Enabling the information society

Data centre Fixed to Variable Energy Ratio metric DC-FVER

An alternative to useful work metrics which focuses operators on eliminating fixed energy consumption

Liam Newcombe	BCS Data Centre Specialist Group
Zahl Limbuwala	BCS Data Centre Specialist Group
Paul Latham	BCS Data Centre Specialist Group
Victor Smith	BCS Data Centre Specialist Group

Contents

1. Summary

In just a few years the data centre industry has embraced PUE as the measure of efficiency for the mechanical and electrical infrastructure of the data centre. The process of measuring and reporting has provided a focus and comparable measure of performance which has allowed many operators to make substantial improvements. However, there is still no generally accepted metric for IT or software efficiency and most energy efficiency measurements stop at the IT power cord. FVER addresses these issues by including the IT equipment and software in the data centre efficiency measurement whilst allowing each operator to select an output productivity measure which usefully represents their data centre and business activity.

1.1 PUE

PUE has pervaded the market to the extent that virtually all vendors now market how their product or service will improve your PUE and many operators compete on who can measure and publish the lowest operating PUE. The success of PUE has however, been such that most operators and designers are now facing rapidly diminishing returns as they have already identified the majority of the economically viable savings available from their mechanical and electrical infrastructure.

This focus on PUE coupled with the absence of any generally accepted IT or software efficiency measures leaves industry extracting ever smaller savings from the mechanical and electrical infrastructure at the expense of many other, potentially greater savings elsewhere in the data centre. Further, operators who improve IT energy consumption face the issue of rising PUE which they must then explain as it incorrectly suggests a reduction in overall efficiency.

1.2 IT metrics

Whilst there have been a number of proposals for metrics to measure IT work or useful work from the data centre none of these has successfully dealt with the underlying problem that different data centres carry out different tasks and there is no single, comparable and fair measure of the “IT work” or “Useful work” from a data centre. The quest for a single catch-all metric of IT efficiency is much like the quest for the perfect meal, entirely subjective to the person making the assessment, neither comparable nor transferable to anyone else.

The significantly shorter lifetimes of the IT equipment and software compared to that of the mechanical and electrical infrastructure also present issues. The optimisation methods and economic balances we have developed and used for the mechanical and electrical equipment with a service life of up to 20 years do not easily translate to software and IT equipment which we may retain for only a couple of years. This difference in lifetime has grown in the last few years as the rate of IT equipment refresh has accelerated with the widespread adoption of virtualisation and commodity computing technologies.

It is now normal for an operator to include efficiency in the selection of new IT equipment, whether in terms of system power requirements and efficiency for their workload or the consolidation ratio they will achieve against their current generation of equipment. There are also well established metrics such as SPECPower[i] and labelling such as Energy Star[ii] which purchasers already use to guide their selection of IT equipment. This means that there is little IT equipment, particularly servers, which is substantially less energy efficient than the market benchmark. This competitive market and high refresh rate leave little value to be added by a metric which reports the “IT work efficiency” of a data centre as this is a constantly moving target and such a metric would inevitably be more a measure of the operators refresh rate which is constrained by many issues other than IT energy efficiency.

1.3 FVER

FVER, the Fixed to Variable Energy Ratio metric provides the required next step in data centre metrics by including the IT equipment and software as well as the data centre infrastructure. FVER targets operating waste in the software, IT and M&E infrastructure and provides a complementary reporting metric to PUE.

The substantial opportunity in most current data centres is not finding additional improvement in the power efficiency when delivering peak business output but how much power they consume to do nothing. Even with our improvements in PUE, the power draw of most data centres is still almost constant despite large fluctuations in the work being delivered. If a car used fuel at the same rate whether driving four passengers on the highway or idling stationary this would be seen as completely unacceptable, however this is how most current data centres behave.

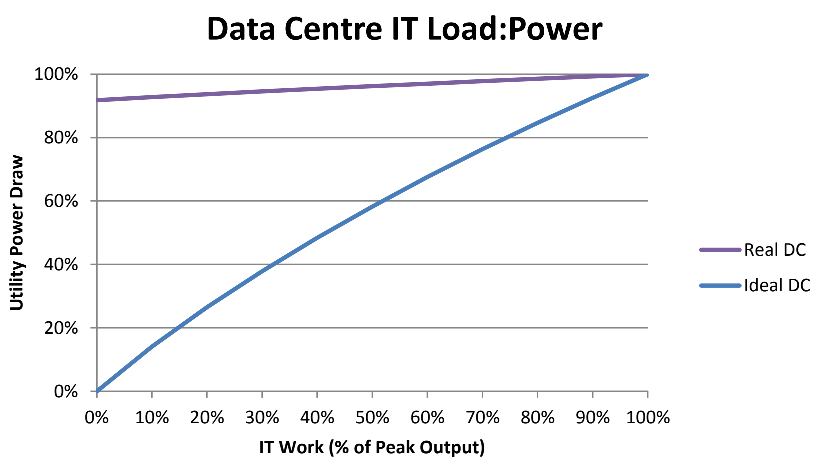

Figure 1 Real vs. Ideal data centre power draw against IT output

The Data Centre Fixed to Variable Energy Ratio metric (DC FVER) measures for the first time what proportion of data centre energy consumption is variable, i.e. related to the useful work delivered, versus what proportion is fixed allowing operators to understand how much of their energy cost is related to the work delivered and how much is a fixed burden to be eliminated. The FVER is calculating using the formula below;

$$FVER = 1 + \frac{Fixed\ Energy}{Variable\ Energy}$$

Data Centre	FVER	Comment
Real Data Centre	10	> 90% of energy consumption is fixed and not related to workload, < 10% is related to work delivered
Ideal Data Centre	~1	~0% of energy consumption is fixed, ~100% related to work delivered

FVER does not prescribe how each data centre operator should measure the output of their data centre as this measurement must vary operator to operator, the same measure will not be effective for a supercomputer, a trading site and a web video operator. FVER instead allows each operator to select an appropriate set of proxy measures for useful work which are meaningful to their business activity. These are then normalised and combined with existing PUE measurement points of IT energy and total Utility energy to allow for a comparable measure of their data centre performance.

Off peak efficiency is of particular value to enterprise data centre operators whose peak IT work period is a relatively small part of the total time. In these cases reductions in off peak energy demand provide a direct and obvious value. FVER also carries value for data centres with less variable workloads though as the optimisation of the site for part load performance through eliminating fixed energy waste is also likely to reduce the energy consumption at high output.

FVER may be measured for the entire data centre or for part of the data centre allowing service providers and customers to measure and report the metric without needing full control of the entire data centre.

2. Introduction

Despite much effort by many organisations around the world we are still no closer to a "useful work" metric for data centres than we were when PUE was adopted. This is because the underlying problem has not been addressed, nobody has managed to devise a universal, portable, time stable and fair way to reduce the many dimensions of IT work down to one number which can be sensibly compared with energy consumed and therefore measure data centre efficiency. Critically no method of measuring the useful work which can be compared between data centre operators has been defined. Given that any metric produced will be compared by operators this inability to provide useful comparison has become the major obstruction to the uptake of any metric above PUE in the data centre stack.

2.1 Fixed to Variable Energy Ratio

This paper suggests that by changing the way we look at IT output metrics we can leverage much of the effort already put into metric development and in doing so we can not only mitigate the weaknesses and perverse incentives of these previous proposals but in most cases turn these weaknesses into strengths. The DC FVER metric leverages both existing measurement protocols developed for DPPE and PUE as well as the effort expended defining and testing the Green Grid productivity proxies (DCeP).

This alternative thereby provides a high value diagnostic metric for the data centre which presents a very strong indication of direction of improvement, fewer perverse incentives and provides some ability for operators to compare their "score" which has been the major part of the success of PUE.

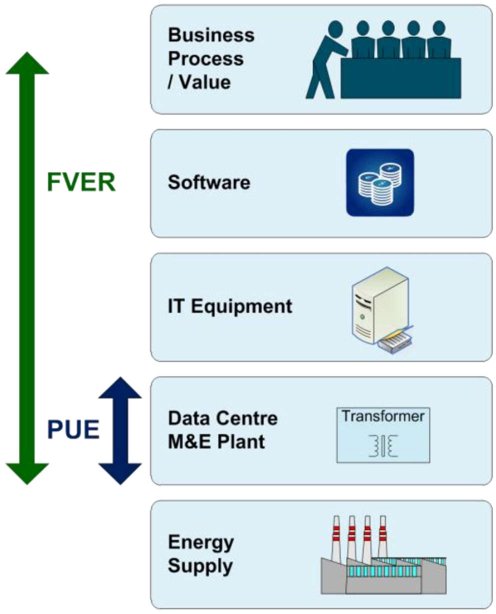

Figure 2 Data centre stack

DC FVER is a normalised metric which provides a way for operators to understand and compare how well their energy consumption is linked to the "useful work" or business value delivered from their data centre, specifically how much of their energy consumption is fixed and how much is variable with load.

2.2 Why does the UPS and utility load not change?

In the last few years there have been significant improvements in both industry awareness of data centre energy efficiency and the skills required to improve that efficiency. This has led to substantial changes in the way data centres are built and operated. The PUE metric is now considered to be a KPI for most data centres. Programmes such as the US Energy Star for Data Centres and the EU Code of Conduct use PUE or the measurement data to track the improvement delivered by the industry. This should indicate that much of the fixed load in the data centre infrastructure has been minimised.

In recent years we have seen two substantial changes in servers, both have come about through the growing awareness of power element of the total cost of ownership and how this relates to the capital cost of the servers. Firstly, servers have been optimised to reduce total energy consumption for a given level of performance. Secondly, servers have also been optimised to reduce their energy consumption when operating at less than full work-load; programmes such as Energy Star have been particularly successful in this area. Given the IT refresh rate there should now be substantial populations of servers with strongly variable power demand installed in many data centres.

Despite these changes in both infrastructure and IT equipment behaviour the UPS output load and utility power draw of data centres still remains remarkably constant. This lack of power variation with workload is now arguably a greater issue than the actual efficiency when at full load. Data centres are still favoured by utility companies for being stable constant base loads when the opposite should be true for most sites, in an ideal data centre, irrespective of how much power is drawn when delivering peak "useful work" when the data centre is idle the power draw would be zero. A metric is required which allows operators to understand, target, improve and compare how well they perform in this key indicator.

2.3 Targeting the waste

A substantial part of this problem is that the clear focus on measured efficiency with PUE stops at the IT power cord. This paper argues that a major part of the potential efficiency improvement in the data centre is now locked up in this fixed "base-load" power draw and that correcting this failure should be the next priority for data centre operators and the industry. In an ideal data centre, whatever the power draw when the IT platforms are at full load it would be zero when the IT platforms are delivering no services. Clearly complete elimination of this fixed load is not possible however, what the DC FVER metric presents for the first time is a way for operators to measure how well their IT equipment and site energy consumption tracks the useful work delivered by their IT platforms.

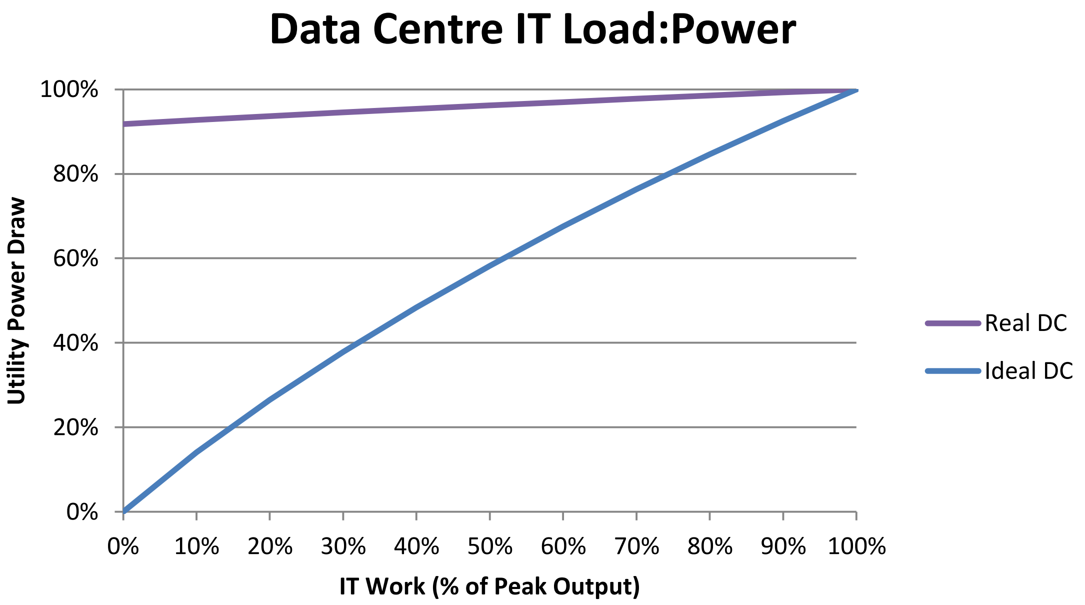

Figure 3 Real vs Ideal data centre power draw against IT output

To use the standard automotive comparison, if we own a vehicle which can transport four passengers on the freeway at 30mpg then that may be better than some and worse than others. What is more important is what happens to the rate of fuel consumption when stationary where we would expect the vehicle to use little or no fuel. Given this, it is somewhat surprising that we are happy to allow our data centres to use almost the same amount of electricity when doing nothing as when delivering full "useful work" output.

2.4 Poor fixed to variable ratio increases even peak energy consumption

Given the low equipment utilisation in most enterprise data centres we can extend this analysis to suggest that a focus on the reduction of fixed energy consumption would also reduce the energy consumption of the data centre when delivering its peak measured load. This is simply because the peak user workload on the data centre does not represent every device in the data centre reaching 100% utilisation. Whilst there may be some components or systems which do represent a choke point most systems will have substantial performance capacity available. This results in the overall utilisation of the data centre being well below 100% even when the measured peak "useful work" occurs.

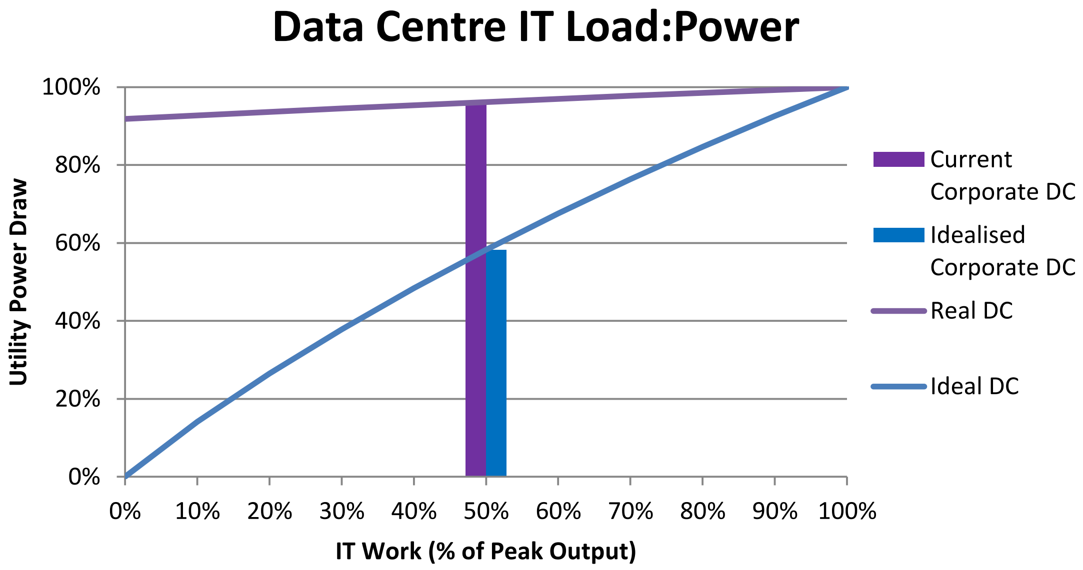

Figure 4 Real vs Idealised data centre at peak "useful work"

This is illustrated in Figure 4 where the measured peak IT work occurs at 50% of what the total installed IT equipment is capable of achieving even if some subset of devices is at 100%. Capacity planning, software and platform architecture prevent the data centre exceeding this value. As shown at this point our current corporate data centre is drawing 96% of its full peak load from the utility feed whilst the idealised version is consuming only 60% of its peak load.

Clearly the blue line represents an idealised data centre and we are unlikely to see a site with a response this steep as there will always remain a certain fixed energy consumption for devices and also the operator may choose to schedule non time-sensitive or low value jobs to run when the primary “useful work” output is low.

3. Existing metrics and current issues

This paper leverages a number of existing and proposed metrics and seeks to make use of their strengths whilst mitigating or cancelling out their key weaknesses. Before discussing the FVER metric a brief reminder is provided of the previous proposals, their strengths on which we should build and their weaknesses which we should seek to mitigate or eliminate. These issues are discussed in detail in the 2010 DCSG white paper on data centre metrics[iii].

3.1 IT work metrics

The IT work metrics proposed so far all rely on some set of weighting factors to combine the many dimensions of IT work or delivered output into a single value to be compared with energy consumption in order to measure "efficiency". Unfortunately these fudge factors mean that they cannot be fairly compared between different data centres or even with the same data centre at different times. Further, only some of these metrics measured the "useful work" from the data centre which may be related to business activity, most measure some aspect of "IT work" which cannot be easily or generally mapped to delivered business value. The "IT work" metrics therefore suffer from a substantial gap which each operator must bridge between the metrics and business justification for changes.

3.2 Productivity proxies

The Green Grid recognised there was no available method for directly measuring the "useful work" of a data centre and instead proposed a series of measures which seek to measure some other value(s) and consider these measurements to be a proxy for the actual useful or IT device work being delivered. None of these proxies fully addressed the underlying issues but they do provide a useful exploration of the practical aspects of collecting productivity data within a data centre and the relative difficulties and competencies of different methods. See the Green Grid paper Proxy Proposals for Measuring Data Center Productivity[iv]

A number of the productivity proxies suggested monitoring or sampling the activity of key applications to count how many useful business actions are performed, these proposals carry the benefit of being relatively easy to implement and record whilst also, at the application actions level being easily mapped onto the delivered business value of the IT platforms. An online auction operator for example might count number of new listings created, number of bids processed and number of item search pages served whilst a corporate user may be more interested in number of customer interactions supported by a call centre system. The use of visible and relevant business supporting actions eliminates the need to attempt to translate an arbitrary productivity measurement unit into business value. This conversion being a challenge at least as great as defining the measurement unit due to the fact that in most data centres equal units of IT work are expended to support services of differing business value.

The major issue with most of the proxy metrics is that they are either user specific or rely on subjective weighting factors or assignments of work units and therefore are not suitable as either absolute measurement or for comparison between data centres. Many of the proxies are also not time stable meaning that they are of little value for tracking performance improvement. An additional hurdle for many of the proxies was the data collection and processing overhead.

3.3 Using IT power as a proxy for productivity

A proposal published and evaluated in considerable depth by the Japanese GIPC is Data centre Performance Per Energy (DPPE)[v]. One of the major issues facing metrics seeking to capture the IT work (even if we had units to measure it) is the data collection required to log how much work each IT device does. The DPPE metric took a novel approach to this issue and instead of attempting to directly measure the work performed it used the power drawn by the IT device as a proxy for the work delivered. For each IT device the peak power draw and peak performance were established and then used the device power draw to indicate the device work. This derived value for IT work was then compared with the overall energy consumption to obtain an overall efficiency measure.

Unfortunately the use of IT power consumption as a proxy for IT productivity presents a number of issues, the two most significant being that;

a) IT device power draw is not linear with work performed – in Figure 5 at zero workload the real server can be seen to draw 150 Watts, 50% of its peak power draw, as DPPE uses the power draw as a proxy for work the sample server work would be reported as 50% of capacity whereas the work is in fact zero
b) The metric creates a strong perverse incentive, in order to increase their score or meet an imposed efficiency target an operator could disable power management on their IT devices, thereby increasing their IT energy consumption and causing an apparent increase in their work delivered whilst clearly the actual impact would be to reduce efficiency and increase energy consumption

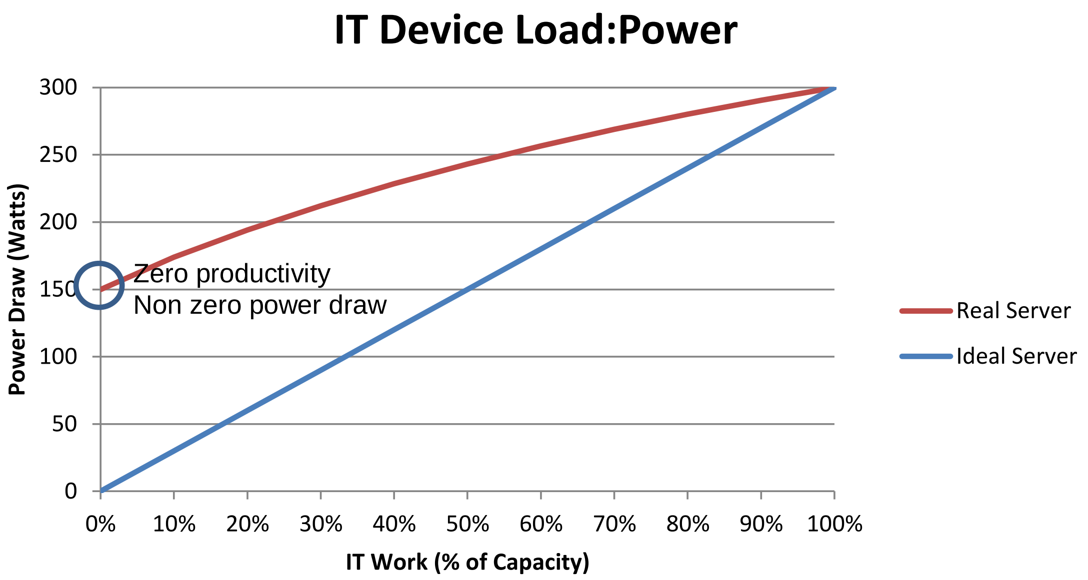

Figure 5 IT device workload to power relationship

Extending the analysis from a single server to a representative mix of server, network and storage equipment in the data centre as described in sections 7.2 and 7.3 below, we can see how significant the error is. The IT Equipment Utilisation (ITEU) for the full set of devices, at zero actual utilisation the fixed power consumption of the mixed devices results in an indicated IT equipment utilisation of over 60%.

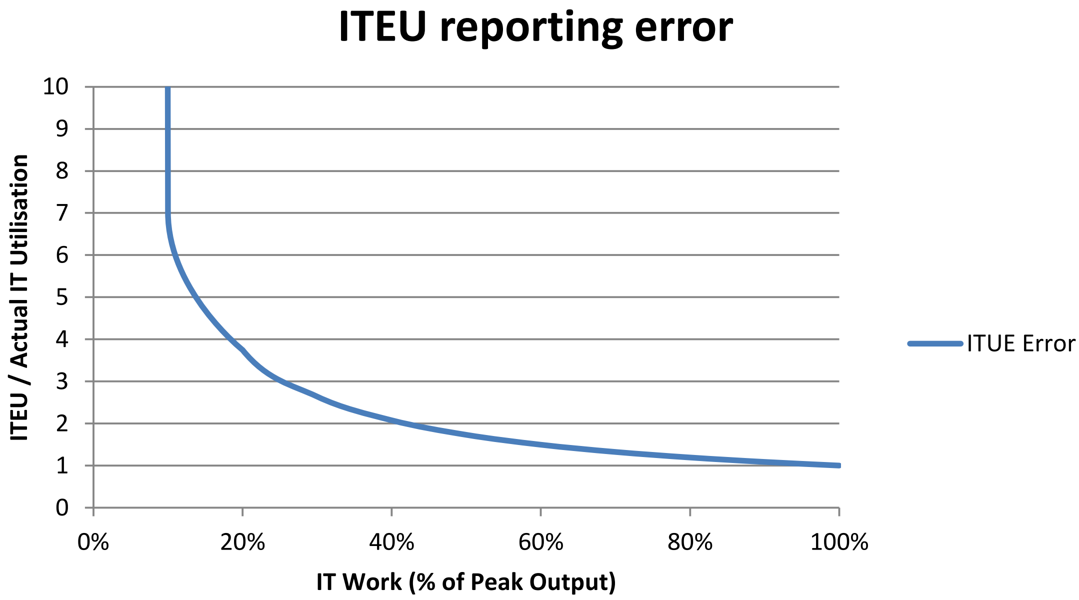

Figure 6 ITUE Reporting Error

Figure 6 shows the extent by which the measured ITEU over-reads against the actual utilisation of the data centre. This is a critical issue for a metric as, despite the recent improvements in the idle power of many devices there is no realistic prospect of devices drawing zero power at zero workload.

In summary metrics which use the IT power draw as a proxy substantially over-estimate the IT work of the data centre and this over-estimation increases as the utilisation drops, effectively biasing the metric in favour of operators with poor utilisation at the expense of those who achieve higher utilisation.

4. FVER metric

To introduce the FVER metric we will examine the behaviour of a sample data centre both before and after efficiency improvement work and explore how comparing the work output with the energy consumption can affect our overall efficiency.

4.1 Measure work output

The first step in measuring the FVER for a data centre is to select one or more productivity proxies with which you will measure the work output. These should be measures relevant to your business which represent the value the data centre delivers. The selection of a proxy measure is discussed in more detail in section 5.1.

The sample data centres represent an enterprise site with high variations between weekday and weekend work outputs, a sample weekday and a sample weekend day are used for both cases. A set of productivity measures has been selected and applied This workload is expressed as a fraction of the peak recorded work output in the measurement period, the weekday and weekend loads are shown in Figure 7 below;

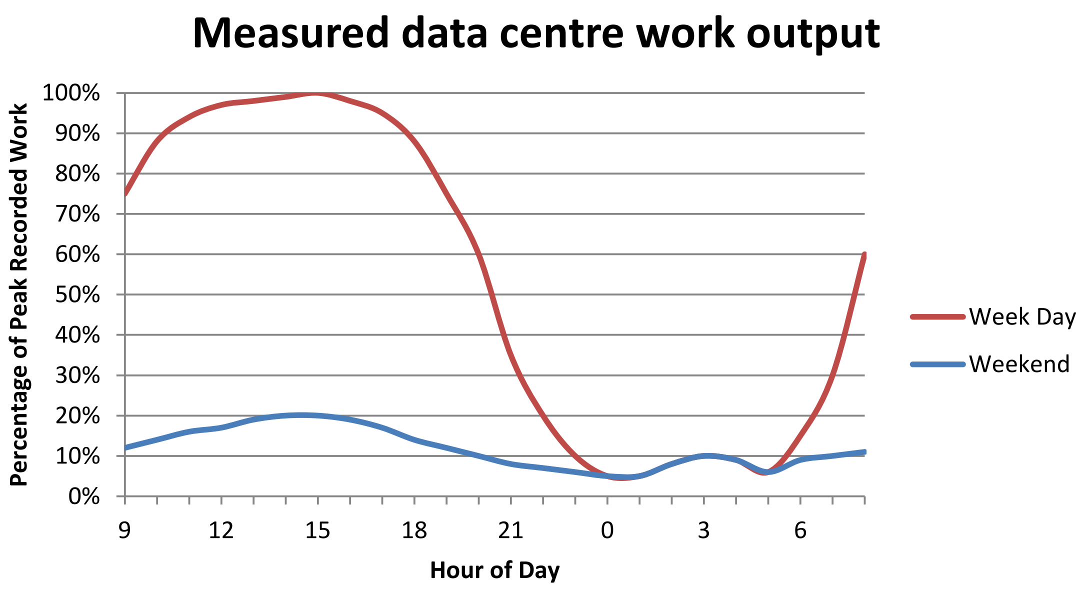

Figure 7 Measured weekday and weekend work output

4.2 Sample old data centre

The first sample data centre represents an old facility without any cooling economiser and a high fixed infrastructure overhead component (poor PUE at low IT load).

Installed IT devices

The installed IT devices consist of compute servers, storage and networking equipment, the power draw of these devices against the applied workload is shown in Figure 8 below;

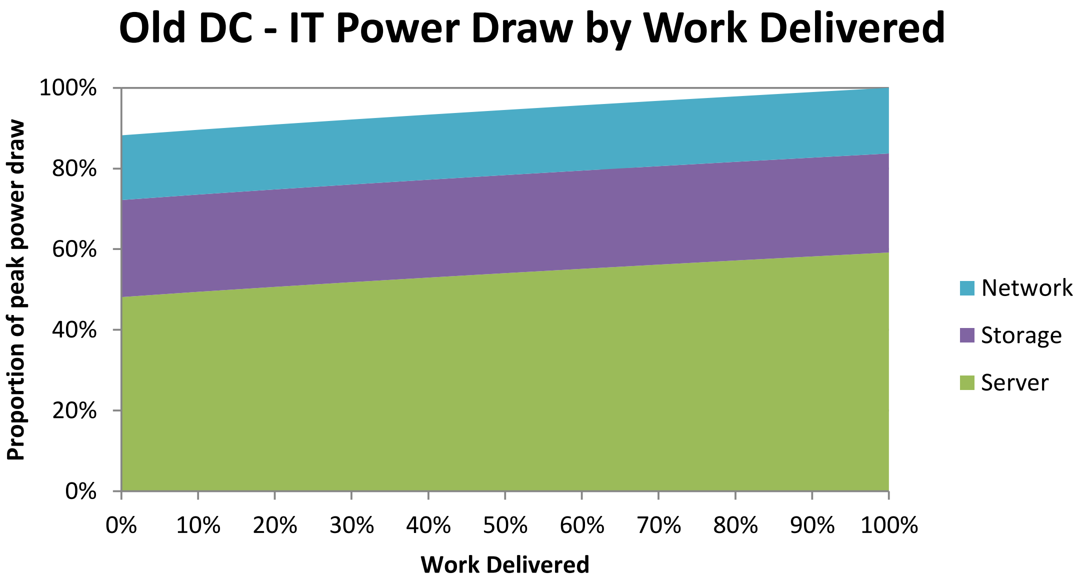

Figure 8 Old DC - IT power draw by work delivered

As shown the power draw is quite constant irrespective of the work delivered, this is due to;

- The peak server processor utilisation being ~ 15%
- The idle power of the storage equipment being > 90% of peak power
- The idle power of the network equipment being > 90% of peak power

Metered power draw

The second step is to capture power draw measurements for the data centre. In this case data for 24 hourly samples across a single weekday is shown;

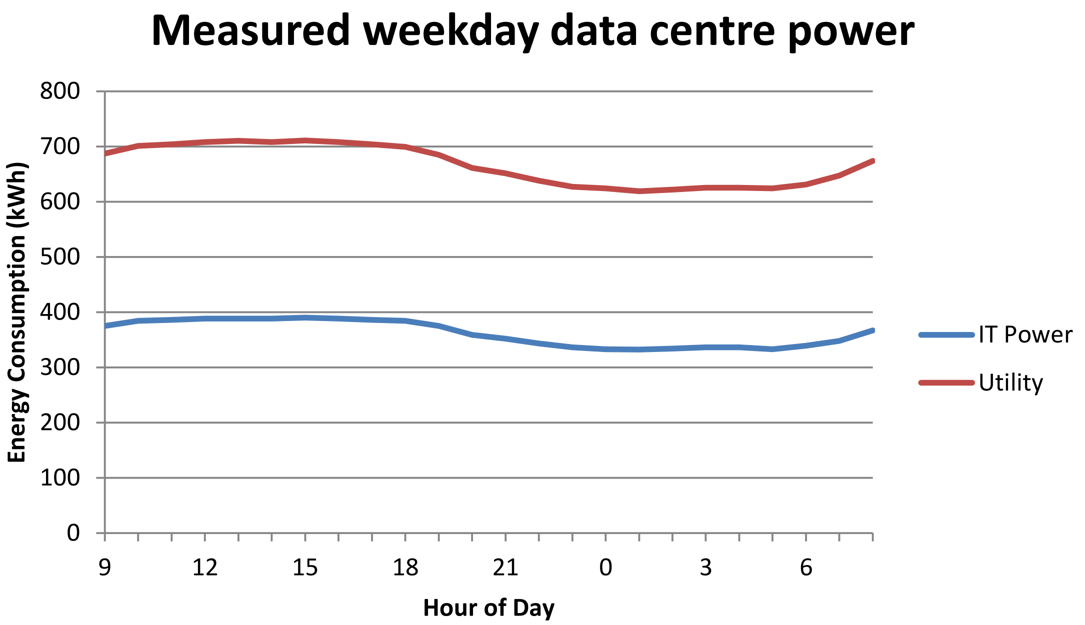

Figure 9 Old DC measured weekday IT and total utility power

As can be readily seen the IT power is relatively constant (although more variable then many real data centres) and the utility power is also quite invariant.

Comparing the work output data from the weekday to the utility power draw we see that the relationship is quite poor;

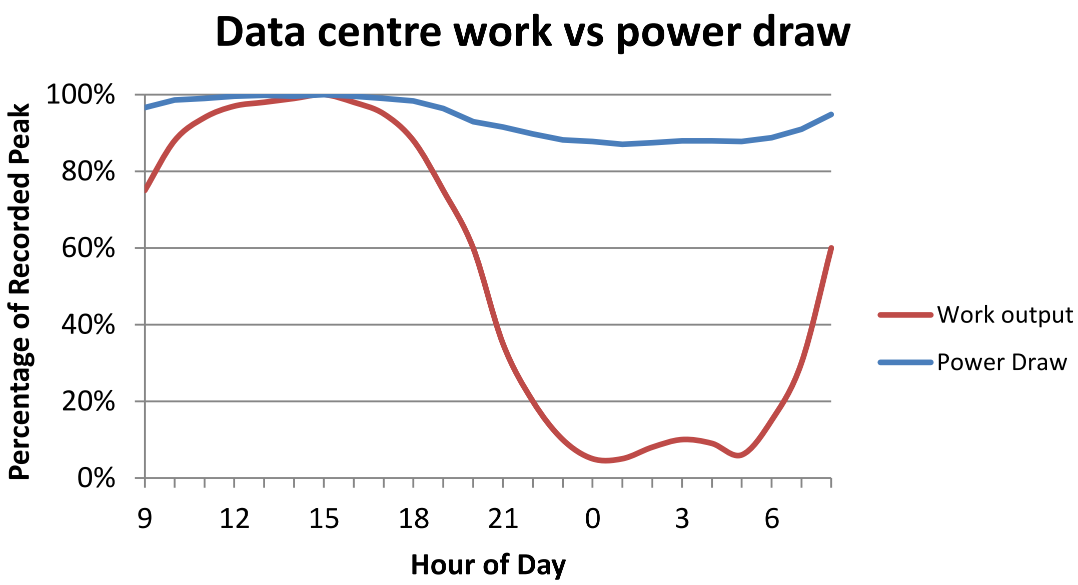

Figure 10 Old DC productivity vs. power draw

FVER analysis

Taking a couple of points from the data shown in Figure 10 above we can immediately see that the data centre is dominated by fixed energy consumption;

Element	Midday	Midnight
Work output	97%	5%
Utility power draw	100%	88%

When compared with the peak power draw from midday at midnight our data centre is still using 88% of that power to produce just 5% of the work output. It appears that more than 80% of our energy consumption is fixed.

We could calculate an initial estimate of the FVER from these two values by;

$$FVER = 1 + \frac{Fixed\ Energy}{Variable\ Energy} = 1 + \frac{0.88}{0.12} = 8.3$$

Of course, this is not a very good measure, we only used two values and we don't have a 0% work output value as there is not a time at which our data centre is doing nothing worthwhile. In order to improve the measure we will take the 24 hourly samples in a weekday and 24 hourly samples from a sample weekend day. Plotting these for comparison produces the data plot of normalised values in Figure 11 below;

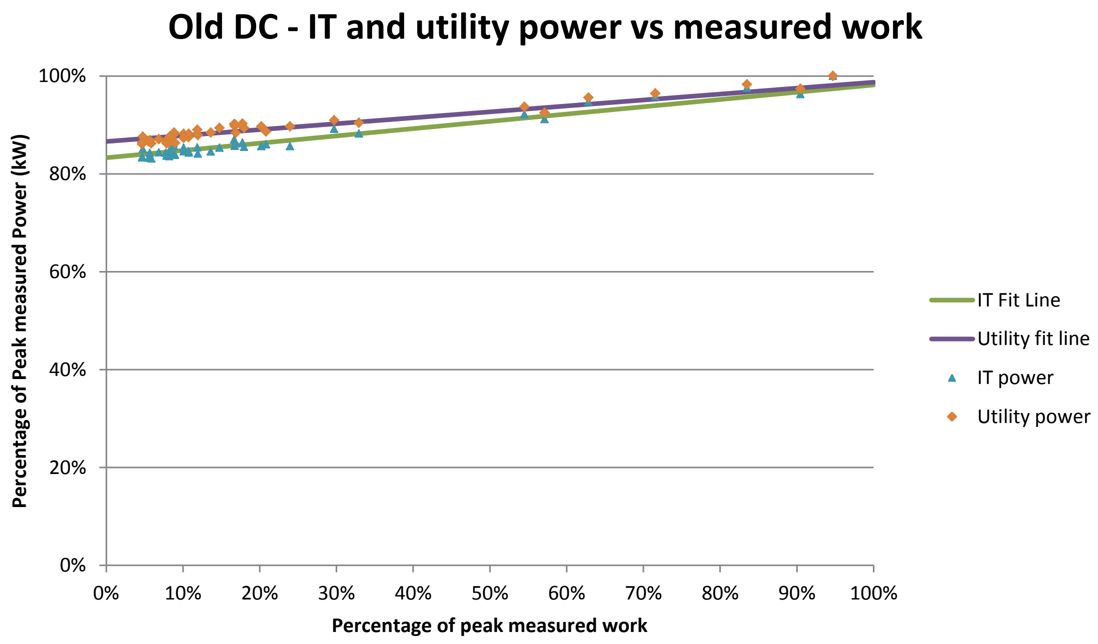

Figure 11 Old data centre FVER plot

Using a simple spread-sheet to analyse the data recorded we obtain the trend lines shown on the chart, reading where these intercept the power axis at zero measured work we obtain the values for FVER (note that an analysis spread-sheet is provided with this paper).

As we have taken both utility and IT power measurements we can calculate two values for FVER, one covering just the IT equipment and software and another covering the software, IT and data centre mechanical and electrical infrastructure.

FVER IT
First we can calculate the FVER for the IT equipment by comparing the IT equipment power draw to the measured work;

$$FVER_{IT} = 1 + \frac{0.83}{0.17} = 6.1$$

FVER Utility
Then we can calculate the FVER for the whole data centre by comparing the utility power draw to the measured work;

$$FVER_{Utility} = 1 + \frac{0.87}{0.13} = 7.6$$

For this old data centre the metrics record the following (quite poor) values for FVER at the IT and Utility;

Metric	Value
FVER (IT)	6.1
FVER (Utility)	7.6

It can easily be seen that having 6.6 (7.6 – 1) parts of energy consumption which are fixed against only one part of energy consumption which varies with work output is less than optimal.

4.3 Sample new data centre

The second example is an updated version of the old data centre which uses the same UPS, CRAC, mechanical cooling, PDUs etc. but uses air flow containment, a free cooling option for the mechanical cooling and VFD control for CRAC fans and water pumps. This site has a substantially lower fixed infrastructure overhead component and a more significant change in PUE with external temperature.

Installed IT devices

The installed IT devices consist of compute servers, storage and networking equipment, the power draw of these devices against the delivered work is shown in Figure 12 below;

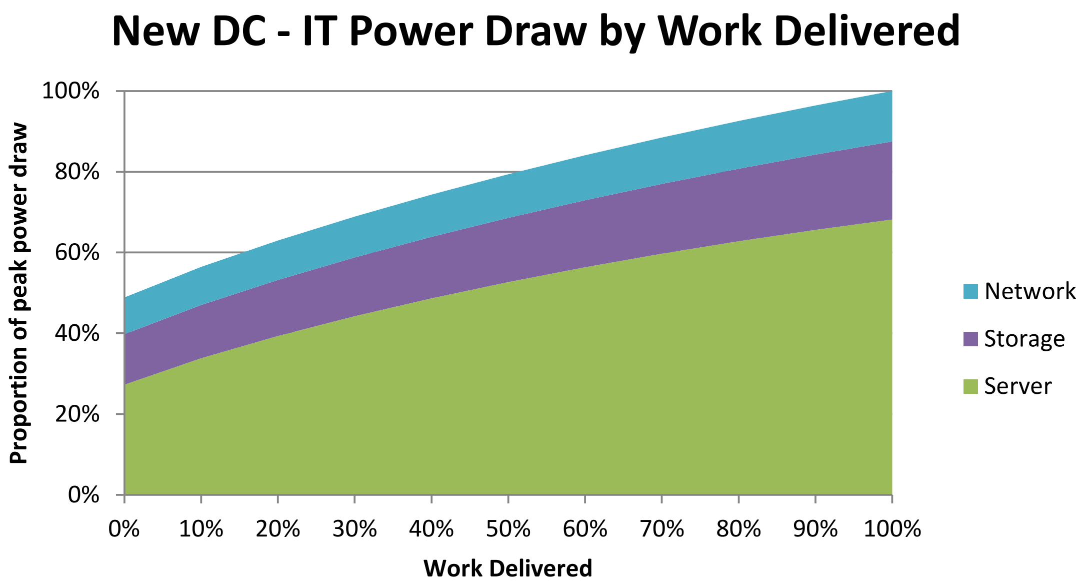

Figure 12 New DC - IT power draw by work delivered

As shown the power draw is more variable with the work delivered than in the old data centre, this is due to;

- The peak server processor utilisation being increased to 50%
- Power management functions being fully enabled on the servers
- The idle power of the storage equipment being reduced to 65% of peak power
- The idle power of the network equipment being reduced to 70% of peak power

Note that the absolute “efficiency” of the IT equipment is not important for this measurement as all power draw and delivered work are normalised.

FVER analysis

Taking 24 hourly samples in a weekday and 24 hourly samples from a sample weekend day produces and applying the linear regression analysis to the data recorded we obtain the data plot and trend lines shown in Figure 13 below;

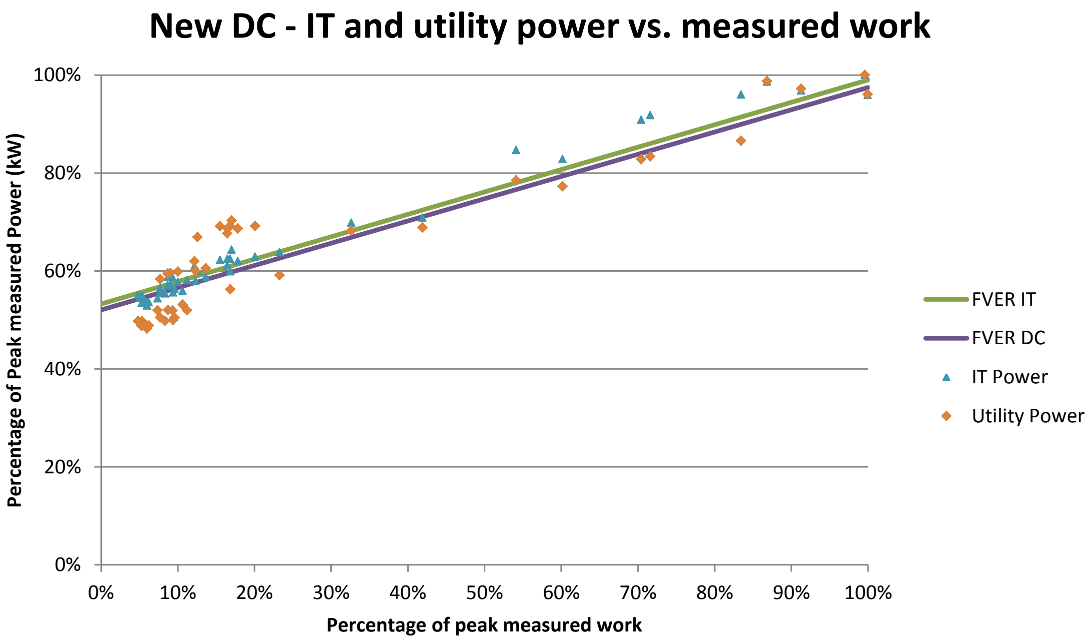

Figure 13 New data centre FVER plot

For this updated version of the data centre the metrics record the following rather better values for FVER at the IT and Utility;

Metric	Value
FVER (IT)	2.2
FVER (Utility)	2.2

It can easily be understood that having only 1.2 parts of energy consumption which is fixed against each one part of energy consumption which varies with work output is a substantial improvement over the previous data centre example.

Whilst it may be expected that the DC and IT lines should meet at 100% workload and 100% power draw these lines are a fit to a sample of data readings taken from a system which is not entirely linear and is influenced to some extent by other factors, therefore some scatter is to be expected.

For some data centres there may be a false correlation between the low measured work overnight and reduced utility power consumption when the cooling economiser is active. This may happen because the lower overnight temperatures allow the mechanical cooling to turn off and the data centre to use entirely “free cooling” during this period which may create a false impression of better FVER performance. To address this the Utility side FVER may be restricted to not exceed the IT FVER value.

4.4 Comparing performance

The changes in the data centre to go from our old example to our new example will have taken some time to implement, quite likely a year or more. In this time many elements of the efficiency metric will have changed. This will include the number and type of services delivered by the data centre, a certain amount of "IT efficiency" improvement will have occurred through the normal replacement and upgrade cycle, also the M&E changes have impacted efficiency. The FVER metric is still able to compare the data centre before and after as the productivity proxies may be updated with the changes to reflect the IT services delivered and the IT efficiency element is normalised.

The energy consumption of the data centre before and after the changes is summarised in the table and Figure 14 below.

Metric	Old data centre example	New data centre example
FVER (IT)	6.1	2.2
FVER (Utility)	7.6	2.2

The table shows the improvement in FVER measurement, both in the IT equipment and software from 6.1 to 2.2 and the parallel improvement in the overall FVER which has allowed the IT and Utility FVER values to match.

The operator outcome is shown in Figure 14 below with the weekday productivity in green and the weekday utility power draw for the old data centre in blue. The power draw for the new version of the data centre is shown in red but scaled to the peak power draw from the old data centre. The updated data centre peaks at only 75% of the old data centre power draw and exhibits a much greater reduction at lower workload resulting in an overall consumption of only 60% of the old version.

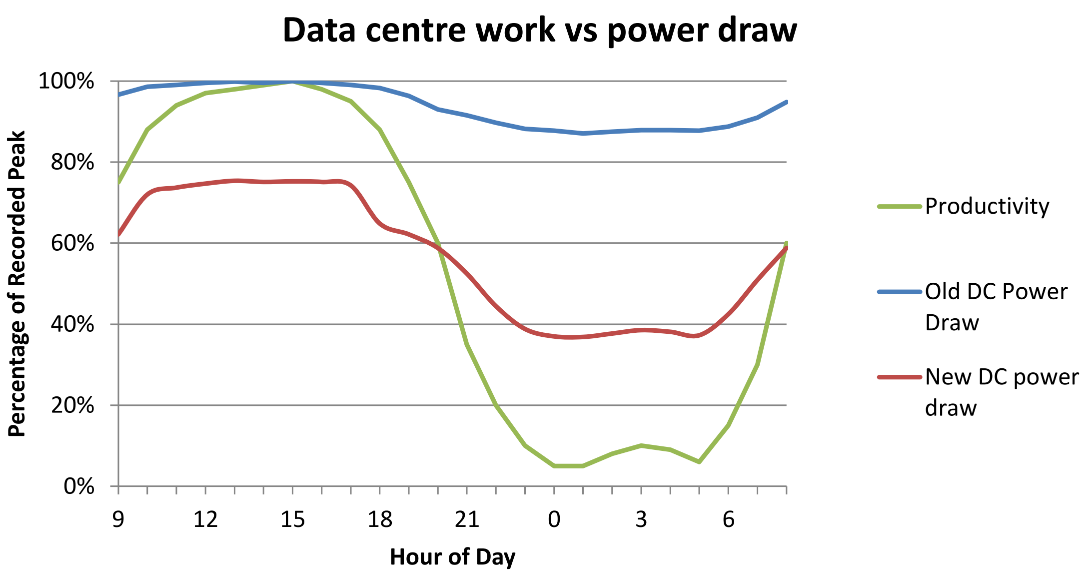

Figure 14 Old vs new data centre example

5. Measuring FVER

This section provides more detail on how to measure the component values and calculate FVER for a data centre. The key steps are the selection of a productivity proxy, measurement of productivity and power over sufficient time to give representative data and reporting of the resulting FVER.

5.1 Selecting a productivity proxy

The recommendation for FVER is that the operator selects and implements a high level measurement of the delivered “useful work” for their data centres which is as closely related to the value actions of the business as possible and not simply a measure of how heavily loaded IT devices are. This is a strong preference as it is known that the same unit of load on different IT devices does not represent activities of equal (or indeed any) value to the business.

The operator should select and record identified activities. Once recorded a relative weighting of value should be assigned to each recorded activity to represent the value to the business and the total activity value aggregated.

Many operators will have a large number of applications which serve a diverse range of needs within the business, some of which may not have an obvious business value. Aside from the clear opportunity to review whether systems for which no business value can be assigned should be eliminated, operators may choose to identify and track only the top few contributors to business value. A small set of applications will be responsible for the major value and these may be considered to be a proxy for the lower value applications

We are not concerned that these “work” values will not be comparable between operators as they are normalised in the process of calculating FVER. The requirement for FVER is that we can establish that the data centre was doing half as much useful work at 02:00 as it was at 14:00.

These measures might include;

- Customer interactions supported (web, telephone, physical)
- Transactions processed
- Web pages served (differing weights dependent upon value)

The productivity proposals contain both proxies which require full measurement of the entire data centre and those which use a sampled subset of machines or activities, this choice is left to the operator.

Combining measured work counters

Once the operator has identified the key activities and services to measure the service measurements should then be combined with a weighting appropriate to their business value for an overall score.

As an example a data centre may support a business with the following identified key activities;

- Telephone customer interaction supported by CTI and VOIP services
- Web customer interaction supported by web services
- Physical customer interaction in a physical store / branch supported by PoS services
- Central finance activity supported by finance platform

(note that these are provided as a simplified example)

To combine the multiple recorded activities which are considered to represent useful business value from the data centre into an overall ‘score’ for each period of measurement the operator should select relative weightings for each activity. These are likely to be specific to that data centre operator;

Activity	Hourly range	Weighting
Telephone customer interaction	0 – 10,000	3
Web customer interaction	100 – 15,000	2
Physical customer interaction	0 – 6,000	4
Central finance activity	0-200	50

For each sampled hour the total activity may be combined by;

$$W_{hour} = a.W_1 + b.W_2 + c.W_3 \dots$$

Where W_{hour} represents the “Total Work” in the measured hour, a, b, c… are the relative weightings and W_1, W_2, W_3 ... are the counted productive actions of different platforms.

Activity	Hourly count	Weighting	Weighted count
Telephone customer interaction	6,745	3	20,235
Web customer interaction	11,475	2	22,950
Physical customer interaction	3,812	4	15,248
Central finance activity	80	50	4,000
Total			62,433

Stability of proxy measurements over time

A key aspect of FVER is the ability to track improvements through changes which would cause a “useful work” metric to incorrectly report a reduction in efficiency. The operator should review the proxy measurements and weightings after (or preferably before) each significant change in services delivered.

As an example we may consider an operator whose primary activity is to deliver video content to Internet users. This operator implements a change to the encoding software they use for video which results in a reduction in the bandwidth required to deliver each customer video stream at the same perceived quality. The cost of this change is an increase in processor load on the video servers which increases power consumption. The overall effect is that the operator is able to serve more customer streams for a small increase in power consumption and a measurable reduction in egress traffic which represents a significant overall gain in energy and financial efficiency.

If the operator were to use a primitive proxy such as bits / kWh then their efficiency would apparently become substantially worse as the egress bits would reduce whilst the kWh would increase. This would falsely indicate that the efficiency had reduced despite the obvious improvement.

Under FVER the operator can assign a value to each complete video delivered and therefore be able to record the increase in total streams delivered. In reality the operator may well choose to assign a different weight to a low versus a high quality stream or a higher weighting to a higher value premium content stream.

5.2 Measure productivity and energy consumption

Once a suitable proxy has been selected and implemented this may then be measured over a sample period along with the power draw measurements for the data centre. In most cases the operator will already measure PUE and have the total IT power draw and the total utility draw available for the site.

All three measurements should be taken hourly for at least one day, preferably a longer period such as a week or month to provide a reporting value. This document will refer to hourly measurements to remove the need to convert between kW and kWh (as the average kW over an hour is the same as the kWh).

5.3 Normalising measured values

In FVER we are not concerned with the absolute productivity per unit energy (although the operator may also choose to track and compare the total productivity score with the total energy consumption if they believe that the productivity score is sufficiently stable and comparable over time). Instead we normalise the recorded values to being between 0% and 100% of the peak recorded in the period. A sample set of readings is shown in the table below;

Hour	Weighted productivity	IT energy (kWh)	Utility energy (kWh)
09:00	54,000	310	442
10:00	63,360	326	512
11:00	67,680	335	524
12:00	69,840	339	531
13:00	70,560	341	536
14:00	71,280	342	534
Peak	71,280	342	536

These recorded values are converted to percentages of their peak to provide the normalised values we will use to calculate the FVER;

Hour	Productivity	IT energy	Utility energy
09:00	76%	91%	82%
10:00	89%	95%	96%
11:00	95%	98%	98%
12:00	98%	99%	99%
13:00	99%	100%	100%
14:00	100%	100%	100%

5.4 Calculating the FVER

As there are two energy measurements collected, the IT energy and total utility energy an operator who controls the entire data centre may calculate their FVER for both the IT equipment and the entire data centre. Those IT operators who do not control the data centre or are not the sole tenant as commonly seen in colocation facilities may still calculate their IT FVER based on metered IT energy consumption.

Supplied spread-sheet

The steps in this section are all demonstrated in the supplied example spread-sheet.

Scatter plot

A useful first step in the representation of the measured data is to plot each point on a graph and manually inspect the pattern. Plot the hourly percentage productivity values as the x axis and the hourly percentage IT and Data Centre energy values on the y axis. A simplified example is shown below in Figure 15;

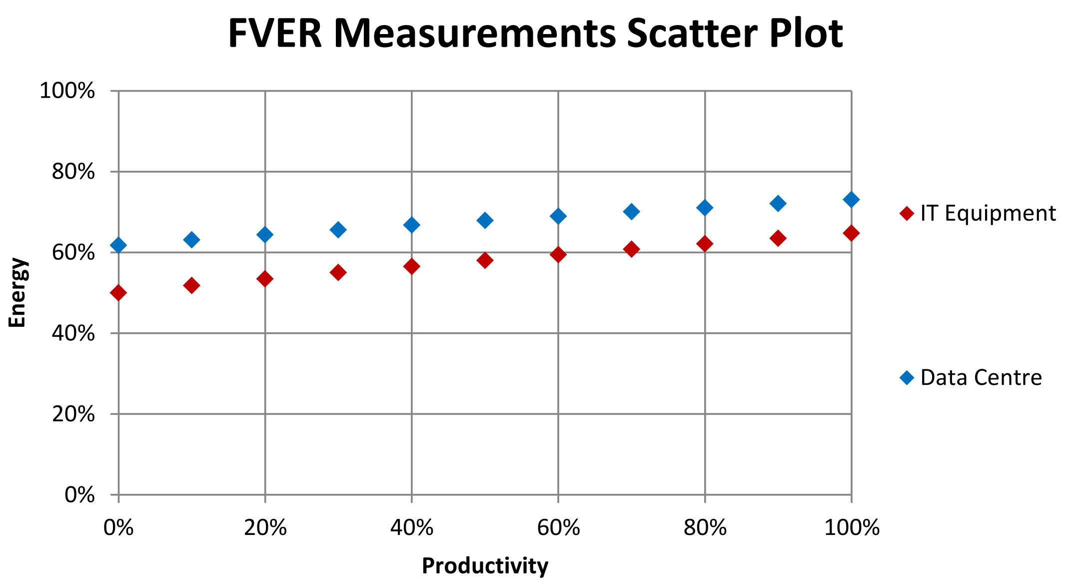

Figure 15 Simple scatter plot of normalised FVER measurements

This shows that there is some relationship between the hourly energy consumption and hourly productivity of the data centre but that the slope is shallow.

Linear regression

To obtain numbers to describe the shallow slope seen in the plot above a linear regression may be performed on the measured values, as stated any spread-sheet or statistical package (such as R [1]) may be used.

The linear regression will produce two key values, the slope and intercept. The slope in this case is how steeply our energy consumption rises with productivity and the intercept is an estimate of the energy at zero productivity i.e. where the energy line meets the vertical axis as shown in Figure 16 below.

In this case the values obtained for the linear regression are;

Measure	Intercept	Slope
IT Energy	0.50	0.15
Utility Energy	0.62	0.11

[1] The R project http://www.r-project.org/

Plotting these lines on the chart provides a strong visual indication of the fixed and variable energy components. The ideal plot would show both the IT equipment and Data Centre lines falling toward 0% as productivity falls to 0% and passing through the origin. This is clearly not possible in a real data centre but nonetheless represents a clear target and establishes the direction of improvement and the magnitude of further improvement available.

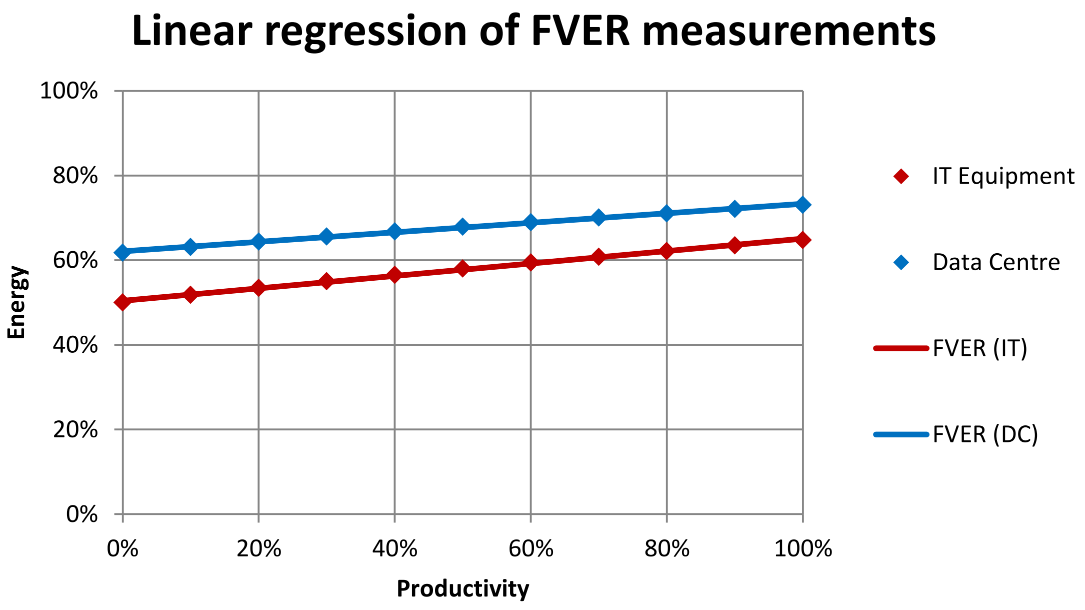

Figure 16 Linear regression of normalised FVER measurements

Both the individual slopes of the lines and the comparative slopes are of interest. It will be the case for many data centres that the FVER(DC) will have less slope than the FVER(IT) due to the fixed overheads of the data centre infrastructure.

Obtaining FVER

Given the linear regression values for the IT and Data Centre energy we can calculate the $FVER_{DC}$ and $FVER_{IT}$ for the site over the measurement interval;

$$FVER_{IT} = 1 + \frac{Intercept_{IT}}{Slope_{IT}}$$

$$FVER_{IT} = 1 + \frac{0.5}{0.15} = 4.\dot{3}$$

Similarly for the data centre energy;

$$FVER_{Utility} = 1 + \frac{Intercept_{DC}}{Slope_{DC}}$$

$$FVER_{Utility} = 1 + \frac{0.62}{0.11} = 6.6$$

This gives the following Fixed to Variable Energy Ratios for the sample site;

Metric	Value
FVER (IT)	4.3
FVER (Utility)	6.6

6. Practicalities of measurement

For these metrics to be measured on a real data centre a number of basic conditions need to be fulfilled;

a) A selected set of proxy measurement points on applications or services and operator subjective weighting factors to identify the "Work" delivered by the data centre as recognised by the business
b) Hourly measurement of the productivity / business action counters
c) Hourly measurement at reasonable accuracy of the IT power draw (as for PUE at the UPS output, IT PDU or IT power cords)
d) Hourly measurement at reasonable accuracy of the utility power draw
e) A reasonable number of readings taken prior to analysis (at least 24 hours across one day, preferably at least one week)

Effective measurement of the metric would require that a reasonable range of IT productivity values had been recorded in addition to a reasonable number.

An operator may wish to report FVER over longer time periods, weekly, monthly, quarterly and observe the short and long term variations.

6.1 Device utilisation is a poor proxy

In most sites the processor or device utilisation is a poor proxy for this measurement due to two key factors;

1. The peak processor utilisation values are commonly still < 25% in many data centres and this low utilisation is one of the causes of the poor relationship between work delivered and energy consumption that the metric seeks to address

2. Processor utilisation does not indicate that a data centre is doing work which is recognised as being valuable to the business

6.2 Impact of external temperature

It is likely that some data centres would record a higher than actual correlation between productivity and overall data centre power draw due to external temperature variation. In many cases the IT productivity will be lower overnight when the external temperature and therefore cooling overheads are lower.

This may be initially remedied by restricting the FVER utility to the minimum value of the FVER IT.

7. Reasons for the problem

In order to understand why the adoption and measurement of the FVER metric will assist operators in reducing, peak, off peak and total energy consumption it is useful to examine some of the apparent reasons for the very flat and continuous energy consumption at the UPS output and utility feed.

7.1 Server power variability with load

As already discussed in section 3.3 the power consumption of most servers is anything but linear with the work performed by the server. This issue was explored in some detail in the 2007 Google paper "Power provisioning for a warehouse sized computer"[vi] .

This issue is further compounded by the low achieved utilisation of many servers, meaning that far from seeing the server power draw swing between 50% and 100% over the full range of applied workload in most data centres as the delivered work changes the low utilisation instead restricts the power variation to a range more like the 50% to 60% of peak power shown in the grey area of Figure 17 below;

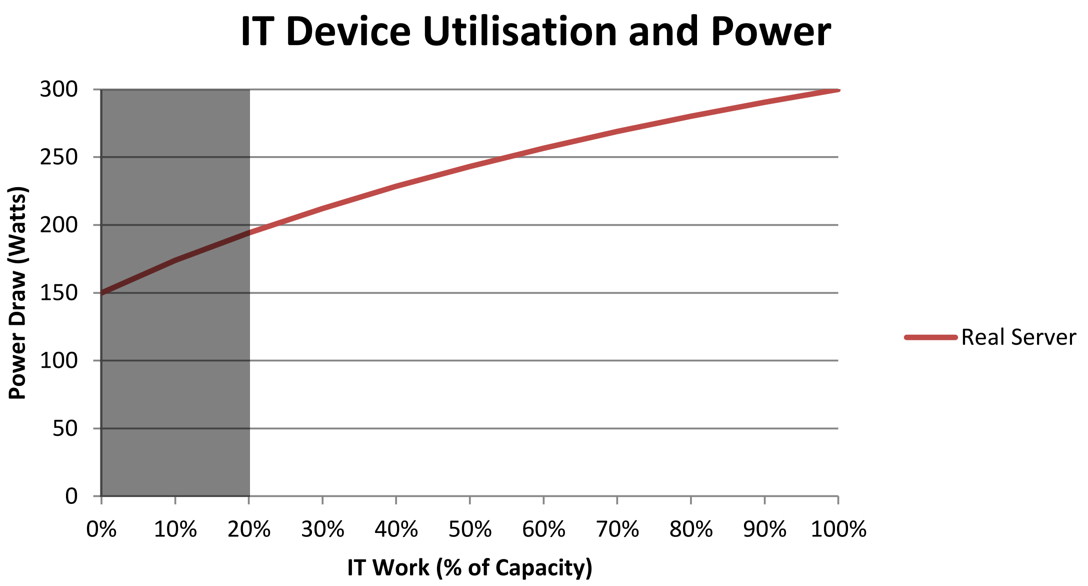

Figure 17 Common real world server utilisation

Virtualisation problems

Whilst many data centre operators have been addressing the problem of low utilisation server energy consumption through virtualisation and consolidation of workloads there is a secondary problem now becoming evident. It can be much harder for the server to manage power consumption at low load when virtualised. This reduces the ability of the servers to react to low overall workload conditions and is compounded by the reluctance of many operators to use power management functionality, particularly once a physical machine has many different service workloads.

The result of this is shown in Figure 18 below, whilst we may well reduce the energy consumption substantially by consolidating some low utilisation physical servers (blue) onto a single virtualised server (green) we lose much of the response to workload. This leaves the virtualised server consuming half of its peak power at zero real workload. The issue is even worse in environments where the operator chooses not to use the full power management capabilities (red).

The result of this is that whilst many operators who have virtualised their compute infrastructure have reduced their power consumption they have not reduced it by as much as they could with effective power management. They have also traded a reduction in the headline power consumption for an increase in the fixed share of their energy consumption making their costs less responsive to load and utilisation.

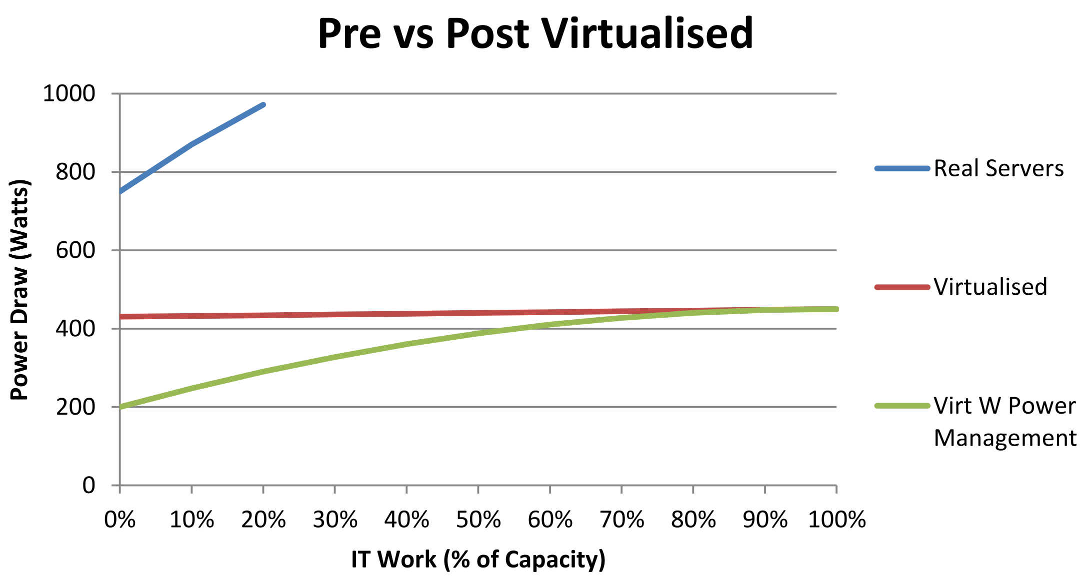

Figure 18 IT workload to power relationship for pre and post virtualised servers

Smarter software

Some virtualisation software allows operators to move running virtual machines across physical machines to allow for entire physical machines to be powered down when workload is reduced. More modern software which is written to a service cloud or resource broker allows for far more dynamic allocation and de-allocation of physical resources. Well known examples of this type of dynamic infrastructure are tool sets such as Cloud Foundry[vii] and service platforms such as Amazon's Elastic Compute Cloud[viii]. Unfortunately most operators are not yet comfortable with the additional complexity and perceived risk of completely dynamic resource allocation and machines which regularly turn themselves off, this is true even of many large cloud providers whose infrastructure is all on an using energy irrespective of the current utilisation. The FVER metric is designed to drive and reward this smarter use of available resources.

7.2 Network and storage

Whilst significant progress has been made in making the power draw of servers more responsive to applied load by programs such as Energy Star, chip manufacturers and IT vendors there has been significantly less progress in network and storage equipment. Much of the currently deployed equipment runs at virtually constant power consumption irrespective of workload. Variable power items such as disks which spin down at idle in storage systems are generally unable to make up for the large fixed power draw of the support systems and particularly the storage networking equipment.

This fixed power draw characteristic of the network and storage equipment has become more important as more data centres virtualise their servers, one of the results of server virtualisation is that the power draw of the servers with their strong power management capability is reduced but the requirements for remote storage increase the largely fixed power consumption of the network and storage equipment. This effectively reduces the variability in the power consumption of the supported servers as they are not separable from the network and storage.

A number of recent technology changes offer the opportunity for this fixed overhead to be addressed, in many cases 3.5" disks have been superseded by lower power 2.5" disks and these, in turn are being replaced with solid state disks whose idle power consumption can be very low. New Ethernet technologies which change speed and therefore power requirements with the traffic demand are also increasingly available; these are likely to be assisted by converged network technologies which dispose of many sets of devices.

The increasing use of software running on cheap, commodity compute nodes containing local storage to replace monolithic centralised storage devices and dedicated Storage Area Networks also presents a strong opportunity to reverse the trend of constant storage energy consumption.

7.3 Scaling this to the data centre

Given this view of the IT equipment behaviour under varying workloads we can add on the characteristics of the data centre to see the overall effect. In many data centres it is well known that there is substantial fixed infrastructure energy overhead and this further hampers the connection between IT workload and total utility power draw. This issue was explored in detail in the 2007 DCSG white paper “Data Centre Energy Efficiency Metrics”[ix]

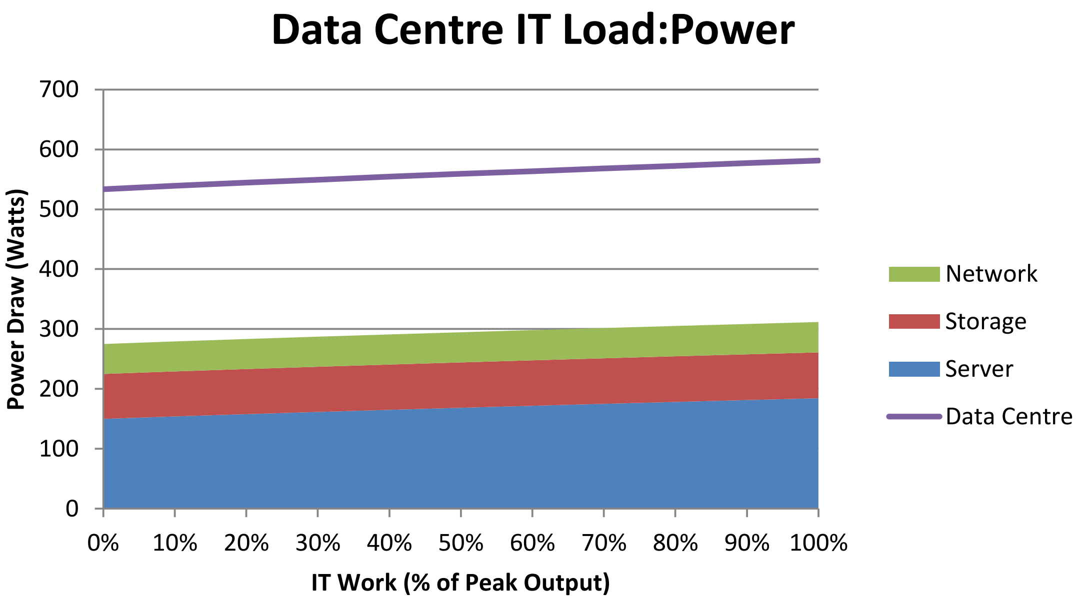

Figure 19 Workload to utility power relationship

As shown in Figure 19 above once the network, storage and data centre overheads are taken into account even the most aggressive power management of servers only results in small changes in the utility load of the data centre. This is borne out in the observed data from most data centres where the UPS load is almost constant and the utility load is largely fixed and generally more responsive to external temperature than IT load.

Whilst this may be seen to make an argument for focussing only on the removal of equipment supporting low value services there remains the opportunity to substantially reduce the power consumption at all levels of IT work through targeting the fixed loads and therefore fixed costs of the data centre.

8. Modifying the Green Grid proxy measurements

To use the FVER metric an operator needs to select one or more measures of the delivered work or value from their data centre. Rather than attempt to redefine these we recommend that operators select from the existing Green Grid productivity proxies. Some small modifications are required however as the DCeP proxies contain a mechanism to relate the indicated productivity to the total energy consumption, for the FVER metric this is not required and so the proxy measures may be simplified. This section describes how the productivity proxy measures may be modified for use in FVER.

8.1 Proxy measurements

The DCeP proxies are listed below in the groups in which they were placed in the Green Grid interim results document[x], those suitable as a proxy for FVER productivity measurement are identified;

Group	Proxy	Comment
1: DCeP proxies	#1 Useful work self-assessment and reporting	Preferred measure
	#2 DCeP subset by productivity link	Lower measurement overhead, still focuses on application output
	#3 DCeP subset by sample load	Not suitable
2: Bits per kW-hr	#4 Bits per kilowatt-hour	Not suitable
3: CPU Utilisation	#5 Weighted CPU Utilisation – SPECint_rate	Weak measure
	#6 Weighted CPU Utilisation – SPECpower	Weak measure
	#7 Compute Units Per Second	Not suitable
4: OS Instances	#8 Operating System Instances per kW	Only suitable in very dynamic environments

Table 1 Productivity proxies

8.2 DCeP Proxies

The DCeP proxies may be modified to be suitable for FVER work measurement as detailed below. Please see the Green Grid document "Proxy Proposals for Measuring Data Center Productivity" [iv] for information on the measurement process.

Useful work self-assessment and reporting

The useful work value W_i is recorded from applications during the measurement window and the proxy calculated using the formula;

$$Proxy_{DCeP} = \frac{\sum_{i=1}^{n}(N_i . W_i)}{E_{DC}}$$

As FVER does not attempt to directly compare the user specific W_i value we must remove the divisor as shown below;

$$W_{hour} = \sum_{i=1}^{n}(N_i . W_i)$$

Note that this is essentially the same as $W_{hour} = a.W_1 + b.W_2 + c.W_3$ ...

DCeP subset by productivity link
The useful work value is calculated by using an instrumentation SDK to record the productivity of key applications on a subset of servers, the values N_{DC} and N_{subset}indicate the total number of servers in the data centre and the number in the instrumented subset respectively;

$$Proxy_{ProdLink} = \frac{\left(\frac{N_{DC}}{N_{subset}}\right).\sum_{i=1}^{n} W_i}{E_{DC}}$$

As before the term for data centre energy is removed to use the raw productivity value for FVER;

$$W_{hour} = \left(\frac{N_{DC}}{N_{subset}}\right).\sum_{i=1}^{n} W_i$$

Weighted CPU utilisation
These measures are weaker for FVER as they measure all activity which loads servers, not just that which provides useful services and is also likely to have a less variability due to the workload floor.

$$Proxy_{SPECint} = \frac{T.\sum_{i=1}^{n}\left(U_{avg}CPU_i.B_i.\left(\frac{Clk\_CPU_i}{Clk\_B_i}\right)\right)}{E_{DC}}$$

As before, removing the term provides for a raw work indicator value;

$$W_{hour} = T.\sum_{i=1}^{n}\left(U_{avg}CPU_i.B_i.\left(\frac{Clk\_CPU_i}{Clk\_B_i}\right)\right)$$

9. Document Information

Version History

Version	Description	Version Updates	Date
1.0	Release	Incremented version number	25/05/2012
0.9	Review draft	Included review comments from Paul Latham, Vic Smith and Zahl Limbuwala	21/05/2011
0.8	Review draft	Updated summary	13/12/2011
0.7	Review draft	Added additional analysis on ITUE in section 0	30/10/2011
0.6	Review draft	Updated based on Stuart Buckland review comments	23/10/2011
0.5	Review draft	Added Impact of reducing the fixed energy component section	11/09/2011
0.4	Review draft	Added abstract, sample data centres, proxy measurements and general updates	11/09/2011
0.3	Draft	Added major content sections and restructured	31/08/2011
0.1	Draft	Initial concept	24/07/2011

Version History

Release History

Version	Description	Authoriser	Date
1.0	Public release	Liam Newcombe	25/05/2012

Release History

---

[i] SPECPower benchmark http://www.spec.org/results.html

[ii] Energy Star for compute servers http://www.energystar.gov/index.cfm?c=revisions.computer_servers

[iii] Future of data centre efficiency metrics http://dcsg.bcs.org/future-data-centre-efficiency-metrics-0

[iv] Proxy Proposals for Measuring Data Center Productivity http://www.thegreengrid.org/Global/Content/white-papers/Proxy-Proposals-for-Measuring-Data-Center-Efficiency

[v] Green IT Promotion Council Data center Performance Per Energy http://www.greenit-pc.jp/e/topics/release/100316_e.html

[vi] Power provisioning for a warehouse sized computer http://research.google.com/pubs/pub32980.html

[vii] Cloud Foundry http://www.cloudfoundry.com/

[viii] Amazon Elastic Compute Cloud http://aws.amazon.com/ec2/

[ix] Data Centre Energy Efficiency Metrics http://dcsg.bcs.org/data-centre-energy-efficiency-metrics

[x] Productivity Proxy Proposals Feedback – Interim Results http://www.thegreengrid.org/~/media/WhitePapers/White%20Paper%2024%20-%20Productivity_Proxy_Proposals_100509.pdf?lang=en