



PRODUCTIVITY PROXY PROPOSALS FEEDBACK – INTERIM RESULTS

EDITOR:

MARK MONROE, SUN MICROSYSTEMS

CONTRIBUTORS:

JON HAAS, INTEL

JESSICA KAISER, EMERSON NETWORK POWER

JOHN PFLUEGER, DELL

JACK POUCHET, EMERSON NETWORK POWER

FREEMAN RAWSON, IBM

PETER SNELLING, SUN MICROSYSTEMS

ROGER TIPLEY, HP

KATHRIN WINKLER, EMC

INTRODUCTION

In January 2009 The Green Grid (TGG) published a paper entitled, "Proxy Proposals for Measuring Data Center Productivity,"¹ which proposed eight candidates for a simple proxy that would be useful for estimating DCeP², but significantly easier to calculate than the DCeP metric is. It was the intent of TGG to gather feedback from users as to which of the proposed proxies they felt would be most useful and worthy of further investigation. Readers were asked to participate in an online survey to evaluate the usefulness of each proxy. This paper reviews the results of the online survey that were submitted in the four-month period following the publication of the paper, as well as the results of two focus groups that were conducted to obtain further user feedback. There were no runaway winners in either the survey or the focus groups, but much useful information was obtained, which will guide The Green Grid in further development of productivity proxies. The paper concludes with a summary of TGG's plans for further investigation and development of proxies for DCeP.



GATHERING FEEDBACK

To date TGG has received feedback from four different groups using two different methodologies of data collection – two qualitative and two quantitative. The feedback process began with the inclusion of instructions in the original paper for readers to go online and participate in a survey.

SURVEY DETAILS

The online survey asked the participants the basic question, "Would you use this proxy?" for each of the eight proxies. It then asked them to evaluate criteria for judging proxies and grade each proxy against DCeP, using those criteria. TGG received 48 responses from this first round of feedback, all from readers of the paper who voluntarily and anonymously chose to participate. A few months later the paper was presented to a group of data center managers at the Data Center Pulse conference and eight additional survey results were obtained. After reviewing the survey process and the feedback obtained, TGG decided more feedback was needed and conducted two focus groups.

FOCUS GROUP DETAILS

In an effort to simplify the feedback process for the focus groups, the eight individual proxies were grouped by type as follows:

Group 1: DCeP proxies:	Proxy #1: Useful Work Self-Assessment and Reporting Proxy #2: DCeP Subset by Productivity Link Proxy #3: DCeP Subset by Sample Load
Group 2: Bits per Kw-hr:	Proxy #4: Bits per Kilowatt-hour
Group 3: CPU Utilization:	Proxy #5: Weighted CPU Utilization – SPECint_rate Proxy #6: Weighted CPU Utilization – SPECpower Proxy #7: Compute Units Per Second (CUPS)
Group 4: OS Instances:	Proxy #8: Operating System Workload Efficiency

The first focus group was held online via WebEx™ in June of 2009. Participants were recruited from data center operators and data center efficiency experts. There were eight participants, representing eight different companies.

Participants were asked:

- What problems do you have in your data center that you are hoping to solve with a productivity proxy?
- What are your thoughts on using a productivity proxy?
- Comment specifically on each of the four proxy groups

The second focus group was conducted during the Datacenter Dynamics Conference in San Francisco in July of 2009. There were 11 participants, all data center operators, again representing eight companies. The nature of the questions and the discussion differed from the previous focus group:

- For each of the four proxy groups, what would you do with the information once it was collected?
- For each of the four proxy groups, what concerns would you have?
- Rank each group as to which proxies were most useful.
- What should TGG do next?



INITIAL FEEDBACK RESULTS

SURVEY RESULTS

After analyzing survey results, it is apparent that there is no conclusive winner. Every group responded differently. For example, when respondents were asked for each proxy, “Would you use this proxy?” the only proxy that received more yes votes than no votes was the “Bits per Kilowatt-hour” proxy and it did not receive an overwhelming percentage of yes votes. See the chart below:

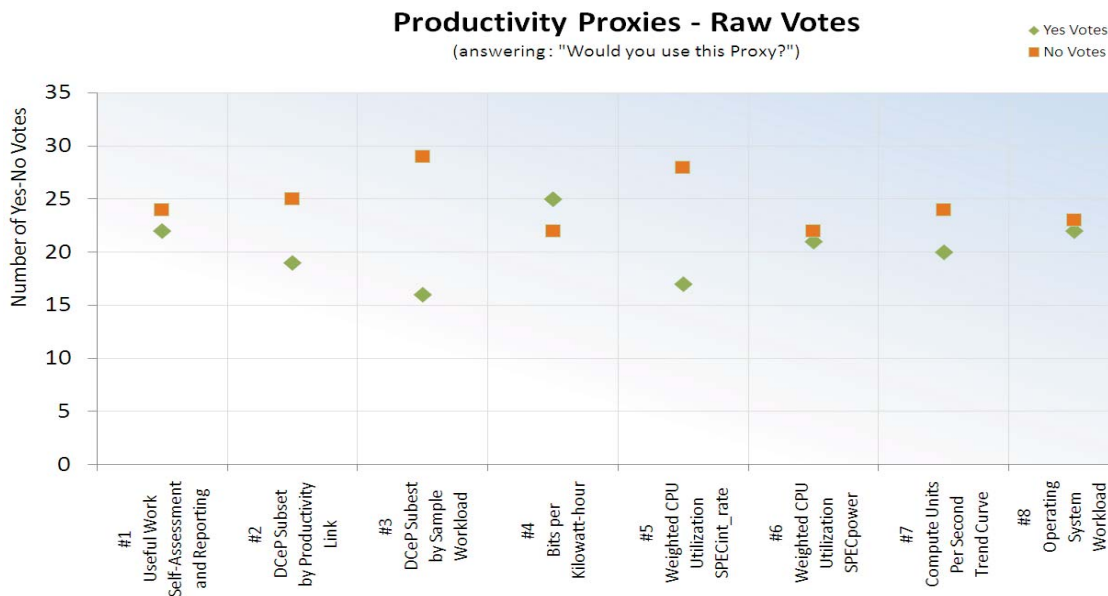


Chart 1 - Number of Yes/No votes

Survey responders were also asked to rank criteria used to evaluate proxies, and then rate the proxies relative to DCEP, based on those criteria. This data was aggregated in a Pugh matrix and the average results, utilizing the five highest ranking criteria, are shown in the chart below. Pugh scores were selected by respondents as “better than DCEP” (scored as +1), “same as DCEP” (scored as zero), or “worse than DCEP” (scored as

-1). The top five evaluation criteria, as ranked by survey responders, were Ease of Use, Cost to Implement, Accuracy, Objectivity, and Invasiveness.

Comparison to DCEP using Top 5 Criteria

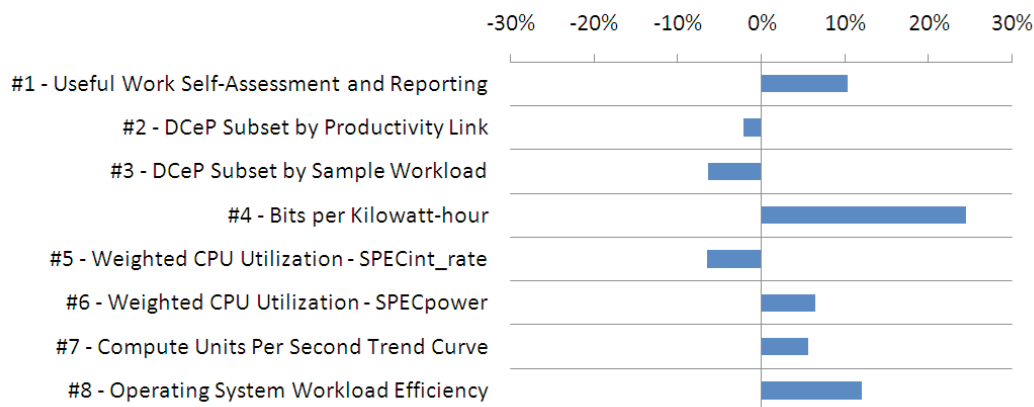


Chart 2 – Comparison to DCEP Using Top Five Criteria

FOCUS GROUP RESULTS

Feedback from the focus groups consisted entirely of comments, many of them repeated. All were considered and resulted in TGG drawing a number of conclusions which are listed in the summary on the following page. Comments that stood out include the following.

What problems do you have in your data center that you are hoping to solve with a productivity proxy?

- We have no way to make business decisions based on productivity.

What are your thoughts on using a productivity proxy?

- Expect that data centers may use multiple metrics depending on kind of data center.

Comments on each proxy group:

DCEP group:

- Not easy to implement.
- Contains too many fudge factors.
- Despite all the concerns, my gut feel is DCEP proxies are good, but need more guidelines.
- DCEP proxies are more accurate, but are more difficult to compute. Thus it makes more sense to use DCEP proxy in a homogeneous data center where only one application is running.

Bits per Kilowatt-hour:

- Simple, well-focused, and useful, but need guidelines as to when it should be used
- Network centric

CPU utilization group:

- Some homogenous data centers with large monolithic applications are already using CPU utilization as a measure of productivity

OS instance per watt:

- Simple to measure, but maybe should be used with other proxies

Please comment in general on the approach of using proxies to measure productivity:

- Definitely want to be able to compare data centers
- Need to incorporate storage

What should TGG do now?

- Do more case studies – examine who has been doing similar work and what their experience has been
- Try some of the concepts out. Do some experiments and report back. How long did it take to implement the program and what was done with the data?



SUMMARY OF RESULTS

After much analysis of the feedback, TGG netted four conclusions:

PROXY PREFERENCES DIFFER ACCORDING TO DATA CENTER CHARACTERISTICS

These studies showed that the industry has different needs for a proxy, based on the type of data center being operated. Some users are looking for the most simplistic, easy to perform measurement tool; some require something more complex and accurate; some are looking for something to analyze a homogeneous data center and others a heterogeneous environment. Most feel that more than one proxy will be needed to meet their corporate needs and to allow for the ability to measure servers, storage, and network productivity. Because of this, there is no general consensus on one best proxy among those that were proposed. The chart on the following page compares how the proxies are ranked by users across all four study groups. Ranking data was pulled from the survey responses by how many votes each proxy received when the participant was asked, “Would you use this proxy?” Focus groups were asked specifically to rank the proxies according to usefulness.

Proxy group	Proxy ranking by study group			
	Online Surveys (48 responses)	DCP Surveys (8 responses)	Focus #1 (2 responses)	Focus #2 (7 responses)
CPU	3	2	1	1
DCeP	4	3	2	2
Bits	1	4	4	3
OS/W	2	1	3	4

Chart 3 – Ranking of proxies by survey and focus groups

Analysis of comments made in the focus groups showed that respondents’ needs and perspectives could be categorized into two groups – homogeneous data centers and heterogeneous data centers. The homogeneous data center operators preferred the group of DCeP-based proxies, while heterogeneous data center operators did not.



Homogeneous data centers are run by companies whose business model is based around one compute-intensive application, such as weather modeling, online auctions, search engines, or streaming video. The data center is their factory, and they need a method for determining the cost of producing their product. If they cannot determine the costs of the data center transactions they cannot understand their business model. The methods these companies use to measure cost of transactions are typically proprietary, but they are already tracking the key metrics that give them this capability. They are also the authors of their own applications, so it is easy for them to instrument an application in a way that best suits their business. These companies are already measuring many of the things needed to calculate DCEP, and their feedback showed that they preferred the DCEP proxies.

Companies that operate heterogeneous data centers do not have the resources to instrument hundreds or thousands of different applications. Operators of heterogeneous data centers are thus not going to be able to implement DCEP or DCEP-based proxies, but want to use something far simpler. Among this group, most preferred CPU-based proxies.

Participants pointed out that every organization makes different determinations about what's important and thus proxies would be useful for comparing within data centers, but not across data centers. At the same time, it was felt that the ability to compare across data centers would be very useful.

Participants also stated that proxies should not be too generic, but that there needed to be several alternative proxies to meet the needs of different kinds of data centers. Related to this, one participant suggested that because data centers are constantly changing and upgrading, and because they are overlaid with various levels of protection relative to disaster recovery, test environments, etc., there should be some attempt to establish benchmarks of efficiency that are correlated to a matrix consisting of type and size of industry.

EASE OF USE IS IMPORTANT

In the online survey, participants ranked ease of use, cost to implement, and accuracy, highest on their list of important criteria for a productivity proxy. Ease of use, however, was the only criterion that was significantly different (outside of upper ANOM limit) from the mean response. See the chart below:

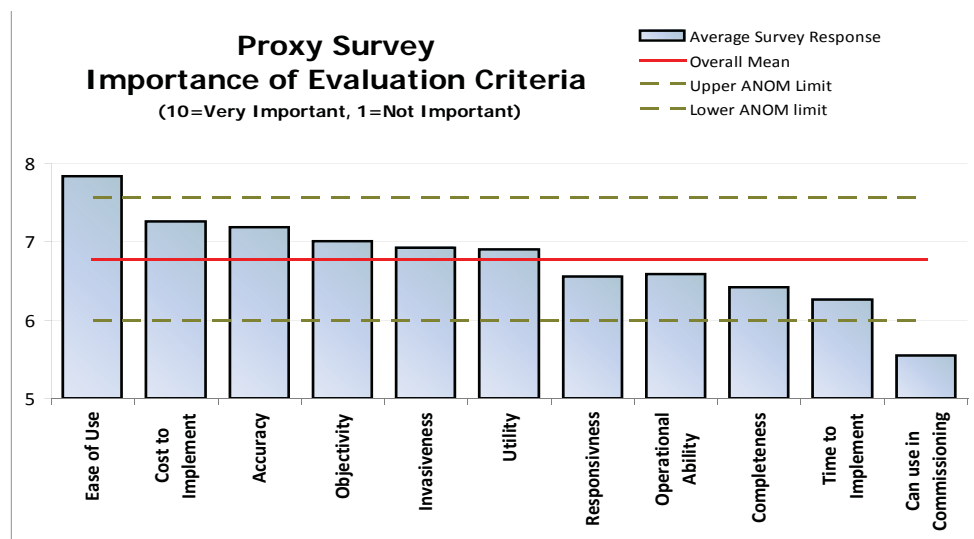


Chart 4 – Importance of Evaluation Criteria

Focus group participants also commented that ease of use was very important. Operators of homogeneous data centers said they preferred DCeP-based proxies because they were more accurate. TGG notes, however, that because many of these operators are already collecting transaction information anyway, and have a small number of applications to instrument, these proxies are easier for homogeneous data centers to implement than for heterogeneous data centers.

The heterogeneous data center operators preferred the CPU-based proxies. Among these, participants pointed out that the CUPS proxy was easy to implement because they were already tracking the age of their servers for asset management purposes.



PROXIES ARE SERVER CENTRIC

Every group brought out the point that, except for Bits per Kilowatt-Hour, the proxies are server centric, “yet only 60 percent of our data center is servers.” CPU-based proxies were the most popular because of ease of use, but many commented that they ignored storage, as did all of the other proxies.

THE INDUSTRY NEEDS ADDITIONAL DATA

Respondents were very enthusiastic about this effort, but it is clear that more work needs to be done. Everyone agreed that the proxy approach is valuable and needed because of the complexity of DCeP. They also understand that it is a hard problem and are appreciative of the fact that TGG is tackling this issue. We heard comments like, “This is really important work you are doing,” and “We know how hard this problem is, but this is exactly what we need for you to be doing. Keep after it.”

While TGG has gained valuable insight into the industry’s need for proxies, it also realizes that many of the responses were based on perceptions and not real-world experience with a proxy. More data needs to be collected from real applications.

FOLLOW-UP WORK

QUESTIONS STILL TO BE ANSWERED

This work has generated perhaps more questions than conclusions. TGG needs to do further work to answer the following questions:

- How do we put into practice what we are learning?
- Once data has been collected from a proxy, what do you do with it to obtain a useful result?
- How accurate can it be?
- What has already been done in this area that can be documented?
- How do we properly account for storage and networking when measuring data center productivity?

NEXT STEPS

At this time TGG is planning follow-up work in three different areas:

- Perform experiments. TGG will pick one proxy from each of the four groups and implement it in a data center.
- Conduct case studies. TGG will study homogeneous data centers which have already begun to track productivity metrics. Even though most of this information is proprietary, TGG hopes to gain insight that will prove useful to the general data center community.
- Incorporate storage. TGG will begin to tackle the problem of how best to incorporate storage into

a productivity metric and will be working with the Storage Network Industry Association³ (SNIA) to accomplish this task.



REFERENCES

¹The Green Grid, "Proxy Proposals for Measuring Data Center Productivity," 2009
<http://www.thegreengrid.org/en/Global/Content/white-papers/Proxy-Proposals-for-Measuring-Data-Center-Efficiency>

²The Green Grid, "A Framework for Data Center Energy Productivity," 2008, p. 8
http://www.thegreengrid.org/gg_content/White_Paper_13_-_Framework_for_Data_Center_Energy_Productivity5.9.08.pdf

³The Storage Network Industry Association, <http://www.snia.org/home>



ABOUT THE GREEN GRID™

The Green Grid is a non-profit trade organization of IT professionals formed to address the issues of power and cooling in datacenters. The Green Grid seeks to define best practices for optimizing the efficient consumption of power at the IT equipment and facility levels, as well as the manner in which cooling is delivered at these levels. The association is funded by four levels of membership, and activities are driven by end-user needs. The Green Grid does not endorse any vendor-specific products or solutions, but will seek to provide industry-wide recommendations on best practices, metrics, and technologies that will improve overall datacenter energy efficiencies. www.greengrid.org