Renewable Energy for Datacenters: Energy Flow in datacenters.

Joel Poualeu
Department of Computer and Information Sciences
The College of St. Scholastica
jpouale2@css.edu

Advised by
Dr. Christopher Stewart
Department of Computer Science and Engineering
The Ohio State University
cstewart@cse.ohio-state.edu

Abstract: While envisaging datacenters going green (100% dependent on clean renewable energy sources), there is need to analyze the required energy consumption in datacenters: the amount of energy flowing at different levels in the datacenter and the type of systems using this energy. The experiment carried out involved doing an inventory of the different systems present in a studied datacenter, finding out their maximum power consumption and carrying out periodic recordings of the amount of current flowing through each screened Power Distribution Unit (PDU) to these systems. These results will later be used to find out the amount of currents flowing through the unscreened PDU to systems connected to these. Servers were the leading types of systems in datacenters both in quantity and in the amount of energy consumed. The energy supplied to servers and all other types of systems by means of the PDU seldom changed; if at all any change occurs, this change is mainly dependent on the time of the day and period of the week. This dependency is based on the amount and type of workload posed on these systems.

I. INTRODUCTION

Datacenters support the 24/7 delivery of web services by ensuring that the servers running such services are continuously powered and function properly. Today, electrical energy is used both to keep these systems on and to ensure their proper ventilation. The functioning of a datacenter could be understood from the structures it constitutes. These structures include uninterruptible power supplies (UPS), panels, racks and power distribution units (PDU).

a. The UPS supports all systems in the datacenter by ensuring continuous power in case of short term power outages and surges.
b. The Panel is a device with switches that control the energy distributed to the racks and the systems it contains.
c. The Rack is a device with multiple mounting slots, each designed to hold a hardware unit secured in place with screws. It sometimes contains a PDU within it (or a slot for one) and a cooling system. They help minimize floor space in the datacenter and ventilate the systems respectively.
d. The PDU: These devices, usually in a rack distribute AC current to multiple systems with varying power requirements.

A vital system found in a datacenter is the ventilation/cooling system. The importance of this system is that the arrangement these structures and all other systems are made around/according to the ventilation system. For instance, racks and the systems within them follow a particular pattern of arrangement in order for them to be properly ventilated. Though the ventilation system might not be seen as one enters the datacenter (it may be set up under tiles), its presence is made obvious by the added cool air and noise it generates. This might suggest that the various systems found in a datacenter produce a significant amount of heat energy even as they consume some energy. A ratio thus exists between the amounts of energy consumed to the amount of energy released by various systems.

The various structures within a datacenter do not only allow a better organization within it but also provide a hierarchal pathway down which electrical energy can be distributed to the different systems being harbored in the datacenter. Each structure and system is thus part of route through which energy flows. The amount of current within this route (which can be referred to as a circuit) varies with the location on the circuit at which the measurement is made, the time of the day and the amount of systems that are parts of this circuit. However, all other things being equal, the amount of current flowing from the PDU of a particular rack seldom changes drastically irrespective of the period of the day.

The goal of my project is to measure down the flow of energy throughout a whole datacenter.

II. RELATED WORKS

In his research on datacenters, Dr. Stewart evaluates the use of renewable energy from wind turbines or solar panels to reduce the dependence of datacenters on costly and less clean energy from the grid. Though renewable energy sources comes with major benefits such as a reduced cost and cleaner environment, their scarcity and quantity are of concern [1]. Because the abundant provision of renewable energy is not controllable, the reliance on such sources for system maintenance expose these systems to greater chances of power failure, which of unacceptable for a datacenter. One of the primary steps toward the resolution of this problem will be the evaluation of the power requirements of a datacenter and how these requirements are met, irrespective of the power source. However, because datacenters are increasingly becoming dynamic – new equipment is added and the configurations of older ones are always modified, thus affecting the energy used for the running of datacenters- it is difficult to determine the actual energy consumption of datacenters based on either size or location. Moreover, the increasing power density in datacenters to reduce cost for space and cabling and the variation in this density makes it difficult to determine exact value for the energy consumption of a typical datacenter in numeric terms [2]. In this part of the research we are evaluating the power distribution in a datacenter at the PDU level as well as the maximum power consumption for each system in a studied datacenter in an attempt to find patterns that may be generalized for most/all datacenters.

III. PROCUDURE

Two datacenters are used for case study: the Ohio State University (OSU) datacenter and the datacenter of the Computer Science and Engineering (CSE) Department at the OSU. I was thus involved with the recording of the amount of current (in Amperes) flowing through each circuit of the datacenter at the level of the screened PDU's. These recordings were done in two way: about 30 PDU recordings were collected per PDU; and PDU readings were recorded at 4 different time periods: morning (between 9:00 and 10:00), afternoon (between 14:00 and 15:00), evening (between 17:00 and 18:00) and night (between 21:30 and 22:30). These recorded values were used to calculate the amount of power delivered by each screened PDU at different times to the various systems plugging into it.

The next part of the task was the identification of all the systems connected to each PDU by their brand, model and serial number. From this information, the maximum power consumption of each system was determined by going through the manuals of the systems and/or contacting their manufacturers.

Subsequently, a chart was made mapping each screened PDU to the total (in watts) of the maximum power consumption of all the system attached to it.

Finally, both the specification of the UPSs and their power readings were taken. Like the PDU recordings, the UPS recording was taken at 4 times of the day and different time of the weeks. In this case, the times of the days considered were morning (9:00 to 10:00), noon (11:00 to 12:00), afternoon (15:00 to 16:00) and night (21:00 and 22:00). This measurement was made on a weekday (Friday) and a weekend day (Saturday). Two methods were used to take down the readings: using a saved log file of a web monitoring system and using a digital camera.

The recording of PDU readings was made easy by the good organization of the datacenter with racks at location defined in a pre-built database. The location of the wanted PDU becomes even easier as the number of records increases. Though monotonous, the collection of such data was thus relatively easy.

Because of the structure of some racks that allow a better arrangement of systems' wires, relating each system's power cord to the PDU to which it connects to may be tedious as these cords interweave with other cord, making it difficult to differentiate between cords. Lengthy cords further facilitate this interweaving and make the process perplexing. Some systems have multiple power cords that may connect to more than one PDU. This makes this process time-consuming.

Most systems in a datacenter are well marked by brand and model for identification purposes. However, there exist systems that lack either one or the other or both of these identifications. Also, some system may not be reported online because they are unpopular or superannuated. These contributed to the uneasiness in getting technical specifications of some system. Inaccurate information are also very likely due to the publication of wrong specifications and/or the interchangeable and misuse of terms such as power supply, input current/voltage and power consumption in manuals and other publications.

Given some power readings at various components levels and behavior models of how each component works, we are calculating the power consumption of all components. We collect these energy usage traces and will use all the data traces in a simulator to solve the problems relating to the flow of renewable energy into datacenters.

IV. RESULTS

Patterns exist at various levels of the collected data:

- Besides servers, many other systems are part of the energy-consuming devices in a datacenter. However, servers are more abundant and have higher aggregate power consumption wattage.

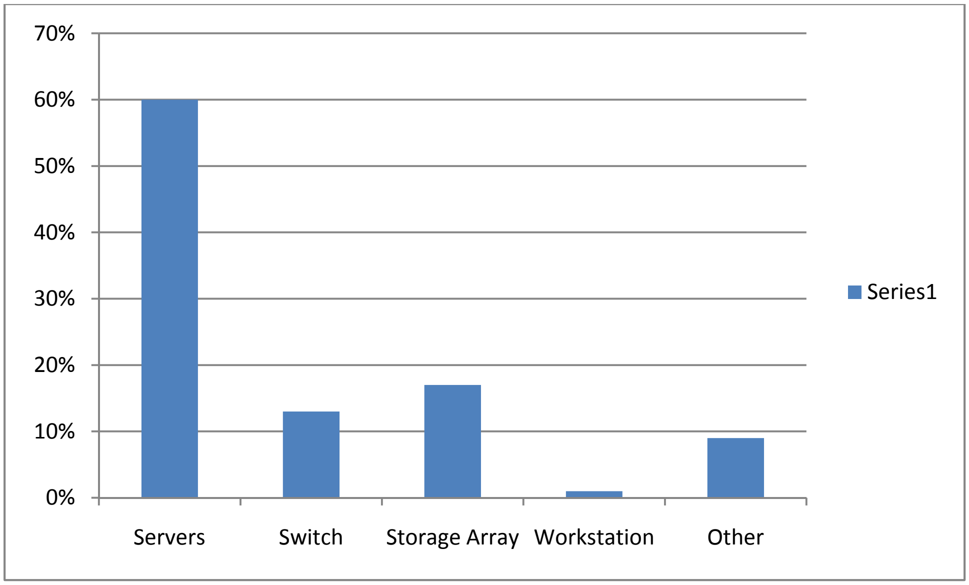

Figure 1.1- Quantity of Each Type of System in a Datacenter

Figure 1.1 shows the percentage of some types of systems in one of the studied datacenters. The types of system present in the datacenter include servers, switches, storage arrays, workstations, firewalls, routers, computers, disk arrays and blade enclosures. As shown in Figure 1.1, the amount of servers in a datacenter is greater than that of any other system. As mentioned earlier, datacenters aim at ensuring a continuous supply of network-based services. However, for it to be effective in this duty, not only must it ensure the powering of the systems launching these services, but also the maintenance of these systems. Though the other system may not launch the network applications as do servers, they compliment to the functioning of the servers and of the datacenter as a whole by providing services like information back-up, information security, or other managerial services for the applications stored in the servers.

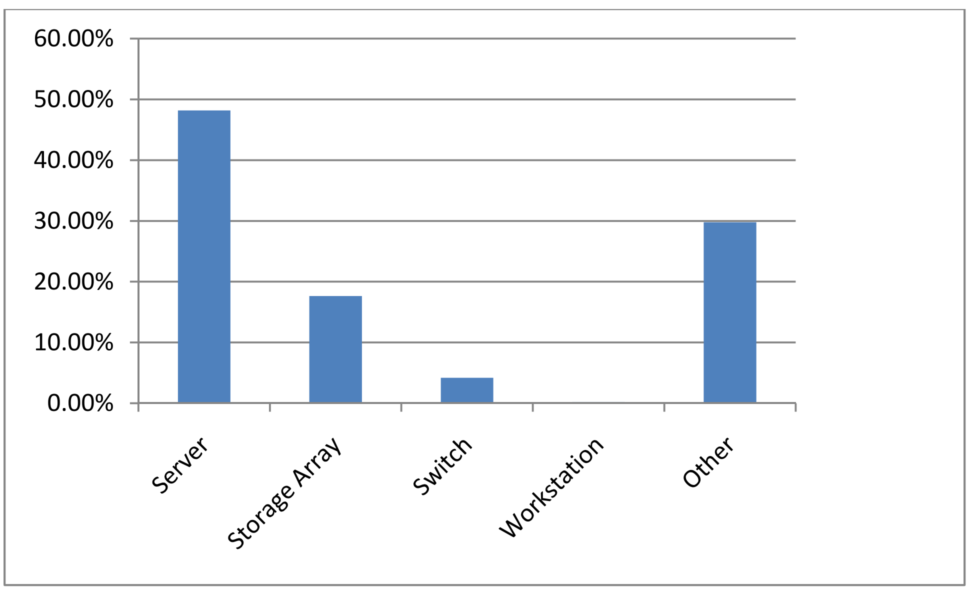

Figure 1.2-Power Consumption of each Type of System in a Datacenter

There also exist variations in the percentage of the total power consumed by each type of system in a datacenter. Because of their out-numbering and the type/amount of workload they handle, servers have a greater percentage consumption of the total energy used by datacenters- as shown in Figure 1.2 Because each other individual system type just compliments to the functioning and maintenance of servers, hence their individual lower energy consumption.

- PDU readings remained more or less constant throughout the observation and for the few whose readings changed, the change occurs at late hours of the evening (around 17:00)

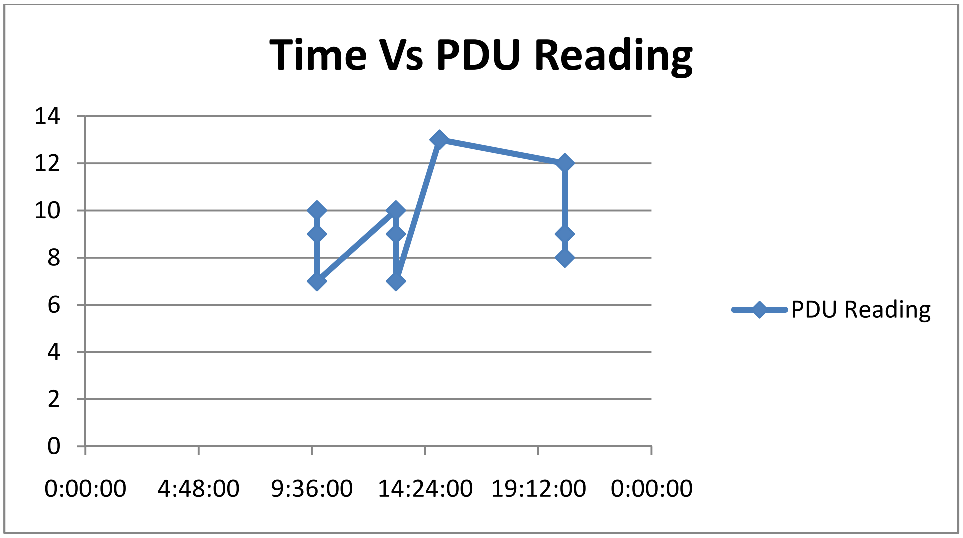

Figure 2.1

In Figure 2.1 above, the PDU reading alternate between 3 constant values- the duration of each value vary in seconds. 3 values could therefore be read each minute. This pattern continues all day till the evening (17:00). At this time, the all-day pattern is broke; a single reading can be gotten per minute and this value is the higher than the previous values. The alternating 3 values pattern starts again at night, but with different values. This cycle repeats itself the next day.

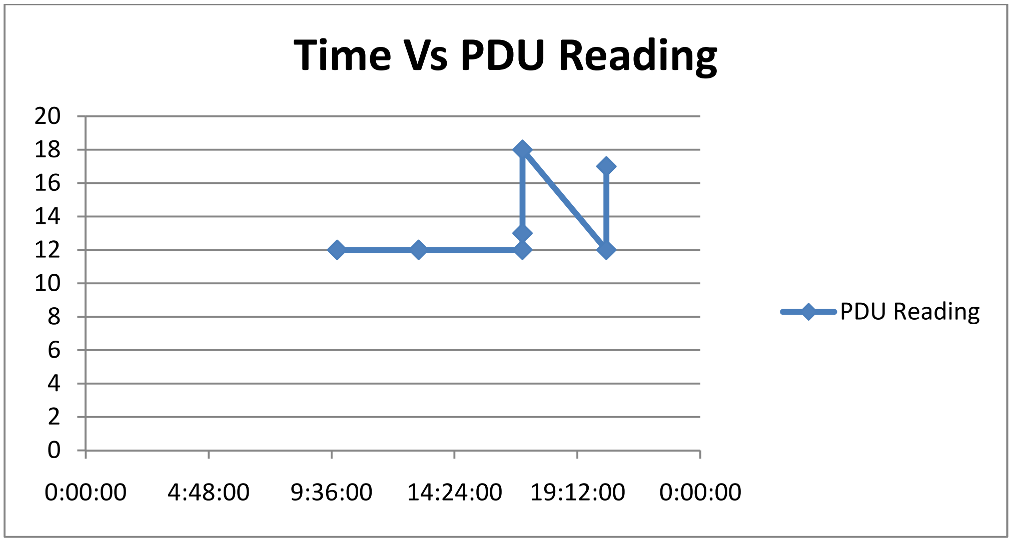

Figure 2.2

In Figure 2.2, the PDU reads a single constant value in the morning and afternoon. In the evening, the PDU reads 3 alternative values; two other values added to the all-day value. Finally, at night, only two values can be gotten per PDU, per minute. All through the day, the PDU has a value which reoccurs irrespective on the time of the day. Other values appear or disappear depending on the time of the day.

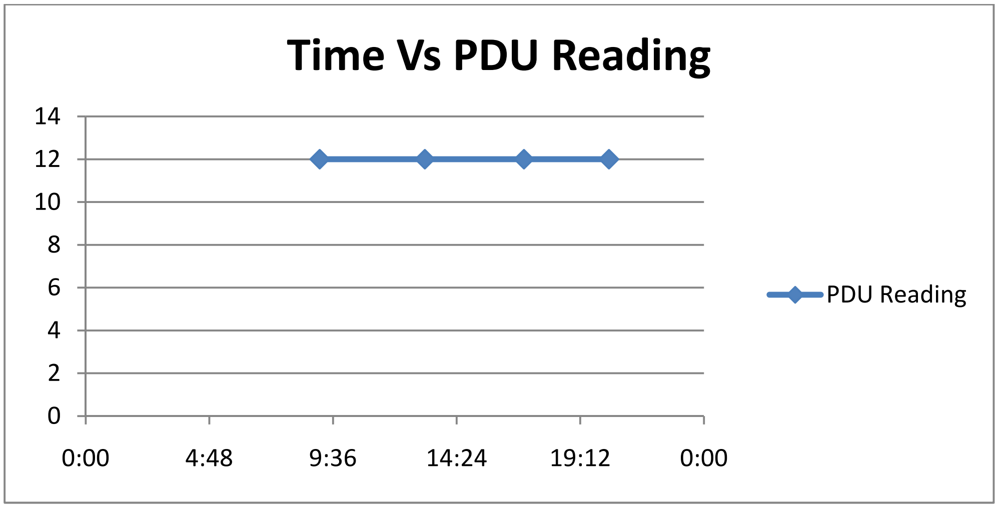

Figure 2.3

In Figure 2.3, the time of the day is not a function of the PDU reading. This remains constant irrespective of time. This pattern is observed by most screened PDU in the datacenter.

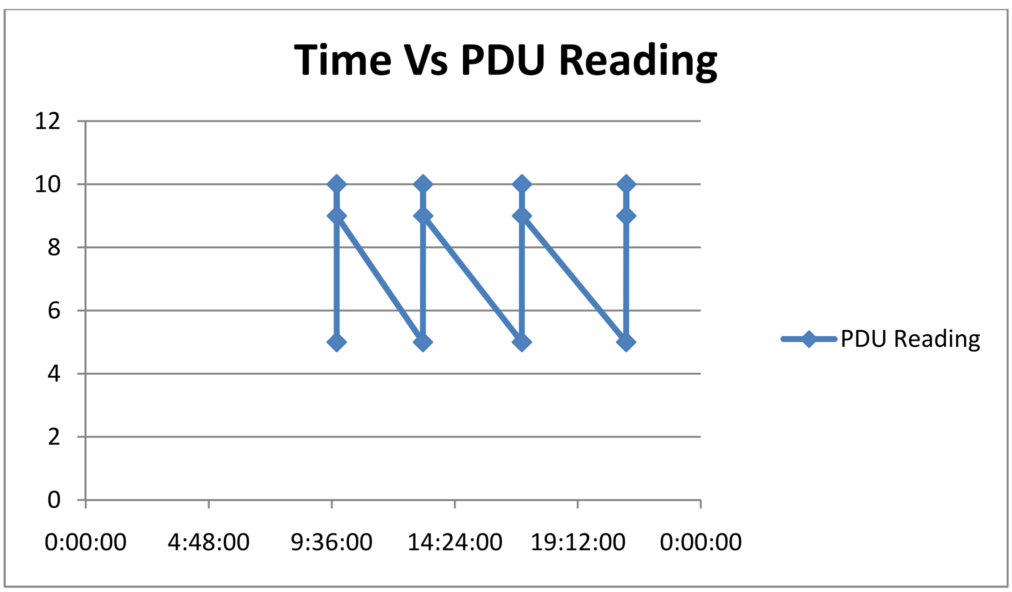

Figure 2.4

As with the PDU reading demonstrated in Figure 2.3, that of 2.4 is not a function of time. The PDU reads a constant series of alternating values all day. This pattern is observed by a greater number of PDU's in a datacenter.

V. CONCLUSION

These patterns are very probable to exist in many datacenters. The typical operation of most systems may require the drawing of a definite amount of current, that is not likely to change unless during adverse conditions (that may be very rare); hence, the generally constant reading per PDU. However, type and amount of workload posed on certain systems may place unusual demands for power consumption. The type of workload varies with the time of the day. Working hours of the day are more likely to have a peculiar demand for power, hence the difference in the in the PDU reading with time.

Datacenters operate at the basis of a constant energy flow through it. Variations may exist in the number and type of systems in it, the type of energy flowing through it or even the amount of energy flowing through it at any given time. However, for it to be effective in uninterruptedly providing network-based services, there is need of a quantitative and constant supply of energy and a well structured pathway that allows the steady and continuous flow of energy to the various systems in the datacenter, through these systems and out of these systems.

VI. FUTURE WORKS

VII. ACKNOWLEDGEMENT

We are thanking the Computer Science and Engineering department and the Information Technology department at The Ohio State University for giving us access into their datacenters and network throughout the research.

VIII. REFERENCES

[1] C. Stewart and K. Shen. *Some Joules Are More Precious Than Others: Managing Renewable Energy in Datacenters*. Workshop on *Power Aware Computing and Systems (HotPower) Big Sky*, MT, October 2009

[2] J. Moore, J. Chase, K. Farkas, and P. Ranganathan. *Data Center Workload Monitoring, Analysis, and Emulation*, 2005.