Water Efficiency Management in Datacenters (Part I): Introducing a water usage metric based on available energy consumption

Ratnesh Sharma, Amip Shah, Cullen Bash, Tom Christian, Chandrakant Patel

HP Laboratories
HPL-2008-206

Keyword(s):
Water Efficiency; Data Centers; Water Usage Metric; cooling Tower; power generation

Abstract:
The demand for data center solutions with lower total cost of ownership and lower complexity of management is driving the creation of next generation datacenters. The information technology industry is in the midst of a transformation to lower the cost of operation through consolidation and better utilization of critical data center resources. Successful consolidation necessitates increasing utilization of capital intensive "always-on" data center infrastructure, reduction in the recurring cost of power and management of physical resources like water. A 1MW data center operating with water-cooled chillers and cooling towers can consume 18,000 gallons per day to dissipate heat generated by IT equipment. However, this water demand can be mitigated by appropriate use of air-cooled chillers or free cooling strategies that rely on local weather patterns. Water demand can also fluctuate with seasons and vary across geographies. Water efficiency, like energy efficiency is a key metric to evaluate sustainability of IT ecosystem. In this paper, we propose a procedure for calculation of water efficiency of a datacenter while providing guidance for a management system that can optimize IT performance while managing the tradeoffs between water and energy efficiency in conventional datacenters.

External Posting Date: November 21, 2008 [Fulltext] Approved for External Publication
Internal Posting Date: November 21, 2008 [Fulltext]

To be presented at International Conference on Water Scarcity, Global Changes, Groundwater Management Responses, December 1st-5th 2008, Irvine, CA

Water Efficiency Management in Datacenters (Part I)
Introducing a water usage metric based on available energy consumption

Ratnesh Sharma*, Amip Shah, Cullen Bash, Tom Christian, Chandrakant Patel
Hewlett-Packard Co. Laboratories
1501 Page Mill Road
Palo Alto, CA

Abstract

The demand for data center solutions with lower total cost of ownership and lower complexity of management is driving the creation of next generation datacenters. The information technology industry is in the midst of a transformation to lower the cost of operation through consolidation and better utilization of critical data center resources. Successful consolidation necessitates increasing utilization of capital intensive “always-on” data center infrastructure, reduction in the recurring cost of power and management of physical resources like water. A 1MW data center operating with water-cooled chillers and cooling towers can consume 18,000 gallons per day to dissipate heat generated by IT equipment. However, this water demand can be mitigated by appropriate use of air-cooled chillers or free cooling strategies that rely on local weather patterns. Water demand can also fluctuate with seasons and vary across geographies.

Water efficiency, like energy efficiency is a key metric to evaluate sustainability of IT ecosystem. In this paper, we propose a procedure for calculation of water efficiency of a datacenter while providing guidance for a management system that can optimize IT performance while managing the tradeoffs between water and energy efficiency in conventional datacenters.

Introduction:

The design and operation of the data center infrastructure is one of the primary challenges facing IT organizations. Unprecedented growth in demand for IT services has led to development of large, complex, resource-intensive IT infrastructure to support pervasive computing. Emerging high-density computer systems and consolidation of IT resources into fewer data centers are stretching the limits of data center capacity [1] in terms of power and resource utilization. The industry is in the midst of a transformation to lower the cost of operation through consolidation and better utilization of critical data center resources. The large number of components in a data center including cooling systems, power systems, and computer systems and the diversity of these components make data center design and operation complex. Successful consolidation necessitates increasing utilization of capital intensive “always-on” data center infrastructure, and reducing recurring cost of physical resources. For the future computing utility, management of physical resources for operation of datacenters will be a requirement from an economic and sustainability standpoint. To improve customers’ RoIT (Return on Information

*Corresponding Author: ratnesh.sharma@hp.com

Technology) [2], it is critical to maximize the resource utilization efficiency of the data center and simplify its management.

Data Center Infrastructure:

Figure 1 shows the basic data center building blocks from utility grid to the cooling tower [3][4]. Switch gear, comprised of transformers and static switches with associated panels, distributes power to the cooling infrastructure and the IT infrastructure. The cooling infrastructure is comprised of chillers, cooling towers, computer room air conditioning (CRAC) units and primary/secondary pumps. The IT infrastructure includes servers, network devices and storage devices housed in standard racks. Uninterruptible Power Supplies (UPS) maintain power quality during normal operation and provide energy storage to operate the IT infrastructure during brown outs or short power outages. Chillers provide chilled water to the data center room that houses the server racks and other IT equipment.

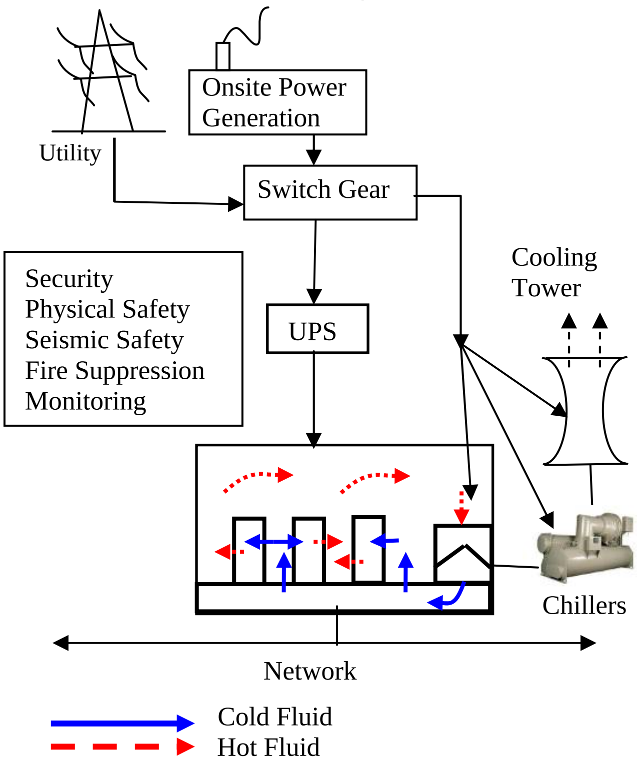

Figure 1: Data Center Building Blocks

Figure 2 shows the detail of the datacenter room including CRAC units, server racks and air flow paths [3]. Data centers are typically air-cooled with a raised floor plenum to distribute cool air, power and networking. Figure 2 depicts a typical state-of-the-art data center air-conditioning environment with under-floor cool air distribution. Computer room air conditioning (CRAC) units cool the exhaust hot air from the computer racks. Energy consumption in data center cooling includes work done to distribute the cool air to the racks and to extract heat from the hot exhaust air. A refrigerated or chilled water cooling coil in the CRAC unit extracts the heat from the air and cools it within a range of 10 ^{o}C to 18 ^{o}C. The air movers in the CRAC units pressurize the plenum with cool air which enters the data center through perforated vent tiles located on the raised floor close to the inlet of the racks. Typically the racks are laid out in rows separated by hot and cold aisles as shown in Figure 2. This separation is done for thermal efficiency considerations. Air inlets for all racks face cold

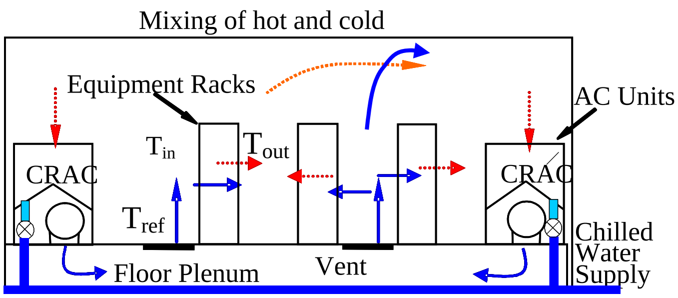

Figure 2: Cross section of the datacenter

*Corresponding Author: ratnesh.sharma@hp.com

aisles while hot air is expelled to hot aisles. A number of other equipment layout configurations and non-raised floor infrastructures also exist.

Figure 3 shows details of the cooling infrastructure of a typical datacenter [3]. Typically, vapor compression based refrigeration chillers provide chilled water to the CRAC units. Exhaust air from the IT equipment (as described in figure 2) dissipates heat to the chilled water inside the CRAC units. The warm water is returned to the chiller for heat rejection. A condenser water loop carries this heat for subsequent rejection at the cooling tower. Water is lost by evaporation to the ambient environment during this process of heat rejection. Evaporation loss depends on the moisture content of the air and the air temperature. In the case of air cooled chillers, the cooling tower is replaced with a heat exchanger for this purpose, thus preventing loss of water to environment. However, since air-to-liquid heat exchangers have lower effectiveness, air cooled chillers have lower coefficient of performance.

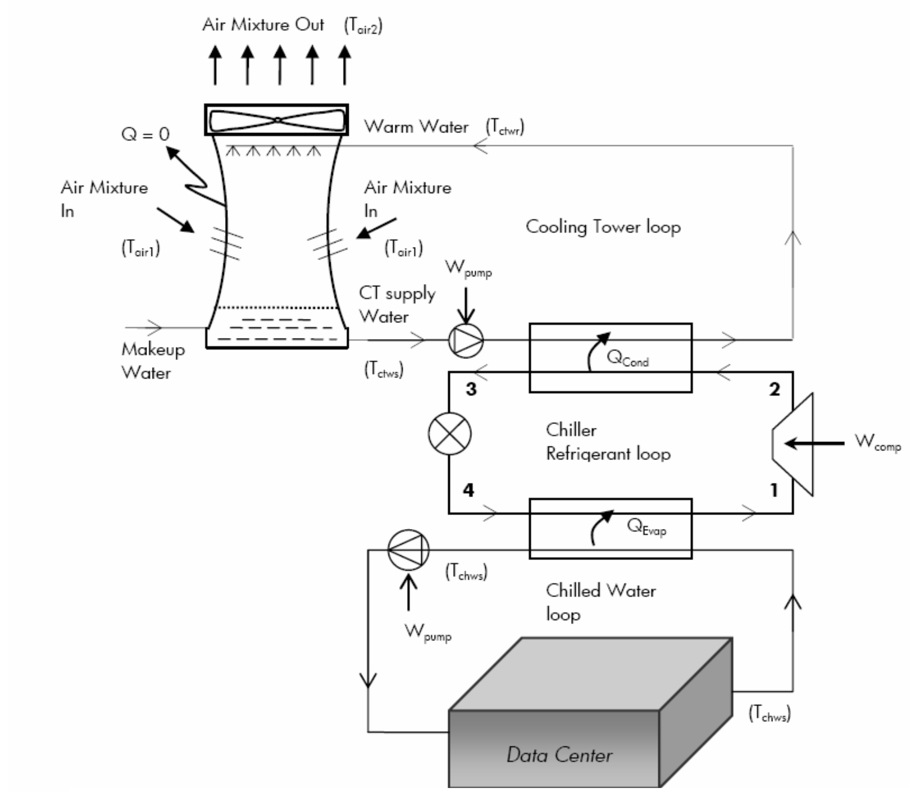

Figure 3: Data Center Facility Cooling infrastructure

Water Usage in Power Generation:

Water is a key component in the physical resource mix that powers datacenters: the generation of electricity usually requires available water for withdrawal and consumption, sometimes up to 30 gallons of freshwater for every kilowatt hour (kWh) generated in the case of some coal plants [5]. Water is mostly used in coal-fired steam plants, natural gas– and/or oil-fired steam plants, nuclear plants, biomass-fired steam plants, municipal-solid-waste (MSW) fueled steam plants, natural gas– and/or oil-fired combined cycle plants, and coal or petroleum residual–fueled gasification combined-cycle plant, with a majority required as cooling water for the condensing of steam. Figure 4 shows the schematic of cooling tower with makeup water supply and a bleed system. Cooling occurs in a tower by the mechanisms of evaporative cooling and the exchange of

*Corresponding Author: ratnesh.sharma@hp.com

sensible heat. The loss of heat by evaporation lowers the remaining water temperature. The smaller amount of cooling also occurs when the remaining water transfers heat (sensible heat) to the air. The rate of evaporation is about 1.5% of the rate of flow of the re-circulating water passing through the tower for every 5.5^{o}C decrease in water temperature achieved by the tower. Make up water accounts for any losses due to evaporation and bleeds. Bleed system is necessary to maintain water quality to prevent fouling and microbial growth. Open-loop, or once-through cooling systems, have relatively large water withdrawal rates and are still widely used, consuming large quantities of fresh water. In power plants using closed-loop cooling systems with cooling towers, the water consumption is generally proportional to the amount of water that is lost to direct evaporation within the power plant. Therefore, power plants with cooling towers require much lower fresh water withdrawal than open loop systems, but tend to consume nearly twice as much water in the power plant itself. Typical water withdrawal rates for Rankine-cycle plants burning coal, oil, or natural gas to be 20,000 to 50,000 gallons per MWh generated [6]. The lower end of the flow rate range corresponds to the higher temperature differential at the cooling tower, and vice versa. Air-cooling systems, which have negligible cooling water demand, can be considered as a replacement where water is scarce, but these tend to be less energy efficient. Essentially, water demand for electricity generation is expected to increase fairly proportionally to the amount of electricity generation.

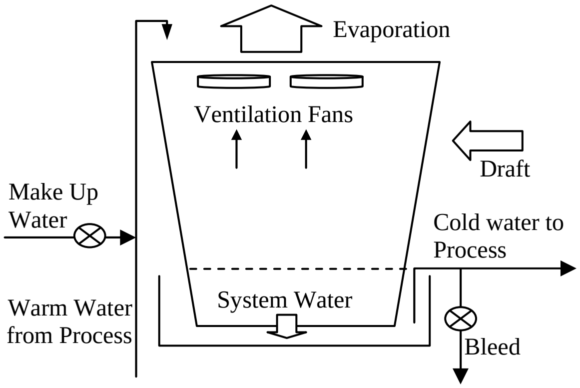

Figure 4: Cooling tower schematic

Apart from water usage in power generation, the usage of water itself also impacts the electricity demand by increasing the load on the power system. For example, water distribution systems and water treatment plants consume 1MWh and 0.25MWh, respectively, for every million gallon of water processed. Electricity use in the water sector could nearly double by 2015, far outpacing population growth [7].

Results

Water efficiency is calculated as water usage for every unit of useful work generated. In datacenters, water usage is critical both in design and operation. We will focus on operational water usage in this paper.

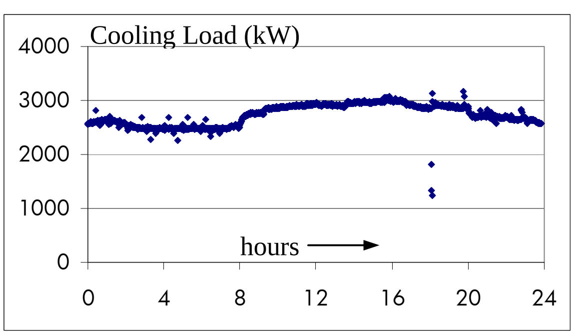

Figure 5: Typical demand profile in an IT facility

Typical cooling demand profile for an IT facility over a 24 hour period is shown in figure 5. The facility consists of a datacenter and office space. The cooling

*Corresponding Author: ratnesh.sharma@hp.com

load includes power consumed by IT equipment like servers, storage arrays, network devices etc. and office space cooling. All datacenter IT power demand is considered critical while office space cooling is considered non-critical. Power consumption by the cooling infrastructure is not included in figure 5 and is discussed later. Both air cooled and water cooled chillers are used to provide chilled water for cooling. The utilization of water cooled chillers varies with time and demand. Operation schedules of the chillers affects both the power consumption and water consumption of the facility. Figure 6 shows the utilization of the water cooled chillers as fraction of the total cooling load for two representative days. Observe that water cooled chillers accounted for a greater fraction of cooling during the period "DAY 1" as compared to that during the period "DAY 2". The demand profiles during the periods of interest were identical. We consider water consumption data for these two representative days to evaluate the tradeoffs between water utilization and power utilization for identical levels of cooling demand.

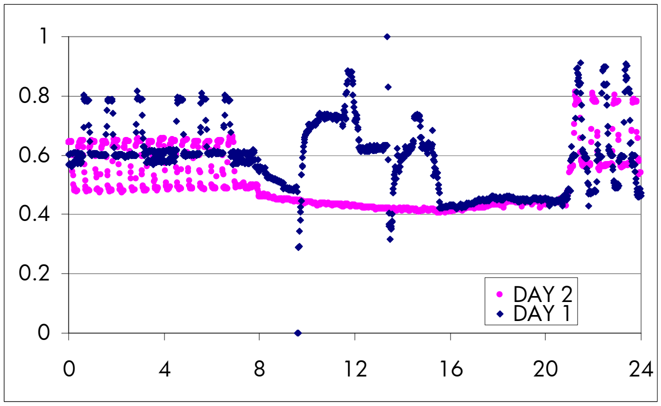

Figure 6: Water cooled chiller load as a fraction of total facility cooling load

Water-cooled chillers are more energy efficient than air-cooled chillers. For a constant cooling demand, increased utilization of water-cooled chillers reduces power consumption by the cooling infrastructure but increases water consumption at the facility. However, reduced power consumption by water-cooled chillers can also reduce water consumption indirectly at the power generation source. A 1MWh reduction in power consumption can indirectly offset 500-600 gallons of water consumption at the power generation source, assuming a fossil-fueled powered source. Figure 7 shows the make up water pump operation for the time periods in question. The pump supplies water to the cooling tower circuit (see figure 4) to compensate for loss of water due to evaporation and blow down (bleed). Observe the higher intensity of pump operation in "day 1" compared to that in "day 2". Operation of the water cooled chillers (see figure 6) increased the consumption of water on "day 1".

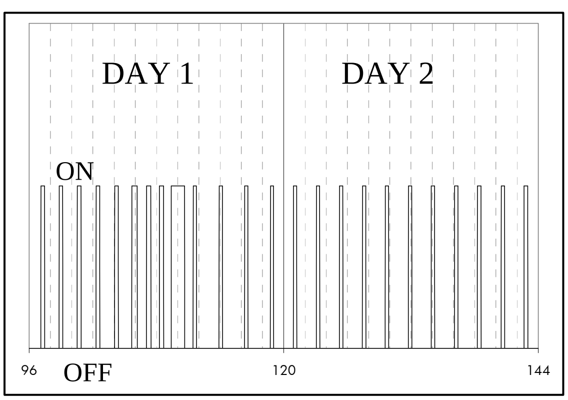

Figure 7: Make-up water pump operation during "day 1" and "day 2"

Analysis of diurnal data revealed that the "day 1" water consumption was higher than that of "day 2" by around 3000 gallons. Power consumption by cooling infrastructure on "day 1" was lower than that of "day 2" by over 1.4MWh. Table 1 shows the overall direct water

*Corresponding Author: ratnesh.sharma@hp.com

consumption and energy consumption for cooling only, between “day 1 and day 2. For the short period of operation of water cooled chillers, the impact on power saving is significant although the water consumption is considerably higher. The electrical (or heat) energy involved in treatment and distribution of cooling tower quality water may offset the energy savings obtained from the operation of water cooled chillers.

Table1: Water Consumption and Electrical Energy Consumption (for cooling) during day 1 and day 2

DAY 2		DAY 1	
Direct Water Consumption (gallons)	Energy Consumption (kWh)	Direct Water Consumption (gallons)	Energy Consumption (kWh)
19980	34240	23010	32850

Discussions

Understandably, water and energy are interconnected in various ways. Electrical power is used in treatment and distribution of water. Water treatment power consumption can vary greatly based on processes involved. Water distribution power consumption can vary greatly as well based on the location of the datacenter. The impact of delivery of IT services on water usage can be analyzed as the energy footprint of its water consumption.

Datacenter water usage can be directly expressed as gallons of water consumed in performing useful IT services. If the power consumption of the IT equipment is used as a proxy for the IT services, then the data center water usage could effectively be represented as the water consumption per unit time, divided by the datacenter IT power consumption. But although easy to use and calculate, such a metric is not dimensionless and does not reflect the energy costs associated with obtaining water in water-scarce regions of the world. Extraction of water from ground or desalination can be highly energy intensive processes. Any water usage metric should capture the energy impact of water usage as a whole.

To capture this holistic energy impact of water use, we propose a *Datacenter Water Usage Energy Metric* (ω), defined as the ratio of the energy “footprint” of water consumption over the power consumption of IT equipment. The energy footprint of water usage is the energy required to treat and distribute the water to the location of demand. Water usage is the total consumption rate of water from a natural source, including both direct (cooling) and indirect (power generation) usage. In our example, use of water-cooled chillers increased direct water usage by 3000 gallons per day and decreased indirect water usage by 700 gallons per day. The energy footprint of water usage is obtained from local utility or water supply district. Datacenter water usage metric is defined as:

*Corresponding Author: ratnesh.sharma@hp.com

$$\omega = \frac{(\text{Power Consumption in direct water usage}) + (\text{Power Consumption in indirect water usage})}{(\text{Power Consumed by IT equipment})} \times 10^{3}$$

where power consumption represents average power over a predefined period of time. For the discussed cases, w for “day 1” was 1.30 while w for “day 2” was 1.25 (see Appendix A). The overall impact is negligible considering small period of activity of the extra water cooled chiller in “day 1” and the contribution from other air cooled chillers as well. Use of water efficient and water-free technologies for operation of datacenters can be evaluated on a common framework. Since water consumption does not scale linearly with IT capacity depending on the mix of cooling infrastructure and operational criteria of water-cooled chillers, the metric should be evaluated over time as the datacenter grows with evolving business needs. In case of grid-supplied power, indirect water usage can be calculated from the power mix of the grid. For example, nuclear power plants have the highest water withdrawal rates (1200 gallons/MWh) [6]. Figure 8 describes the methodology for calculation of energy metric. Some of the key steps involve measurement of direct water consumption in cooling infrastructure, obtaining the indirect water usage for the power mix for the datacenter and the energy impact of water distribution and treatment.

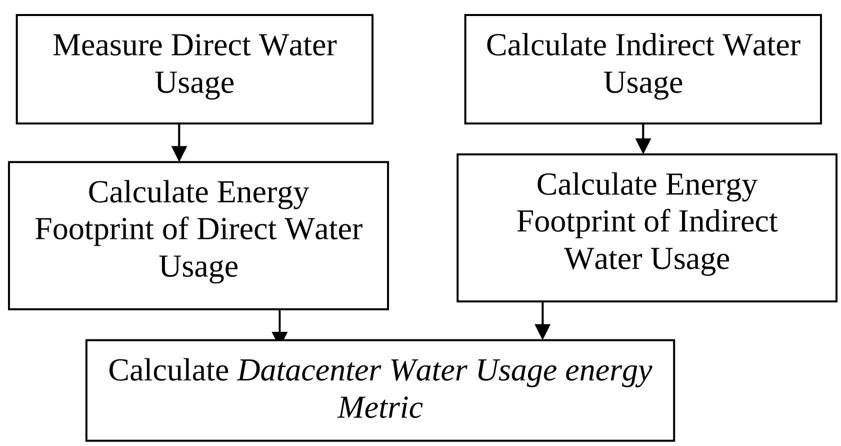

Figure 8: Flowchart for calculation of data center water usage energy metric.

Data center water efficiency can be managed by minimizing the water usage energy metric for a given IT power consumption. The metric can also be used to compare water efficiency across datacenters. There is a need to do seasonal benchmarking of this metric to capture the effect of regional weather patterns. As a preliminary approach, one could define a design curve for the metric in time and create operational policies for chiller and IT operation to meet the targets. Such a design curve could be a part of service level agreements with local utilities or administration. Another approach may be to include the metric as a part of coefficient of performance of the datacenter ensemble [4]. Needless to say, management policies are an open area of research at the present moment.

Conclusions

Water usage is closely coupled with energy usage. The current energy efficiency framework for datacenter can be extended to manage water efficiency. A *data center water usage energy metric* is proposed that provides a common framework for

*Corresponding Author: ratnesh.sharma@hp.com

datacenters while capturing the diversity in water availability across geographies and seasons. Sustainable datacenters need to be designed around the local supply constraints of physical resources. Such constraints can be a function of service level agreements associated with IT services or with local administration/utilities.

References

1. The Uptime Institute, "Heat density trends in Data Processing, Computer systems and Telecommunications Equipment", White Paper issued by The Uptime Institute, 2000
2. Patel, C. D, Shah, A., "Cost Model for Planning, Development and Operation of a Data Center", Hewlett-Packard Laboratories Technical Report, HPL-2005-107(R.1), June 9, 2005
3. Sharma, et al., "On building next generation datacenters: Energy Flow in the Information Technology Stack", ACM Compute 2008, Paper ID 129
4. Patel et al., "Energy flow in the information technology Stack: Introducing the coefficient of performance of the ensemble", Proc. ASME Intl. Mech. Eng. Cong. Exp. (IMECE), IMECE2006-14830.
5. King, C., Duncan, I., Webber, M., "Water Demand Projections for Power Generation in Texas", Bureau of Economic Geology, University of Texas at Austin, Apr 2008
6. Goldstein, R., Smith, W., "Water & Sustainability-U.S. Water Consumption for Power Production—The Next Half Century", EPRI Topical Report, March 2002
7. Trask, M., "Water-Energy Relationship", Integrated Energy Policy Report 2005, California Energy Commission.

*Corresponding Author: ratnesh.sharma@hp.com

APPENDIX A:

Ref. Table 1

DAY 2		DAY 1	
Direct Water Consumption (gallons)	Energy Consumption (kWh)	Direct Water Consumption (gallons)	Energy Consumption (kWh)
19980	34240	23010	32850

Ref. definition of water usage energy metric

$$\omega = \frac{\text{(Power Consumption in direct water usage)} + \text{(Power Consumption in indirect water usage)}}{\text{(Power Consumed by IT equipment)}} \times 10^3$$

<u>Calculation for "day 1</u>":

Direct Water Usage:

Direct Water consumption = 23010 gallons per day

$$\text{Energy Consumption in Water Usage} =$$

$$\left(\text{Energy consumed in distribution and treatment of water}\right) \mathrm{X} \left(\text{Water Usage}\right) \qquad (1)$$

$$\text{Energy consumed in distribution and treatment of water} = 1.3\text{MWh/million gallons}$$

Therefore, energy consumption in direct water usage = 29.9 kWh per day

Power Consumption in direct water usage = 1.25kW

Indirect Water Usage:

Energy Consumption in cooling infrastructure = 32850 kWh

Energy Consumption in IT infrastructure = 79200 kWh

Indirect water usage in generation of 1MWh = 500 gallons (assumes fossil fuel fired power plants)

Therefore, Indirect water usage = 56025 gallons per day

From (1), Energy consumption in indirect water usage = 72.83 kWh per day

Power Consumption in indirect water usage = 3.03 kW

Power consumed by IT Equipment = 3.3MW=3300kW

Datacenter Water Usage Energy Metric= $((1.25+3.03)/3300) \times 10^3$ = **1.30**

<u>Calculation for "day 2</u>":

Direct Water consumption = 19980 gallons per day

Therefore, energy consumption in direct water usage = 25.9 kWh per day

Power Consumption in direct water usage = 1.08kW

Energy Consumption in cooling infrastructure = 34240 kWh

Energy Consumption in IT infrastructure = 79200 kWh

Therefore, Indirect water usage = 113440 gallons per day

From (1), Energy consumption in indirect water usage = 73.73 kWh per day

Power Consumption in indirect water usage =3.07 kW

Power consumed by IT Equipment = 3.3MW=3300kW

Datacenter Water Usage Energy Metric= $((1.08+3.07)/3300) \times 10^3$ = **1.25**

*Corresponding Author: ratnesh.sharma@hp.com