management and configuration options. The ACPI specification is central to the Operating System-directed configuration and Power Management (OSPM) system, an implementation for ACPI which removes device management responsibilities from legacy firmware interfaces via a UI.

Intel, Microsoft and Toshiba originally developed the ACPI standard, while HP, Huawei and Phoenix participated later. In October 2013, the ACPI Special Interest Group (ACPI SIG) -the original developers of the ACPI standard - agreed to transfer all assets to the UEFI Forum, in which all future development will take place. The UEFI Forum published the latest version of the standard, 'Revision 6.3', in January 2019 (UEFI, 2019).

Processor power management technologies are defined in the ACPI specification and are divided into two categories or states[4] (Microsoft, 2018):

Power performance states (ACPI P states)

P-states provide a way to scale the frequency and voltage at which the processor runs so as to reduce the power consumption of the CPU. The number of available P-states can be different for each model of CPU, even those from the same family.

Processor idle sleep states (ACPI C states)

C-states are states when the CPU has reduced or turned off selected functions. Different processors support different numbers of C-states in which various parts of the CPU are turned off. To better understand the C-states that are supported and exposed, contact the CPU vendor. Generally, higher C-states turn off more parts of the CPU, which significantly reduce power consumption.

[3] The OS is meant here in its broader term: Hypervisors like VMware ESX ® or Microsoft Hyper V ® are just as valid as any MS Windows ® , Linux or Unix OS.

[4] The current document is not intended as a detailed technical discussion on ACPI states, an excellent detailed description/explanation can be found on the web as "A Minimum Complete Tutorial of CPU Power Management, C-states and P-states" (metebalci.com/, 2021)

Server power management has several **steps**:

- High Performance (HP) – causing ACPI C-states to be locked in C0. This means that little energy is saved when the server's CPU is idle. In many cases adjustments will still be made to the clock frequencies. These adjustments fall under the so-called ACPI-P states. These adjustments happen if a CPU is not entirely idle, but underloaded.
  Some hardware manufacturers offer the option of a 'static high performance', in which case all CPU cores will run at a fixed frequency, the so-called thermal design frequency.
  Paradoxically, the high-performance setting does not guarantee maximum performance. It is not widely known that disabling ACPI C states above C0 will block the use of turbo frequencies, limiting the performance of single CPU cores.

- Additional power management steps - many servers have multiple power management settings. These can be specific per brand and type of server. They aim to achieve increasing energy savings, as the need for CPU capacity decreases further and / or one or more core (s) are switched off (deeper). For example, CPU states:
  - C0: active;
  - C1: least aggressive form of CPU sleep state. The Wake-up time of a switched core is roughly 0.5 micro-seconds;
  - C6: heaviest C-state, All caches are flushed and the CPU has no power at all.
    Wake-up time from C6 is roughly 40 micro-seconds.

If the CPU 'runs' at 3.3 GHz, then wake-up from high C-states to C0 is within 1.650 – 13.200 clock cycles (Schöne, 2015).

The first thing to realize is that for a CPU to be put into a C6 state, this particular CPU must not have been used for a considerable amount of time. ACPI-C states apply to idle CPU's only, the quoted wakeup latency is a delay that only happens once, when an inactive CPU needs to be added to the pool of active CPU's.

The second thing to realize is that various other delays occur during the run time of an application. The response time of a hard disk is in the range of 10 ms and network traffic can also introduce micro-seconds of delays. Even without any handling logic (send/receive), the round-trip time over 100 meters of optical fibre is 1 micro-second. It is fair to conclude that it would be impossible for an end-user to detect an additional 40 micro-seconds delay in the response time of an application.

The working of ACPI states (power management) seems particularly useful when taking into account the total network traffic profile, e.g., such as those published by the Amsterdam internet exchange (AMS-IX) (see Figure 1).

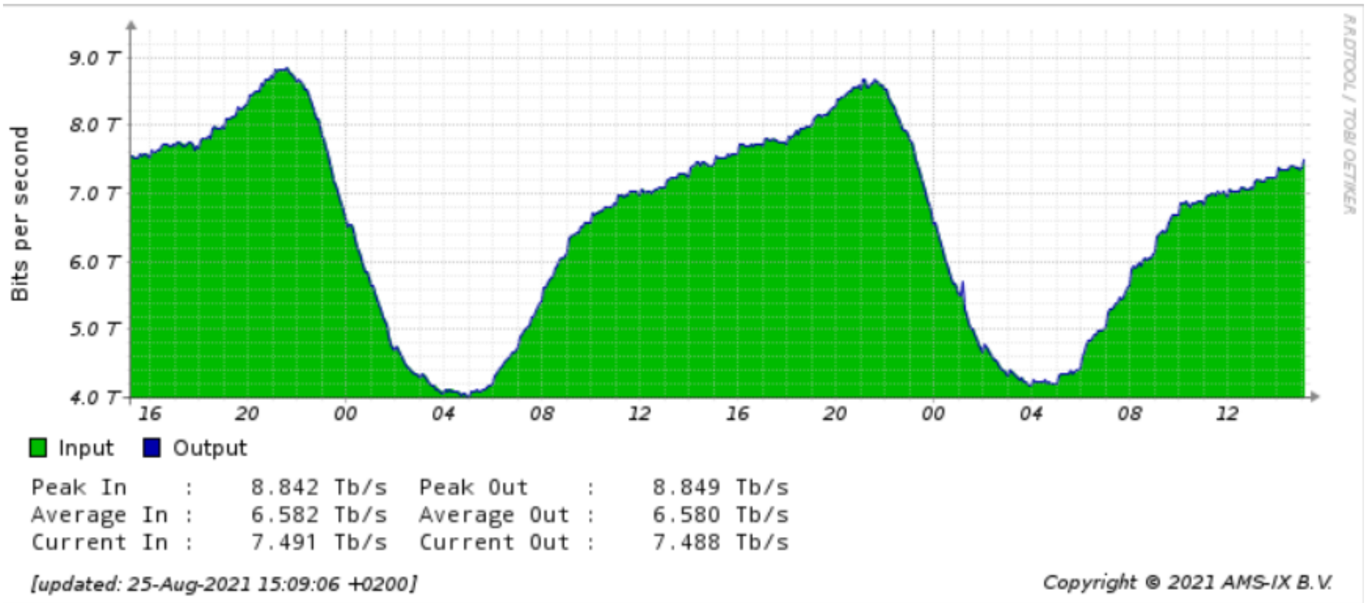

Figure 1: Amsterdam internet exchange(AMS-IX) daily traffic (AMS-IX B.V., 2021)

This graph shows network traffic over two periods of 24 hours and demonstrates the huge difference in internet traffic over one given day. The peak of 8,8 Tb is more then double the dip in throughput of close to 4 Tb per second. Network traffic can be seen as a proxy for activity and it is therefore reasonable to assume that a similar variation in server CPU load accompanies the dramatic variations in network traffic. Through the use of the server power management states, servers can lower their energy use when the workload decreases.

2.2 DEFINING THE SIC

When targeting the energy use of the IT equipment, the concept of 'work' re-enters the discussion. However, restating the difficulty number 2 in the introduction:

> There is no generic metric for the useful work realized by the data centre[5]. There is no consensus on the definition of 'useful' nor the definition of 'work'. As a result, it is impossible to define an efficiency metric that would by definition be of the form; Unit of Work per Unit of Energy.

The efficiency of servers can be expressed by benchmark workloads in a unit of Work per unit of Energy, but the workloads in an operational data centre are diverse and importantly, do not tax the IT equipment to its maximum. While benchmarks do help to distinguish various models of servers, they give no indication of the efficiency of the data centre deploying these server models.

Since "useful work" is not defined, dr. D.H. Harryvan, during a power management pilot run in the Netherlands, used the absence of work, idleness, to indicate when server resources are not being used effectively. (D.H.Harryvan, 2021) If a server cannot schedule tasks (workload) from applications, the CPUs run an idle loop. **These idle cycles still cost energy, but do not produce any results**. All servers, regardless of make, model and even architecture carry out these idle cycles, and are able to report CPU utilization and thus idle time, making idle cycles the most universally identifiable workload of all ICT workloads.

Servers can, by virtue of power management features, reduce the electrical power draw of the CPUs that are partially loaded by adjusting P-states, or completely idle by adjusting C-states (paragraph 2.1). The use of these power management features is highly recommended but inhibits the use of simple CPU load as an indicator of effective use of resources. Instead of solely targeting CPU load, it makes more sense to

[5] While this is true in general, single purpose data centres exist and in such cases a unit of work can be defined. Similarly, when assessing a specific workload, the efficiency of that workload can be expressed as unit of work per unit of energy.

target energy use of the server, specifically that part of the energy that does not produce any results, namely the energy use when idle.[6]

It is thus logical to split the total energy use of a server into two parts (equation 3):

$$\boldsymbol{Server\ energy\ (total) = Server\ energy\ (Idle) + Server\ energy\ (active)}$$ **Eq 3: Total server energy**

It is possible to determine the energy used for running these idle cycles from easily obtainable data, and express this as a percentage of the total energy use of the server (equation 4):

$$\boldsymbol{SIC = \frac{Server\ Energy\ (Idle)}{Server\ Energy\ (total)}}\ .\ 100\%$$ **Eq 4: Server Idle Coefficient**

The SIC is, like the PUE, a measure of effectiveness, expressed as a percentage. The ideal value would be zero, meaning that all energy is used for active cycles, the worst value is 100% meaning that servers systems are using energy without any (scheduled) result being produced.

3 THE DATA CENTRE IDLE COEFFICIENT (DCIC)

As the name suggests, the server idle coefficient is determined for a single piece of equipment, an individual physical server. The server being monitored can also be “virtualized”, running multiple virtual servers, but all data collected concerns physical components, physical CPU load, and power draw as determined by the base operating system and the power supplies of that server.

The complete ICT hardware architecture inside a data centre is comprised of servers, storage and networking equipment. The ICT impact study (VHK and Viegand Maagoe, 2020) concludes that servers were, and are expected to remain the most energy demanding components of the ICT architecture. In the year 2020, servers used up 81% of the energy used by the entire ICT infrastructure (see Table 1),

It is possible to extend the server idle coefficient into a KPI that covers the entire data centre; by means of a simple addition of the idle energy over total energy use of all servers in a given data centre. This resulting KPI is the Data Centre Idle Coefficient - DCIC (equation 5):

$$\boldsymbol{DCIC = \frac{\sum Server\ Energy\ (Idle)}{\sum Server\ Energy\ (total)}}\ .\ 100\%$$ **Eq 5: Data Centre Idle Coefficient**

The formulation of the DCIC ignores the energy use of storage and networking equipment. However, since the servers’ energy usage contribution is known (81% of the total ICT energy use), the DCIC can be used to estimate the data centre energy savings potential, achievable through the optimization of servers´ usage, as defined in equation 6:

$$\boldsymbol{Energy\ Savings\ Potential = server\ energy\ fraction\ .DCIC\ .\ E(DC)}$$ **Eq 6: Energy Savings Potential**

Where the *server energy fraction* is the fraction of the total ICT energy used by servers.

The DCIC defined in equation 5 would add all contributions of different users in a colocation data centre. There is however a possible other use of the DCIC that employs the exact same calculation. Instead of adding the contributions of all servers in a data centre, the equation can be used to add all contributions of servers from a single owner (user) or single application.
This methodology can thus be used to determine the ineffectiveness of a distributed IT architecture. Since the same calculation method is used, distinction between the use in an entire data centre and use for evaluation of a distributed architecture could be indicated by use of a superscript, $DCIC^{distributed}$ where the superscript can be adjusted to identify what particular environment has been evaluated.

[6] A simple comparison can be made with the fuel use of an automobile, when standing still, a running engine still consumes fuel. Total fuel use is the fuel consumption when moving plus the fuel consumption when stationary.

3.1 DATA AND CALCULATION METHOD FOR THE DCIC

The server idle coefficient is calculated using data available from the server hardware platform. Two parameters need to be obtained from each server at a regular interval and recorded for further analysis. The data to be recorded is:

- Total power draw [Watt];
- CPU utilization [%].

The interval of measurement shall be between 1 minute and 1 hour. The it equipment owner shall decide the interval based upon server operating conditions such as provisioning cycle, speed of change of the server load, change of server load in a day, and characteristics of the application. Similarly, the length of the measurement period over which the data is required should be determined based upon workload characteristics. The minimum suggested length is one week, so that day/night and workday/weekend patterns are included. For continuous measurements, a rolling time window of the same length can be used.

The total power draw, in Watt, can be obtained from the systems management console, and is available to the system administrator of the server.

The CPU utilization is collected either from the master (host) operating system, or from monitoring software. The CPU utilization is expressed as a percentage of available CPU capacity.
Table 2 shows an example of such server hardware data, where the interval was chosen at 15 minutes.

Time stamp	CPU %	Power [W]
28/05/2020 12:16	24,16	364
28/05/2020 12:31	28,2	359
28/05/2020 12:46	53,57	408
28/05/2020 13:01	24,54	351
28/05/2020 13:16	24,43	356
28/05/2020 13:31	28,85	372
28/05/2020 13:46	35,7	377
28/05/2020 14:01	45,36	392
28/05/2020 14:16	29,22	367

Table 2 : Example output of server hardware data with a 15 min. interval [anonymized data from leap pilot (D.H.Harryvan, 2021)]

As shown in equation 3, the SIC is determined by dividing the energy usage when in idle, by the total energy use of a server. The energy value is obtained from the measurements of power (P, in Watt) by applying the basic formula from physics (equation 7):

$$\boldsymbol{Energy} = \int \boldsymbol{P(t)dt}$$ **Eq 7: relation between power and energy**

Given that the data is collected in time intervals (t), the assumption is that the power and CPU utilization represent **an average value over these time intervals**. With this approximation, the integral can be replaced by a summation. Calculating both, idle and total energy, thus involves a summation over the number of recorded intervals, as shown in equation 8

$$\boldsymbol{Server\ Total\ Energy} = \sum_1^N \boldsymbol{P(n).t(n)}$$ **Eq 8: Server Total Energy**

Where
- n is the time interval number;
- P(n) the recorded power for interval (n);
- t(n) the duration of the interval (n).

Determining the idle energy involves both the CPU idle **time** as well as the **power draw** of the server when idle. Most monitoring software solutions record/provide CPU load, but not idle. Therefore, the CPU idle energy needs to be calculated with a difference to a 100% load (100%- CPU load%); which results in the following formula (equation 9):

$$\boldsymbol{Server\ Idle\ Energy} = \sum_1^N [\mathbf{100\%} - \boldsymbol{CPU\%(n)}].\boldsymbol{P_{idle}}.\boldsymbol{t(n)}$$ **Eq 9: Server Idle Energy**

Where
- n is the interval number;
- CPU%(n) the recorded CPU load (%) for interval n;
- P_{idle} is the idle power of the server;
- t(n) is the duration of the interval (n).

As can be seen in equation 9, there is one parameter that still needs to be determined and plays an important role in this calculations, namely P_{idle}, the power draw of a server when it is in an idle state.

3.2 DETERMINING SERVER IDLE POWER

Determining the server power draw when the server is in Idle mode (P_{idle}) is essential for determining the SIC (see equation 9) but the determination of the P_{idle} is not trivial.
The ideal situation is to have a fully installed server, including the virtualization layers, OS and applications installed, but without any user programs running.
The power draw is recorded with the system turned on, but without any programs running, yielding P_{idle}.

This ideal situation to determine the P_{idle} in active servers is impossible to reach, because running equipment cannot be isolated, and the user programs cannot be stopped to carry out a measurement of idle power.

A series of other options exist for determining the server idle power draw:

1) When a server has a **static power setting**, the active and idle power are identical. In this case, the calculation for determining the server idle coefficient is simplified, and the SIC equals the average CPU idle percentage;
2) When a server has a **dynamic power setting** and shows a period in which CPU utilization is below 1%, the average power draw over this period can be considered a reasonable approximation to P_{idle};
3) When a server has a dynamic power setting and is never completely idle, the linear extrapolation of the power vs CPU utilization curve towards 0% utilization will yield an acceptable value for P_{idle}.

Each of these options has been used for the analysis of the Dutch pilot results (D.H.Harryvan, 2021).

3.3 INCORPORATING IDLE COEFFICIENT WITH EXISTING KPIS

When assessing performance of a data centre, currently, the most quoted KPI is still the PUE (ISO, 2016). While ISO 30134 includes a variety of KPIs , these other KPIs are not commonly used.

Since the SIC and thereby the DCIC as well as the PUE are energy focused KPIs, it is not hard to define a new metric that incorporates the available KPIs and data. There is no formal name for such a metric but an idleness corrected PUE can be formulated as (equation 10):

$$\boldsymbol{ICPUE} = \frac{\boldsymbol{E(DC)}}{\boldsymbol{E(IT)}}.\mathbf{1}/(\mathbf{100\%} - \boldsymbol{DCIC})$$ **Eq 10: ICPUE**

where
- *E(*DC) is the total data centre energy consumption (annual) in kWh;
- *E(*IT) is the IT equipment energy consumption (annual) in kWh;
- ICPUE is the power usage effectiveness corrected for idleness (annual);
- DCIC is the data centre idle coefficient.

The result of the correction is a number that is higher than that what is currently customary in PUE reporting. If, for example, the current PUE of the facility would be 1,5 and idleness accounts for 50% of the energy use, the ICPUE would yield 3,0.

An ICPUE can now be improved on two fronts, by improving the efficiency of the facility and by lowering the energy used for idle cycles. The last can be achieved through the use of power management, but more effectively by consolidating workload on a limited number of servers, and turning off unused infrastructure. By increasing the utilization of the infrastructure, idle energy is decreased. The ideal value of such an ICPUE would be 1, and worst case would still be infinity.

The ISO 30134 defines two KPIs, not often used in practice, which have some overlap with the SIC, these are ITEEsv (ISO/IEC) and ITEUsv (ISO/IEC).

The IT Equipment Energy Efficiency for servers (ITEEsv) according to,ISO 30134-4 is a KPI which describes the maximum performance per kW of all servers or a group of servers in the data centre, based upon a specification or potential performance of these servers. ITEEsv reflects the energy efficiency **capability** of servers, not the energy efficiency at a real operating situation of the servers.
True efficiency is not addressed by the ITEEsv, but ITEEsv accounts for capability and is used to quantify the effects of introducing servers which have high capability per unit energy. Data centres with larger ITEEsv values indicate, on average, installation of servers with higher energy efficiency.

Determination of the ITEEsv uses two more parameters that characterize a server capability:

- SMPE(i) is the maximum performance of a server i; and
- SMPO(i) is the maximum power consumption of a server i in kW.

$$\boldsymbol{ITEEsv} = \frac{\sum \boldsymbol{SMPE(i)}}{\sum \boldsymbol{SMPO(i)}}$$ **Eq 11: ITEEsv**

The performance (SMPE) of a server can be expressed in any desirable quantity, and as such the unit for ITEEsv depends on this number. The ISO 30134-4 states requirements for the benchmark used:

The benchmarks used to calculate ITEEsv shall have:
- a SMPO that is collected using a precise and highly reproducible power and performance measurement methodology;
- a run to run variation of effectiveness score of ≤ 5%;
- workloads with a high correlation to server power consumption;
- a benchmark that measures and reports power and performance during execution of included workloads;
- a benchmark that is a generally accepted tool or its results are used for the class of server being tested.

The relationship between the ITEEsv and the SIC is minimal, like the SIC, an ITEEsv concerns servers only, but the ITEEsv is a one-time measurement based on a benchmark score and does not reflect the operational conditions when the server is in use.

The IT Equipment Utilization for servers (ITEUsv) is a logical extension of the ISO 30134-4 in the sense that the actual utilization of the servers is added into the reporting of this KPI.
ITEUsv accounts for utilization aspects and, according to its description, is used to quantify the impact of one or both of the following:
- improving utilization ratio of servers by using such technologies as virtualization and server consolidation for sharing use of servers;
- reducing the number of servers to achieve the same level of information processing.

$$\boldsymbol{ITEUsv(t)} = \frac{\sum \boldsymbol{CUSi(t)}}{\boldsymbol{N}}$$ **Eq 12: ITEUsv(t)**

Where:
- CUSi(t) is the CPU utilization ratio of server i at time t (%);
- ITEUsv(t) is the average CPU utilization of all servers or a group of servers in a data centre at time t;

- N is the number of servers in a data centre or in a group running at time t.

$$\boldsymbol{ITEUsv} = \frac{\boldsymbol{1}}{\boldsymbol{a}}\sum[\boldsymbol{ITEUsv}(\boldsymbol{t0} + \boldsymbol{e}.\boldsymbol{i})]$$ **Eq 13: ITEUsv**

Where the summation is over the interval number i ranging from 1 to a and

- a is the number of ITEUsv(t) measurements intervals over a year (all intervals should be the same length);
- t0 is the starting time of measurement;
- e is the interval of measurement, where e × a = one year.

ITEUsv depends on a periodic registration of CPU utilization, identical to what is needed to calculate the SIC and DCIC. The important difference being that the current ITEUsv simply reports the average CPU utilization while the SIC and DCIC are energy centric measurement KPIs where through the energy, the DCIC scales the contribution of larger and smaller servers to appropriate contributions.

4 EXAMPLES OF IDLE KPIS FROM PILOT PROJECTS

As part of "LEAP", a project commissioned by the Amsterdam Economic board, a number of servers running production workloads recorded power draw and CPU utilization over a period of a week.
The complete report is available from the website of the Dutch RVO website (D.H.Harryvan, 2021). Actual data collected during this pilot was used for the examples shown in this chapter to illustrate the calculation method, the usefulness of the metric as an overall indicator as well as the usefulness of the raw data in determining a course of action for optimization.

4.1 THE MUNICIPALITY OF AMSTERDAM VIRTUAL DESKTOP SERVER

One of the participants in LEAP was the municipality of Amsterdam, the server shown in this example is a typical example of a well-designed virtual desktop server.

Server description:
- ProLiant BL460c Gen8, power management setting high performance, allowing CPU P-states;
- Intel(R) Xeon(R) CPU E5-2680 v2 @ 2,80GHz, 512 GB RAM;
- Application: Multiple Citrix servers on top of VMware ESXi, 6.5.0.

A small sample of a week worth of measurements is shown in the table below:

Server ID	Time stamp	Power (Watt)	CPU load (%)
citrix.amsterdam	27/01/2020 05:37	148	1,8
citrix.amsterdam	27/01/2020 05:52	147	1,92
citrix.amsterdam	27/01/2020 06:07	151	3,84
citrix.amsterdam	27/01/2020 06:22	171	3,38
citrix.amsterdam	27/01/2020 06:37	155	2,77
citrix.amsterdam	27/01/2020 06:52	179	8,43
citrix.amsterdam	27/01/2020 07:07	181	16,68
citrix.amsterdam	27/01/2020 07:22	211	28,06
citrix.amsterdam	27/01/2020 07:52	226	36,12
citrix.amsterdam	27/01/2020 08:07	202	22,88
citrix.amsterdam	27/01/2020 08:22	218	32,13
citrix.amsterdam	27/01/2020 08:37	235	43,08
citrix.amsterdam	27/01/2020 08:52	232	34,25
citrix.amsterdam	27/01/2020 09:07	236	39,48
citrix.amsterdam	27/01/2020 09:22	251	49,36
citrix.amsterdam	27/01/2020 09:37	265	46,41
citrix.amsterdam	27/01/2020 09:52	253	50,92
citrix.amsterdam	27/01/2020 10:06	266	65,73
citrix.amsterdam	27/01/2020 10:22	275	63,93
citrix.amsterdam	27/01/2020 10:37	269	61,39

Table 3 : Excerpt of actual monitoring data for a virtual desktop server.

In order to calculate the SIC, the first order of business is the calculation of the P_{idle}, the power draw of this particular server when idle. As stated in paragraph 2.5, several methods can be used, the choice for the method depending on the available data.
To evaluate the data, a graph is made, plotting the power draw (Watt) as function of the CPU load.

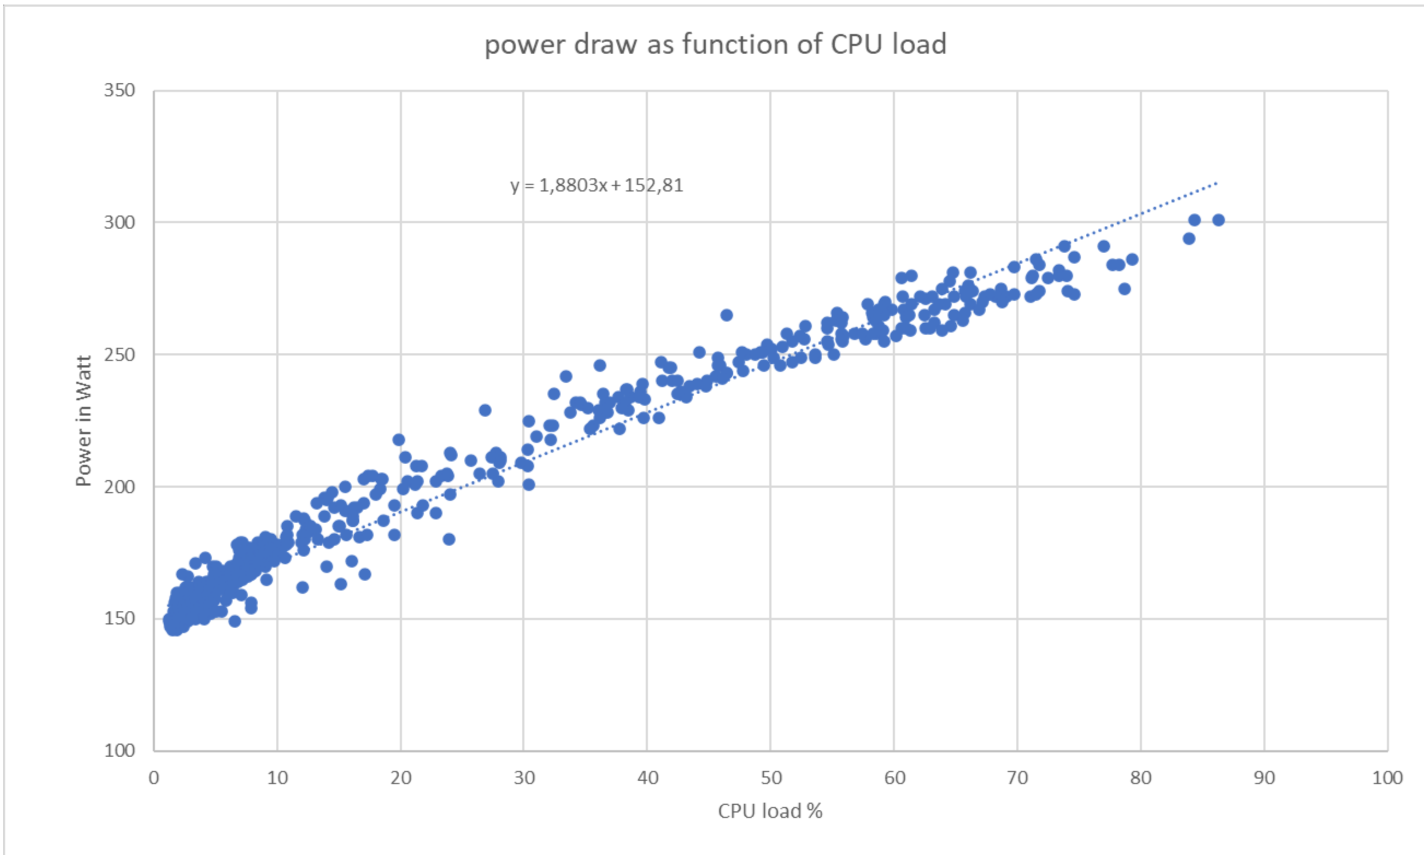

Figure 2 : Server power draw as function of CPU load (Data collected during the LEAP pilot at the Municipality Amsterdam)

The graph in Figure 2 contains much information, but for the determination of the P_{idle}, in this case the power draw at very low utilization can be used, yielding a P_{idle} of approximately 150 Watt. Alternatively, a linear approximation can be used to obtain a value for a P_{idle}. Although in this particular case there seems to be enough data for higher order or logarithmic fitting of the data, in most cases the range of different CPU loads is not as wide and a linear regression is most appropriate. The linear regression of this data yields a P_{idle} of 152,8 W which underlines that the determination of P_{idle} is not always exact, an error of 2% is a reasonable margin.

From the entire table, the equations 8 and 9 allow to calculate the idle energy use and the total energy for the week of measurement :

Idle energy:20,2 kWh
Total energy: 31,9 kWh
SIC: 63,3%
Average CPU idle: 80,1%

The first important conclusion is that almost 2/3 of the energy used by this server is spent on a process without result, 20 kWh are wasted each week.
A secondary observation is that the server power management delivers some energy savings. Over 80% of CPU cycles are idle, but the energy use in idle is dampened by the use of ACPI P states. This observation underlines the usefulness of the SIC which focusses on the energy use during idle cycles, not on the cycle count directly. Depending on the server technology and power management settings, the energy lost in idleness can be considerably lower than the simple observation of the average CPU load might suggest. To further explore the options for limiting the idle energy, a different graph, as shown in Figure 3 can be of help:

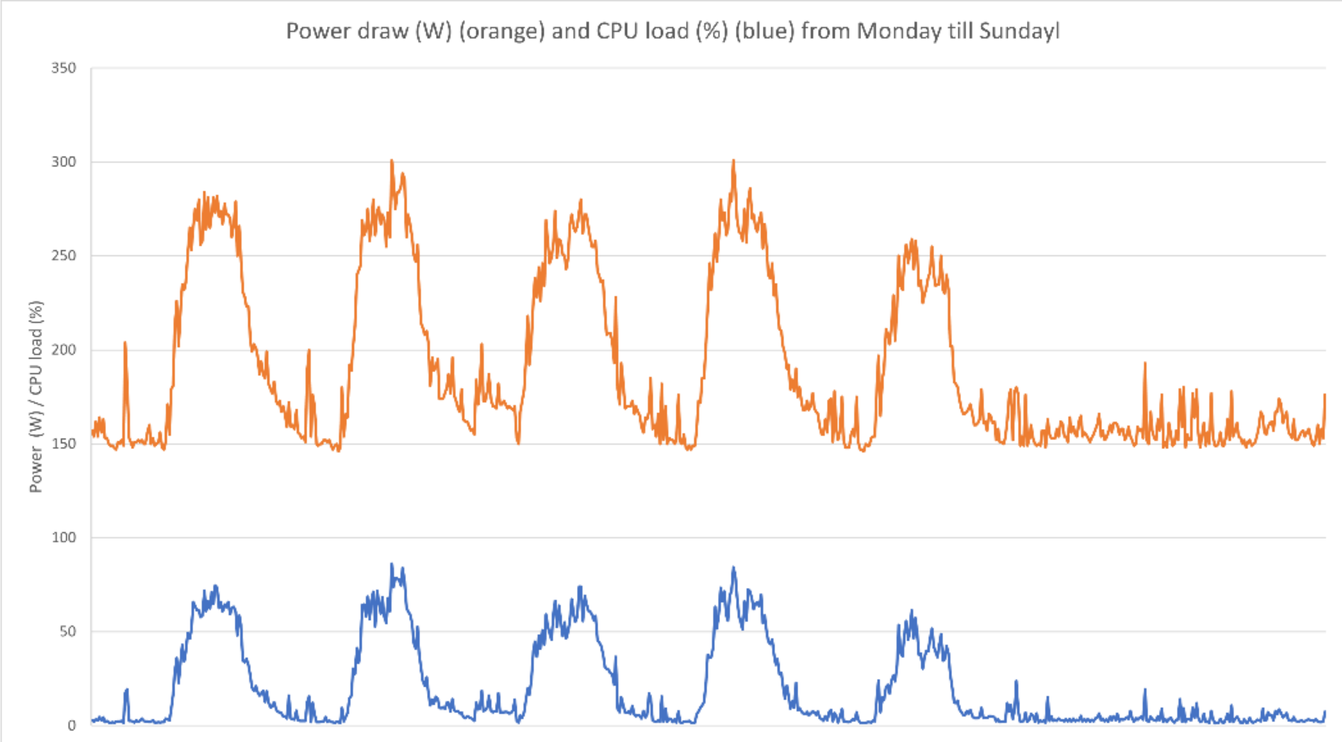

Figure 3 : CPU load (%) and Power draw (W) as function of time (Data collected during the LEAP pilot at the Municipality Amsterdam)

Figure 3 shows the variation in demand on the server over an entire week and the associated power draw. A number of observations can be derived from the graphs, the first being that the system is properly sized. During office hours, recognizable by the peaks in CPU load (Monday through Friday) CPU load reaches 80% and higher. Downsizing the system to reduce energy use is therefore not an option. It is also clear from the graph that the highest energy loss is during the time that the system is not used at all. In the evening and during the weekend there are no users, but power draw is still 150 Watt. It has been discussed with system management that it is a viable option to turn this system off at the end of business (e.g., on Friday) and turn it on at the start of business (e.g., Monday morning).
Such an action would save 9 kWh (30% of total energy) and lower the SIC to below 50%.
Energy savings can also be obtained from simply allowing CPU C-states by changing power management to a 'balanced' setting. In that case, P_{idle} will be lowered, which results in lower energy use during all hours of limited demand (see chapter 4.3).

4.2 EXAMPLE FOR SHARED SERVICES

Measurements performed in a financial institution during the same LEAP pilot program illustrate the effect of both, resource sharing as well as the effect of static power management settings. The load profile of the servers shown in Figure 4 and Figure 5 is very different from the example in Section 4.1, exhibiting a much more evenly distributed character. The server hardware is very similar to the hardware described in the previous example, but as will become clear later on, the power management setting in this case was set at 'static high performance'. This setting does not allow changes in either CPU P or C states.

Server description:
- ProLiant BL460c Gen8, power management setting static high performance;
- Intel(R) Xeon(R) CPU E5-2680 0 @ 2,70GHz;
- Application: VMware server.

Figure 4 and 5 below show the performance for two servers that function as cluster nodes in a large VMware cluster:

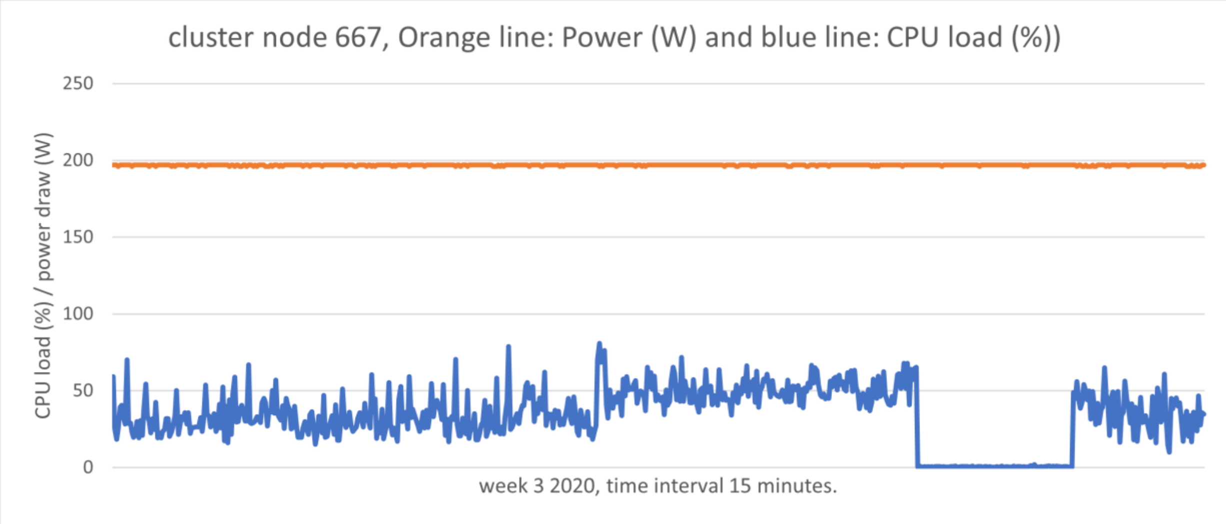

Figure 4 : CPU load (%) and Power draw (W) as function of time (server 667).

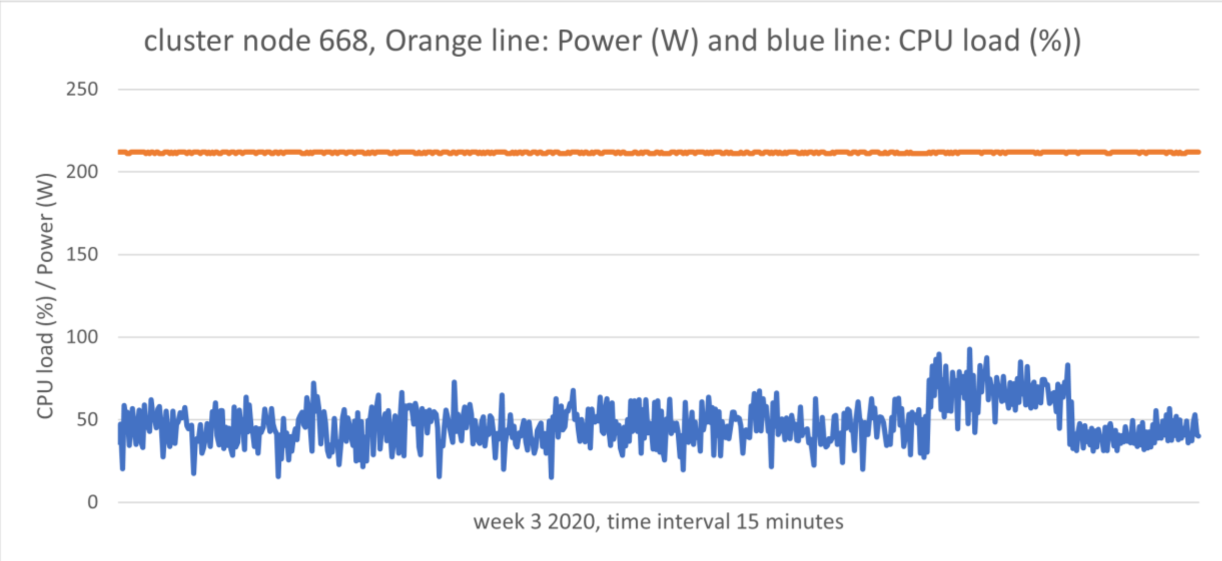

Figure 5 : CPU load (%) and Power draw (W) as function of time (server 668).

As is visible in these graphs, the power draw of both servers is completely independent of the CPU load. The effect of these static settings is most clear during a maintenance period at the end of the week. All workload from node 667 is removed and spread over the remaining cluster nodes. CPU load on 667 is near zero, the CPU load on 668 is raised by almost 20% to an average near 70%. None of these actions have any impact on the power draw of the servers.
Because power draw is static, taking P_{idle} in the calculation of the SIC to be identical to P_{total}, this simplifies the calculations so that the SIC is equal to the averaged CPU idle%.

For these servers, this results in:

Node 667: SIC = 66,7%
Node 668: SIC = 52,5%

The same measurements show that in theory, it is possible to turn off cluster node 667 completely, the redistribution of workload would result in a SIC of approximately 30%.

4.3 EXAMPLE ON THE INFLUENCE OF POWER MANAGEMENT

The LEAP pilot, the source of all data shown in this chapter, was actually conducted to demonstrate the use of power management in a server as a means of saving energy. During the pilot a limited number of participants changed the power management setting on the running server. The data generated allows the evaluation of a possible secondary use for the raw data for evaluation of a single server model.

The VMware cluster being studied was built with on a HP blade enclosure containing:

- HPE BL460 Gen9;
- 2 x Intel(R) Xeon(R) CPU E5-2697A v4 @ 2,60GHz.

The following 2 graphs shown in Figure 6 and Figure 7, were created on the basis of the data, which in this test had an interval time of 5 minutes. The two graphs show the results for a single server, without changes in workload. During the week of measuring, the power management setting of the server under study was changed from high performance to balanced mode.

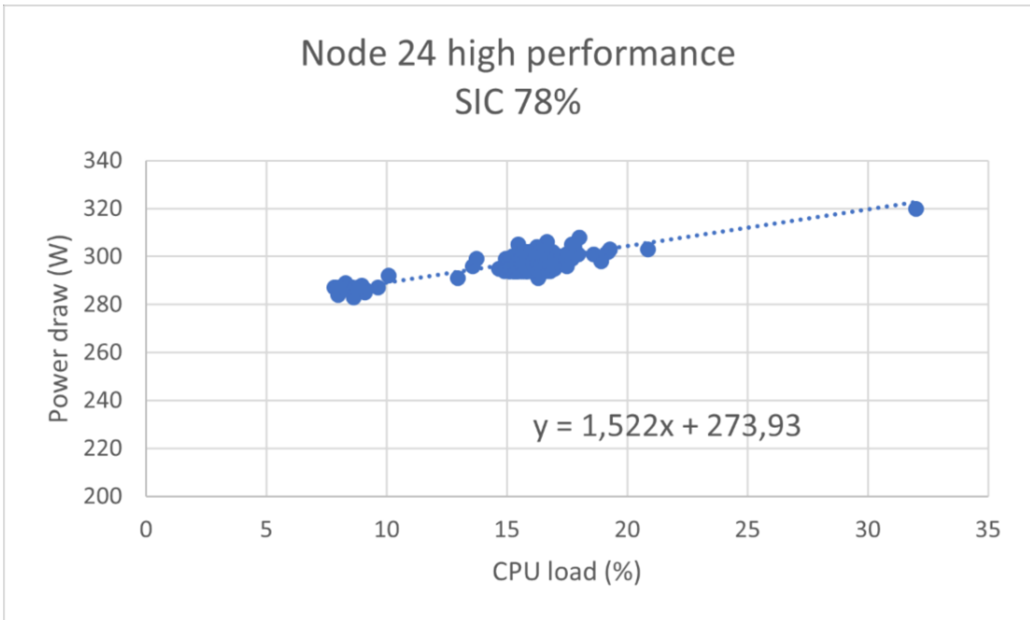

Figure 6 : Server power draw as function of CPU load (P-states only).

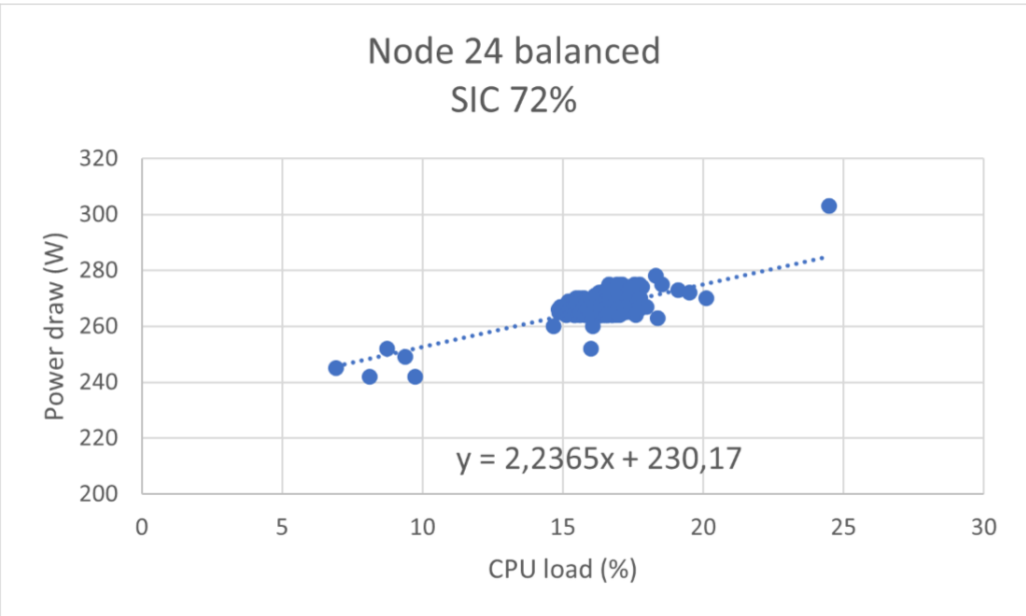

Figure 7 : Server power draw as function of CPU load (P-states and C-states).

The linear regression lines with their mathematical function are shown in each graph.

The following conclusions can be drawn from the raw data and these graphs.

1) Changes in power management setting take effect immediately without requiring a reboot (as is described in hardware and OS documentation);
2) The linear regression line of high-performance mode, allowing adjustment of CPU P-states only, is less steep than that of the balanced setting, which also allows CPU C-state changes;
3) Allowing both P and C-states results in a lower P_{idle} than when only allowing P-state changes (as is expected);
4) Taking into account that there were no changes in application load when the power management setting was changed, the change in setting results in direct energy savings, as the average total power dropped by 11%;
5) The SIC is shown to be sensitive to changes in power management and can be improved by allowing both, P and C-states.

The regression lines offer a possible novel insight for comparing different server models. When the SIC would evolve into a commonly used KPI, the data collected might be used to identify servers with deviations or unexpected behaviour. Deviant behaviour can indicate a wrong setup or the need for maintenance. As such, the data collected benefits the system owner not only by giving insight into resource allocation and usage but also as an indicator for maintenance purposes.

5 USABILITY OF THE IDLE COEFFICIENT

Although examples in chapter 4 show SICs that are reasonably close to 50%, the same LEAP pilot revealed that a large number of servers have much lower utilization and thus much higher server idle coefficients. The LEAP pilot was carried out using servers that were selected by the participating companies and not using a random selection, still average CPU idle of 90% and higher were not uncommon, which suggests that there is a large energy savings potential. Actual resource usage data for ICT infrastructure at global, country or even data centre level are either unknown or unpublished, and this lack of reliable data is an important issue.

In order to get an indication of the future use of a new KPI such as the DCIC, it is useful to look at the history of a well-known KPI that is in common use in the data centre industry today, the PUE.
Taking the development of the PUE as an example; In the first period of the KPI usage, adoption was low and there was a reluctance to disclose the value for many data centres. In those first reports, very high PUE values were not uncommon, and a PUE of 2 was considered to be very reasonable. Over time, broader use of this KPI, e.g., in marketing statements and minimum requirements in government tenders has resulted in a marked improvement of reported PUE values. More importantly, the improvement in PUE is indicative of a more efficient control of ambient conditions within the data centres, and is therefore an indicator for very substantial energy savings. Off course, one cannot attribute all of the innovations in cooling technologies to the existence of a metric such as PUE, however the adoption of these innovations by data centres have most certainly been accelerated by a universally recognized method for measuring and reporting a KPI, in this case the PUE.

Where the PUE was targeted at the energy use of the facility infrastructure of a data centre (power and cooling), the SIC and DCIC are targeted at the energy use of the IT equipment. There are very compelling reasons to target **energy use of the ICT infrastructure itself**, mostly in light of the continuous growth in demand for ICT services and partly because of the success of the PUE.

Worldwide figures obtained by the uptime institute show a stalling of PUE improvements (Uptime institute, 2020), information that is confirmed by the observations in the ICT impact study from VHK (VHK and Viegand Maagoe, 2020):

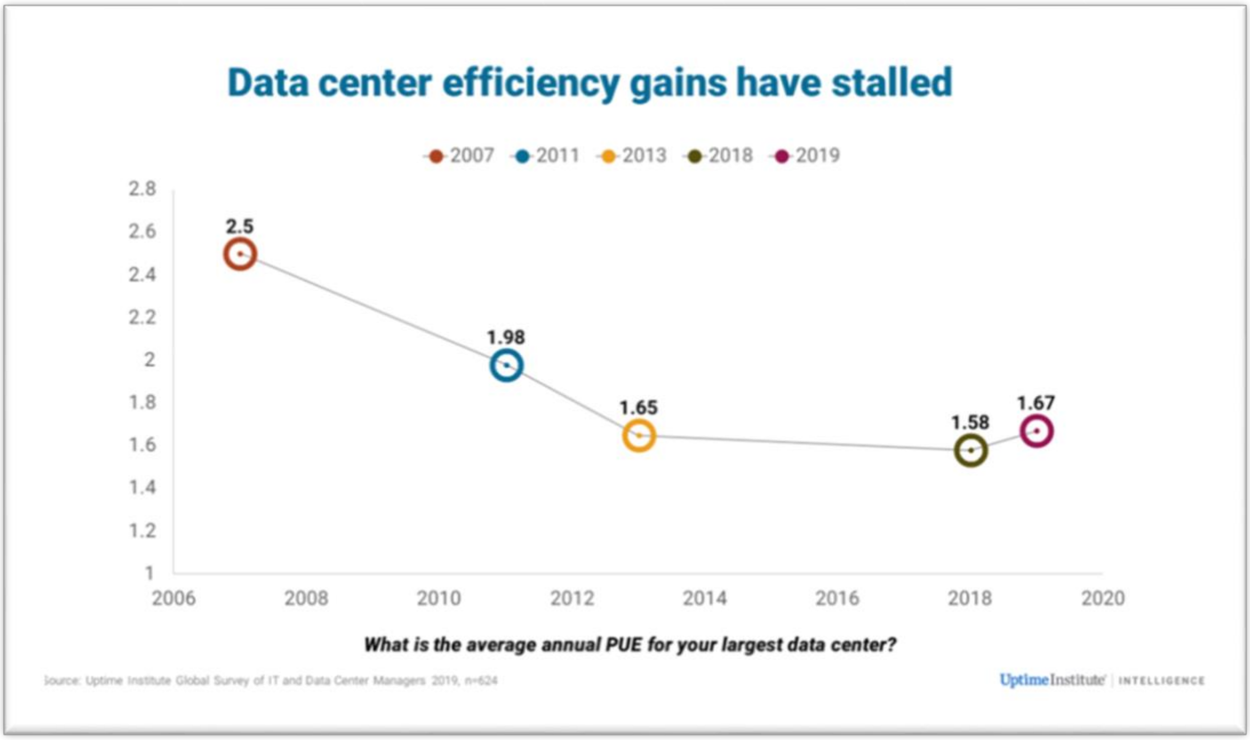

Figure 8 : Trend in worldwide PUE numbers ((Uptime institute, 2020)).

Over the past decades, the ICT industry has benefited from the continuous development of faster and importantly, more energy efficient chips. The computational power per kWh of electrical energy has increased exponentially with a factor of 2 each 1,6 years, generally recognised as 'Koomey's law'. But also, this trend is slowing down (Koomey, 2016).

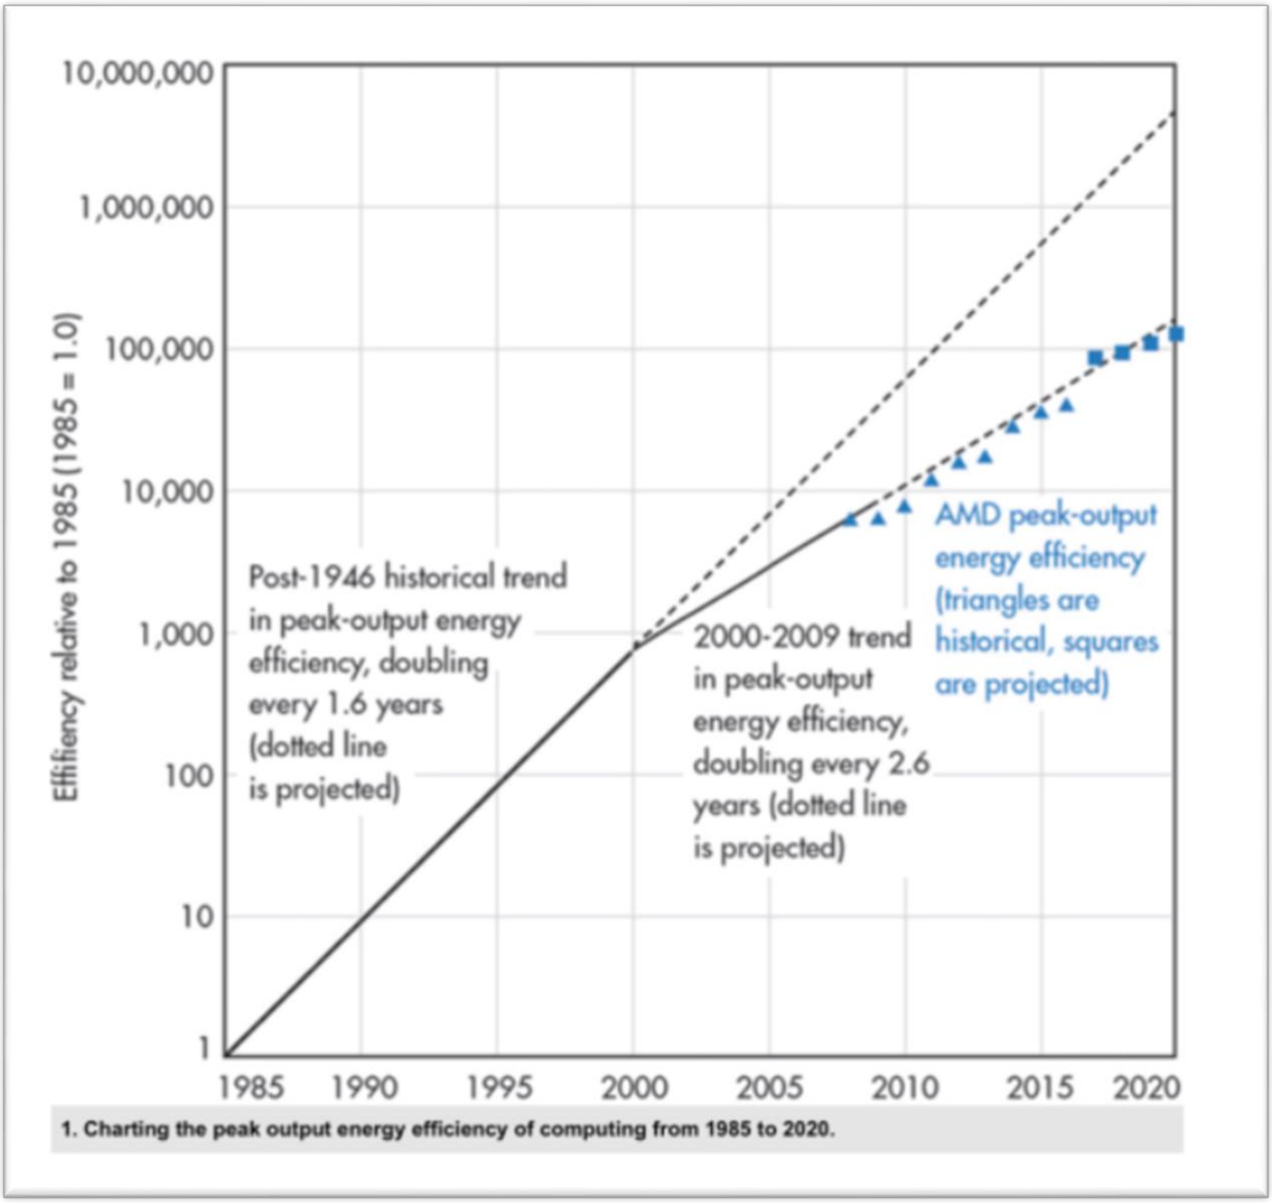

Figure 9 The maximum computing efficiency of commercially available servers; (Koomey, 2016)

Figure 9 shows the efficiency of CPU's expressed as the number of computations per unit of energy at peak performance, relative to the efficiency attained in 1985. The historical trend in this "peak-output energy efficiency" showed that this energy efficiency used to double every 1,6 years, but the trend is slowing down. With the slowdown in energy efficiency improvements in new generations of servers, absorbing the ICT demand growth by replacing older servers with newer ones will be become less effective as an energy saving method. All of this could result in a substantial growth in DC energy demand.

An increase in utilization of existing/installed equipment can help to curb this growth in energy demand and as with the PUE, the SIC and DCIC KPIs that clearly show the fraction of the energy that is wasted in idle, could accelerate the adoption of existing technology and spark innovations that are available to reduce this wasted energy.

Existing technology includes power management and virtualization utilizing virtual machines, containers and serverless computing. But innovations are needed such as software defined data centres where the resources can be matched to the current workload, as well as and economic incentives such as price differentiation to shift workload from peak to off hours, in order to get a more even workload profile.

The focus on the energy used for idle instead of energy for active use might seem semantic, but in decades of discussions the debate on what constitutes activity in a data centre has not yielded a clear definition. The definition of what does not constitute work is simpler and traverses not only different Data Centre infrastructure topologies but also different processor architectures.

Work still remains to be done in order to create an actual standard, but this work is similar to, and can benefit from, the work done in ISO 30134-5, the IT Equipment Utilization for servers (ITEUsv). A clear benefit of the metrics SIC and DCIC above the use of the ITEUsv is that there is a clear link with energy, and as shown with the measurements in paragraph 4.3, the SIC is sensitive[7] to changes in power management settings.
Another aspect of an energy focus is that there is the option of extending the metric to include the energy use of networking and storage equipment. These contributions are smaller than that of the servers, but still significant. The definition of the idle coefficient need not be changed, it would remain E_{idle}/E_{total}, where idle energy for storage and networking equipment would have to be defined.

5.1 CHALLENGES FOR ADOPTION

There are several challenges for successfully implementing the of the idle coefficients as accepted KPIs. The most important aspect is that the data needed for the calculation, while easy to obtain and record, can only be obtained by those with correct permissions on the server that is being monitored. This means that in a colocation data centre, the tenants must all cooperate and provide data to a central location which can, based on this data, calculate the KPI for the data centre and all individual servers.
In current contracts, there are no provisions for recording such a continuous stream of data, and many data centre operators prefer to keep responsibility for the IT contained in the data centre separated from the responsibility for the actual data centre and the facility equipment needed for a secured electrical supply and climate control.
For policy makers and policy enforcement, this source for data can also pose problems. If the authorities would want to monitor the idle coefficient for a data centre, it will complicate matters if a list of tenants using the data centre needs to be obtained, and each of them would need to be contacted to request the needed data and information.
Another consequence of the distributed data sources will be the difficulty in validation of the data. It is relatively easy to introduce errors in the SIC calculation, for instance by underestimating the P_{idle}, but such errors can be eliminated once the data collection and interpretation is automated.
Headway is being made on both the data collection; the LEAP pilot attracted the interest of the Dutch environmental authorities. Environmental law in the Netherlands obliges large energy consumers to take energy saving measures and obliges reporting on both energy use and the implementation status for the mandatory measures.
For data centres in the Amsterdam region, additional requirements from the municipality places limits on the PUE of a data centre. Since reporting structures are already in place in the Netherlands, and the law states that the data centre operator can be held accountable for the behaviour of its tenants, the authorities are exploring the inclusion of the Idle coefficients in the obligatory reporting.
The way the Dutch authority is currently trying to address this situation is based upon Voluntary Agreements. While the big challenge is that the primary data source is in the hands of the ICT owner, this also constitutes an opportunity for a data centre owner to intensify the relationship with its tenants. As discussed in chapter 4, the data does not only yield the idle coefficient, but it also provides insights on capacity utilization and optimization and could lead to cost reductions. The DC operator can, based upon experience with many of its clients, function as an advisor on top of providing the services that operator already offers.

[7] Different from the ITEUsv that claims it is influenced by server power management but doesn't the SIC is actually shown to correctly reflect a change in a power management setting. The SIC is improved when switching from high performance to balanced.

For large companies and shared services providers that are also data centre owners, it is more straight forward to calculate the KPIs. The obstacle remains that as the Idle coefficient is new, openness about use by companies remains unclear (i.e., companies might not want these numbers to be made public).
Again, taking the Dutch example, energy and PUE reporting by a data centre to the appropriate authorities is mandatory, but these authorities cannot disclose this information before it is anonymized. In this way insight is gained in the total impact of the data centres on the energy infrastructure and on the possibilities for energy saving measures, without any image damaging information being disclosed publicly.

The last challenge is a technical one, data integrity and validity must be guaranteed, especially as it might be obtained from multiple sources. The logical method would be through monitoring agents placed upon the servers, but issues around security and privacy must be addressed before large scale applications collecting data from all servers within a data centre can be used.
These issues are challenging and important but can be solved. For the LEAP pilot for instance, a script collecting and storing the required data was created for VMware users by VMware; the data was stored locally and forwarded to the LEAP team later for processing. The support of major parties such as the hardware manufacturers and operating system providers (e.g., Microsoft, VMware and others) will be needed, and with the increased use of a KPI comes the development of tools and software to determine it.

5.2 NEXT STEPS

As mentioned in the previous section, the current lack of information on the actual energy use of servers and data centres, and the split in what is an effective use of ICT and what not, is a serious issue. With respect to development of energy efficiency policies, it seems that a first step might be to require reporting from large energy users. Legislation that requires reporting on energy use and on the adoption of energy saving measures exists in the Netherlands and has provided valuable insight into areas for energy savings. **For data centres** such reporting could include total annual energy use and an annualised PUE report such as is already required for participants in the EU code of conduct for data centres (EU-JRC, 2016). Such a report can then be extended with the DCIC calculation.

When such data is compiled, analysis of raw data will suggest courses of action to reduce the total energy spent on idle cycles, both by minimizing the fraction of CPU cycles that are idle as well as by reducing the P_{idle} . This reduction in the energy spend on idle can diminish the total energy used by data centres or diminish the rate of growth of this energy use. Preferably, the organizations analysing the energy data would supply the data centre and/or ICT equipment owners with advice on measures that result in a more effective use of resources and leave it to the industry to adopt these measures on a voluntary basis.

The European Code of Conduct [EU CoC] for Data Centres (12) is an example of such a scheme where participants voluntarily commit to the adoption of the measures described in the best practices guide. The best practices guide of the EU CoC already contains practices that allow determination and optimization of the SIC, specifically:

- Enable power management features: formally change the deployment process to include the enabling of power management features on IT hardware as it is deployed. This includes BIOS, operating system and driver settings.
- Energy and temperature reporting hardware: select equipment with power and inlet temperature reporting capabilities, preferably reporting energy used as a counter, in addition to power as a gauge.
- Control of equipment energy use: select equipment which provides mechanisms to allow the external control of its energy use. An example of this would be the ability to externally restrict a server's maximum energy use or trigger the shutdown of components, entire systems, or sub-systems. Consider the use of user defined policies.
- IT Equipment utilization: set minimum or average targets for the utilisation of IT equipment (servers, networking, and storage).
- Control of system energy use: Consider resource management systems capable of analysing and optimising where, when and how IT workloads are executed and their consequent energy use.

When adoption of these best practices is below a desired level, policies that are more restrictive can be necessary to push adoption. One example of such local policy is the permit system such as it is used in the Amsterdam region. Data centres will need to show a (design) PUE of 1.2 or better in order to obtain and retain a permit. How such a PUE can be obtained is not prescribed, a number of solutions exist, it is up to the data centre operator to choose the solutions that fit this purpose best.

Setting the PUE level that must be attained is the result of many years of monitoring, combined with the input of experts. A similar approach can be envisioned for the use of the idle metrics. First monitoring, data analysis and expert input, resulting in maximum idle levels that can and should be attained by the industry.

In the case of the idle metrics, setting a maximum DCIC goal might not be realistic, in which case it is feasible to specifically mandate the implementation of best practices. Dutch environmental law uses such an approach, which is industry specific. The data centre industry in the Netherlands is subject to specific mandatory measures. The list of such measures (RVO, 2021) is published on the website of the Netherlands Enterprise Agency and describes measures that have a proven return on investment of less than 5 years. The use of power management features for example is listed as a mandatory action specific to improving the SIC.

Aside from direct actions, compiling data on the use of resources by the data centre industry is as important as the sharing of this data. Anonymised data can be useful both, for identifying best practices and for the proliferation of these best practices at a faster pace than is currently the norm.

Government organisations can play an important role as a trusted data hub where all of this data is collected and analysed without direct commercial goals, to be used for meeting National environmental goals as well as the strategic goals of the ICT industry, so crucial to our economy.

6 REFERENCES

D.H.Harryvan. (2021). *analysis LEAP track 1 'power manageent'.* Retrieved from https://www.rvo.nl/sites/default/files/2021/01/Rapport%20LEAP_EN_TG.pdf

EU-JRC. (2016, 6 13). *Code of Conduct Data centres*. Retrieved from Data Centres Energy Efficiency: http://iet.jrc.ec.europa.eu/energyefficiency/ict-codes-conduct/data-centres-energy-efficiency

ISO. (2016). *ISO 30134-2*. Retrieved 2021, from https://www.iso.org/standard/63451.html

ISO/IEC. (n.d.). *ISO/IEC 30134-4:2017.*

ISO/IEC. (n.d.). *ISO/IEC 30134-5:2017.*

Koomey, J. (2016). *blog post*. Retrieved from https://www.koomey.com/post/153838038643

NEN. (2017). *NEN-EN 50600-4-2.*

RVO. (2021). *Informatieplicht energiebesparing uitgevoerde erkende maatregelen Bedrijfstak 06 - Commerciële datacenters.* Retrieved from https://www.rvo.nl/sites/default/files/2020/04/erkende-maatregelenlijst-commerciele-datacenters-april-2020.pdf

Schöne, R. (2015). Computer Science - Research and Development volume 30, pages 219–227 (2015). (30).

Standard Performance Evaluation Corporation. (2021, 8 28). *SPECpower_ssj2008 Results*. Retrieved from SPECpower_ssj2008 Results: https://www.spec.org/power_ssj2008/results/

UEFI. (2019). *Advanced Configuration and Power Interface (ACPI) Specification 6.3.* Retrieved from https://uefi.org/sites/default/files/resources/ACPI_6_3_final_Jan30.pdf

Uptime institute. (2020). *UI Intelligence report 34 : Beyond PUE.* uptime institute.

VHK and Viegand Maagoe. (2020). ICT impact Study.