

A Holistic Approach to **Energy Efficiency in Datacenters**

Dileep Bhandarkar, Ph.D. Distinguished Engineer, Global Foundation Services **Microsoft Corporation**

Published: August 2010



A Holistic Approach to Energy Efficiency in Datacenters

Datacenters have caught the attention of the computing industry over the last few years, fueled mainly by the tremendous growth in online services and cloud computing. Global consortiums of IT companies and professionals like The Green Grid were formed a few years ago to standardize on a common set of metrics like Power Usage Effectiveness (PUE), processes, methods, and new technologies to improve energy efficiency and business computing ecosystems in datacenters. Some companies like Microsoft have embraced PUE as useful metric, but other purists have raised objections. Like many metrics, PUE is not perfect and can be misused or gamed. There can be inconsistencies in how it is calculated. In spite of this, we have found it to be a useful tool. For a given datacenter, using a consistent measurement technique it can be useful to identify deviations from the norm. More careful analysis can then pinpoint the reasons for the deviation.

When comparing two similar datacenters owned and operated by the same company, it can be used to identify issues. Microsoft has used PUE as a metric to improve the energy efficiency of one of our older datacenters over a period of time.

The figure below shows the PUE value for a Microsoft facility from August 2004 to August 2007. After two years of energy improvements, the PUE for the site improved by 25 percent. Examples of non-intuitive changes made to the site, which proved effective in reducing PUE, were cleaning the roof and painting it white, and repositioning concrete walls around the externally-mounted air conditioning units to improve air flow. Both of these changes were validated as effective by measuring the effect on PUE.



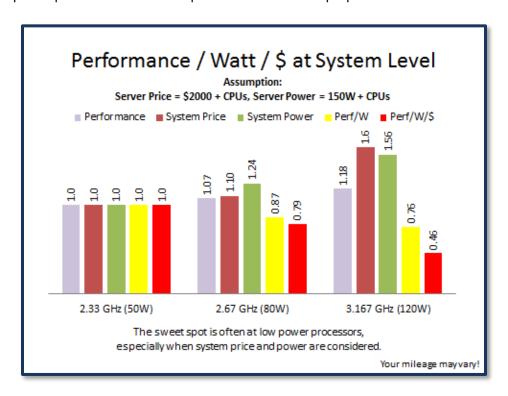


But just improving PUE should not be an organization's goal. The real goal is to eliminate waste and pack as much compute capability in the available power budget. PUE can be a useful indicator of energy efficiency, but it can also mislead you if used blindly. Take for example a scenario in which the fans in a server can be removed without impacting its performance. The elimination of fan power reduces the IT power (fan power is part of IT power), improves energy efficiency, but it also increases the PUE!

Another historical trend in the industry that has held over from the early days of colocation has also caused problems in the current environment. The colocation providers did not have knowledge or control over the servers to be deployed and they tended to build guard bands to guarantee in their level of service. The equipment owners often wanted to protect their interests by assuming that they could not count on the colocation providers to meet what was promised. Two levels of guard bands can create incredible waste!

We believe that careful and accurate monitoring of all aspects of your datacenter combined with thoughtful and holistic design of the infrastructure and the servers together can lead to significant improvements in energy efficiency.

Our measurements of server performance under load and workload analysis has enabled us to rightsize our server platforms. We have eliminated unnecessary components, used higher efficiency power supplies and voltage converters and bounded the expandability of server platforms to achieve significant power savings. We remain very much focused on performance per dollar per watt as an additional means of achieving higher energy efficiency. The chart below illustrates how low power processors can indeed provide a better value proposition.





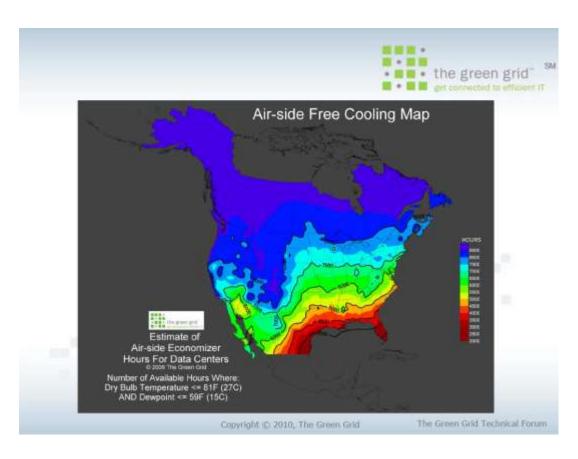
The performance, price, and power numbers were based on public information available in early 2009. In a truly distributed computing environment, where each server does not need to deliver the maximum possible performance, the servers with the best performance per dollar per watt are the best choice.

More than a year ago, we made a big bet on container technology for our Chicago datacenter. We chose water cooled technology for this initial indoor deployment. We developed an external specification that listed the amount and type of power, water temperature and flow rate, and a few other parameters and challenged the industry to innovate. We received very different engineering solutions from different vendors, but our goal was to maximize the total available computation while using up the allocated power budget. We did not want to strand power. We allowed our vendors to stack containers if needed to ensure that they were not space constrained. I am happy to say that we achieved PUE levels around 1.2, which is much lower than our published worldwide average of 1.6!

We have learned a lot from our container experience in Chicago and have now moved more towards containment as a useful tool. We do not see ISO standard containers as the best way of achieving containment in future designs. Our experience with air side economization in our Dublin datacenter has taught us a lot too. Our discussions with server manufacturers has convinced us that we can widen the operating range of our servers and use free air cooling most of the time. Our early proof-of-concept with servers in a tent, followed by our own IT Pre-Assembled Component (ITPAC) design (shown in the photograph below) have led us to a new approach for our future datacenters.



Climate data from The Green Grid shows that free air cooling can be used exclusively in many places. Because many systems can run at higher temperatures without problems, using outside air to cool computing equipment is viable. The addition of evaporative cooling allows operation when the outside temperature goes about the 90 to 95 degrees fahrenheit inlet that the equipment needs. Evaporative cooling is especially well suited for climates where the air is hot and humidity is low. Nature's most efficient means of cooling is through the evaporation of water. Evaporative cooling works on the principle of heat absorption by moisture evaporation. The evaporative cooler draws exterior air into special pads soaked with water, where the air is cooled by evaporation, and then circulated. When considering water evaporating into air, the wet-bulb temperature, as compared to the air's dry-bulb temperature is a measure of the potential for evaporative cooling. As the difference between the two temperatures increases, the evaporative cooling effect improves. When the temperatures are the same, no net evaporation of water in air occurs, thus there is no cooling effect. Hence, evaporative cooling is only efficient when the relative humidity is low, restricting its effective use to dry climates. Evaporative coolers are colloquially referred to as swamp coolers in the U.S. In other places they are known as desert coolers. We can add secondary cooling with air or water cooled chillers in places where free air cooling is not practical all the time. This additional mechanical cooling does add to capital expenditure, but its infrequent use reduces overall energy consumption.





Microsoft views its future datacenters as driving a paradigm shift away from traditional monolithic raised floor mega datacenters. Using modular pre-manufactured components yields improvements in cost and schedule. The modular approach reduces the initial capital investment and allows you to scale capacity with business demand in a timely manner.

Our combined datacenter and server engineering team is very excited to embark on this exciting journey to deliver highly optimized compute infrastructure for our online services.

Terms & Conditions

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, this document should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2010 Microsoft Corporation. All rights reserved.

Microsoft and the Microsoft Logo are either registered trademarks or trademarks of Microsoft Corporation, in the United States and/or other countries.