

APPLICATION NOTE

Addressing/Calculating Efficiency Challenges in the Data Center

Power Conversion Plays a Key Role

Introduction

With the digital expansion over the past few decades, increasingly demanding requirements for data centers to store, transmit, and manage data of all sorts are reaching a critical level. Companies are feeling the crunch as data centers are filling up quickly with information technology (IT) equipment, real estate in the data center remains at a premium, and power demands and energy prices are skyrocketing. Several issues are confronting data centers today. First is the total cost of ownership (TCO) — which consists of the IT equipment as well as the data center infrastructure used to support operations. The infrastructure includes the power distribution, cooling, cabling, and several other components. Second, is the fact that data center real estate is at a premium. Each inch of space within a rack enclosure and the data hall needs to perform at maximum efficiency, including computing performance, power, and cooling. Efficient use of power is also a key ingredient in achieving green and sustainable operations along the path to carbon neutrality of data centers.

Background

The number of servers per rack enclosure is increasing dramatically and to help with the lack of real estate in data centers, servers are becoming smaller with increased processing capabilities. These servers are also becoming more power dense. In 2010, the average rack enclosure housed 20 servers. This number is expected to be routinely exceeding 40 servers today. At the same time, because servers pack so much computing power into a small space, the power and cooling challenges intensify. Therefore, on one hand, these dense IT solutions are helping alleviate data center real estate issues, yet simultaneously, they are also contributing to the power and cooling challenges.

Data Center Power and Cooling

Power and cooling costs skyrocket

For a data center to be considered “efficient” the highest proportion possible of the power used by the total facility should be used by the Information Technology (IT) equipment.

The term Power Usage Effectiveness (PUE) has been coined, with the definition of:

$$\text{PUE} = \text{Total Facility Power} / \text{IT Power}$$

The most desirable value of PUE is 1.0, but values from 1.2 to 2 are currently typical. Having a facility that uses 100kW of total power of which 80kW is used to power your IT equipment, would generate a PUE of 1.25.

Similarly, Data Center infrastructure Efficiency (DCiE) is the inverse of PUE

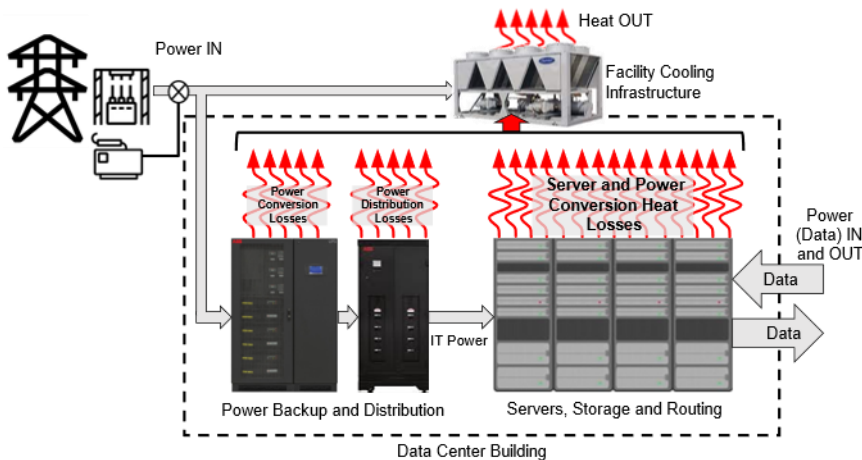
$$\text{DCiE} = (\text{IT Power} / \text{Total Facility Power}) * 100\%$$

APPLICATION NOTE

The ideal DCiE is 100%. Having that same facility that uses 100kW of total power of which 80kW is used to power your IT equipment, would generate a DCiE of 80%

Where does all the power go?

Power costs for the average data center come from the utility grid, and primarily consist of two parts; a) powering the IT equipment (IT Power above) and b) powering the rest of the facility infrastructure, including cooling the facility (Total Power – IT Power).



The data center shown in Figure 1 shows how power from the grid powering the IT equipment (servers, storage and routing) through a UPS, and removing the waste heat through the cooling infrastructure, which itself, uses a large amount of power .

Figure 1 – Where the power goes in a traditional UPS style Data Center

In simple terms, if the IT power is 100KW, almost all that power will be dissipated as heat, with very little actual power being transferred with the movement of data in and out of the facility. This means that almost 100KW of heat from the IT equipment, plus the heat given off by the power train shown, must be removed from the facility by the cooling apparatus. Cooling apparatus can be active (Air conditioning) or passive (if the climatic conditions are conducive) or a mixture of the two. Looking at active cooling, such as traditional refrigerant air conditioners, it takes a certain amount of energy to remove heat, leading to an efficiency rating of air-conditioner systems. The simplest of these is the Coefficient of Performance (COP). The larger the COP, the more efficient the heat removal.

$COP = \text{Cooling Capacity (W)} / \text{Power Consumption (W)}$

COP Example: If a system requires 2KW of utility power to remove 5KW of heat the COP would be:

$$COP = 5KW / 2KW = 2.5$$

Typically, air conditioning systems run with COP in the 2-4 range, although high performance units may be slightly higher.

In our example, if we ignore the heat in the rest of the power train, simply removing the heat from 100KW of IT load will require an additional 25KW of power if the air conditioning has a COP of 4. This illustrates how significant the cost of cooling can be in the operating costs of a data center.

In practice we cannot ignore the heat developed by the system's power train, because despite best efforts the efficiency of this infrastructure is not 100%. Let's look at this portion of the heat load.

See footnotes on page x

APPLICATION NOTE

Power Train Efficiency and Heat Generation

Critical Power

Data centers are critical infrastructure and require continuity of operation. This requires a power train that can maintain operation even if the utility power supply should fail. This is usually achieved with an Uninterruptible Power Supply (UPS). There are two types of UPSs, one that supplies alternating current (AC) and one that supplies direct current (DC). Figure 2 shows both types for comparison. Both types use a battery or other energy storage device to maintain output during a utility outage. The UPS may or may not be supplemented by a local generator for extended outage management. Figure 2 compares the power architectures used in a typical data center, using either AC or DC UPSs.




Figure 2 – Data center power architecture with AC and DC UPSs and power conversion steps

In the AC UPS architecture, incoming utility power is converted first to 480VAC, then to 540VDC for intermediate (battery) storage, then back to 480VAC. The power distribution unit (PDU) then converts down to 208VAC; the power supply in the server converts down to 12VDC for use within the server itself. This example illustrates a total of five separate conversions, each of which incurs its own losses due to the inefficiency of the various electronics.

The DC UPS has a similar power train, except that once the power is converted to DC it is not converted back to AC, since the IT equipment universally requires DC for the electronics operation. The equivalent number of conversions in this architecture is reduced to 3.

In the traditional telecom data center environment, the DC voltage (DC1) is 48V, but other voltages have also been used, including 380VDC. The use of higher DC voltages reduces the distribution losses in what may be long (100's of feet) distances. If losses are incurred due to distance, one can decrease the cable losses (decrease resistance) or decrease the distance by moving the conversion step to lower voltage closer to the load. This is the subject of another paper covering the OmniOn Power's Edge Distributed Data Center Power Architecture.

Figure 3 shows typical efficiency numbers for AC and DC UPS power train elements and the equivalent power (heat) releases.




Figure 3 – Data center power architecture with AC and DC UPSs with typical power conversion efficiencies

Both systems in Figure 3 power a 100KW IT load. Using typical efficiency data enables the calculation of the power required from the utility and the power lost as heat at each conversion point. These calculations do not take into account the losses in cables and distribution, which are assumed to be small and similar in both cases. The total losses in the AC power train are almost double those in the DC case.

APPLICATION NOTE

The table in Figure 4 shows how the power train in Figure 3 gives off heat because of the inefficient conversion of power. In the AC case the power train generates 11.8KW vs the DC case with 6.3KW. The total heat dissipated in each case is 112KW vs 106KW respectively. All heat dissipated in the facility must be removed by the cooling system, and as previously shown, a cooling system with a COP of 4 will require an additional 25% of that heat power in energy to remove the heat.

Figure 5 shows how many modern servers are equipped with one (or two for redundant operation) plug-in power supplies to convert incoming power to the requisite DC voltages. In many cases the manufacturer has both AC/DC and DC/DC versions of these power supplies, simplifying the use of the equipment in both AC and DC powered architectures.

	Efficiency	AC Power	Heat KW		Efficiency	DC Power	Heat KW
Input Power		111.79		Input Power		106.26	
Utility Transformer	99%		1.12	Utility Transformer	99%		1.06
		110.68				105.20	
AC to DC	98%		2.21	AC to DC	98%		2.10
Battery		108.46		Battery		103.09	
DC to AC	98%		2.17	Server Power Supply	97%		3.09
		106.29		IT Power		100	100
PDU	98%		2.13				
		104.17				Total	106.26
Server Power Supply	96%		4.17				
IT Power		100	100				
		Total	111.79				

Figure 4 – AC to DC Power Supply vs DC to DC Power Supply




Figure 5 – Server Power Supplies

Power Conversion Efficiency

AC to DC

The AC to DC conversion of power involves several key parameters, however, we are primarily looking at the efficiency of that conversion in this context. Both the AC and DC UPS power trains compared in Figure 3 utilize an AC to DC (rectifier) step.

To achieve the highest levels of reliability demanded by today's data center operations, redundancy is often utilized to ensure operations even if a single module fails. The use of parallel power modules enables building large power systems while maintaining ease of maintenance with plug-in replaceable modules. Using more modules than strictly required gives the redundancy level required.

Rectifier Size:	4	KW	
Load:	20	KW	
Redundancy	Qty	System Capacity	Utilization
N+0	5	20 KW	100%
N+1	6	24 KW	83%
N+2	7	28 KW	71%
N+3	8	32 KW	63%
N+N	10	40 KW	50%

Figure 6 – Power Supply Redundancy

APPLICATION NOTE

Redundancy Implications

When redundant power modules are provisioned in the power supply system, the overall system must operate at less than 100% of its' total capacity. The chart in Figure 6 shows how the utilization decreases as the redundancy increases. A fully redundant system, with two complete power trains - each capable of supporting the load - will only operate at 50% of its full capacity, at most.

Efficiency Implications

Operation of a system at 50% capacity has implications on the efficiency as can be seen in Figure 7, which illustrates how efficiency may vary over variations in load, shown as percentage of total capacity. The red line is an older generation rectifier with the peak efficiency occurring at very high loading, around 95%. Modern rectifiers have been improved to achieve higher efficiency, close to 98%, but also to achieve that efficiency at 50% utilization.

IT equipment is subject to large variations in utilization as processing demands vary over time. In addition to the utilization (and efficiency) decreases due to redundancy, there will be temporal variations in load power usage, which will also decrease efficiency.

As can be seen in the black efficiency curve in Figure 7, if the utilization can be kept above 20%, there will be minimal degradation in power supply efficiency, a significant improvement over the older red line performance.

What is not quite so obvious is that the lack of efficiency, (i.e., the amount below 100%) manifests itself as heat release . The graph and table in Figure 8 show the equivalent heat released for the power supplies in Figure 7. These values are calculated for a power supply capacity of 10.8KW. Intuitively, one would expect the heat released to increase as the loading increases. What is not so intuitive is that the heat released by the red rectifier is 3.5X the heat released by the black rectifier, even though the efficiency difference is only 8% (at full load). At 50% load the difference is 4.4X for a 9.8% efficiency difference.

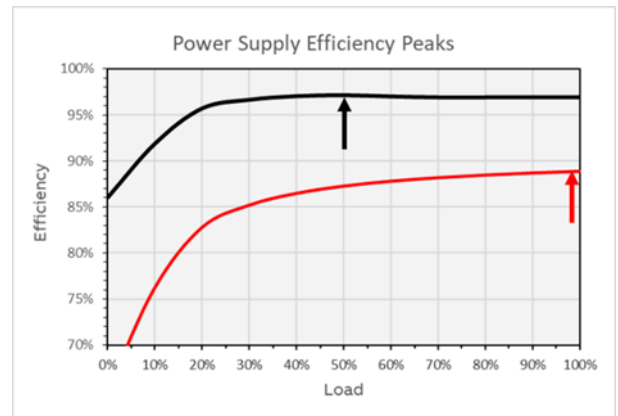
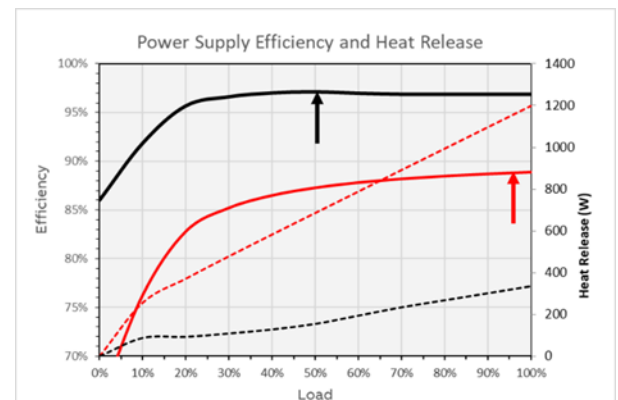


Figure 7 – Power Supply Efficiency



Load	Efficiency	Heat Release	Efficiency	Heat Release
0%	65.0%	0	86.0%	0.0
10%	76.2%	257	91.9%	87.8
20%	82.8%	372	95.7%	92.9
30%	85.2%	479	96.6%	109.1
40%	86.5%	583	97.0%	128.2
50%	87.3%	686	97.1%	154.8
60%	87.8%	789	97.0%	194.4
70%	88.2%	892	96.9%	234.4
80%	88.5%	994	96.9%	267.8
90%	88.7%	1097	96.9%	301.3
100%	88.9%	1198	96.9%	334.8

Figure 8– Power Supply Efficiency and Heat Release.

APPLICATION NOTE

Supporting Server Power Supply Performance Information

Looking at a popular single rack unit (1RU) server datasheet the maximum heat release is calculated based on the power supply rating. The data sheet states that when a server with an AC power supply unit (PSU) rating of 1,400W is used there can be a total of 5,459BTU/Hr heat release. $5,459\text{BTU}/\text{Hr} = 1,599\text{Watts}$. If 1400w is from the PSU output, through the processor etc., then 199W is from the PSU. If the PSU is producing 1,400W of output and dissipating 199W of heat the efficiency of the PSU is ~88%

When the same 1RU Server is equipped with a DC PSU rated at 1,100W the data sheet shows 4,266BTU/Hr of heat release. $4,266\text{BTU}/\text{Hr} = 1,250\text{Watts}$. If 1,100w is from the PSU output, through the processor etc., then 150W is from the PSU. If the PSU is producing 1,100W of output and dissipating 150W of heat the efficiency of the PSU is ~88%

Both specifications confirm the estimation made earlier that virtually all of the power input to the IT equipment is converted to heat with a minimal amount of power being transferred out of the facility in the data transfers.

Summary

Based on the data that we have analyzed; there are inefficiencies in each component of the data center power train. These inefficiencies lead to greater facility power consumption and recurring electricity costs that can substantially outweigh the original purchase price of the IT equipment. Furthermore, heat generated by inefficient power conversion results in additional costs to cool the facility. Low efficiency results in a larger carbon footprint for the facility.

Our analysis shows, that in terms of efficiency, there are several areas for improvement. All power conversion stages offer opportunities for improvement, either by elimination or performance improvement. The use of DC-powered equipment offers the opportunity to eliminate several power conversion stages. OmniOn Power offers many high efficiency AC to DC and DC to DC front-end power supplies. OmniOn Power has concentrated on these products to increase power densities and true operational efficiencies. This approach enables data centers to leverage existing power architectures while experiencing substantial reduction in their TCO through reduced power consumption, physical size, and cooling costs.

Data from OmniOn Power high-efficiency power supplies has been used as examples to demonstrate the significant gains that can be achieved in a data center by addressing components of the total power train.

While the power train efficiency is critical to the overall performance of the datacenter and the reduction of the heat load in the facility and its attendant cost to remove, the lion's share of power consumption is actually in the IT equipment. And the actual data processing portion of the facility offers the biggest opportunity for power reduction through processor efficiency.

OmniOn Power Inc.
601 Shiloh Rd.
Plano, TX USA

omnionpower.com

We reserve the right to make technical changes or modify the contents of this document without prior notice. OmniOn Power does not accept any responsibility for potential errors or possible lack of information in this document.

Rev 3.0

We reserve all rights in this document and in the subject matter and illustrations contained therein. Any reproduction, disclosure to third parties or utilization of its contents – in whole or in parts – is forbidden without prior written consent of OmniOn Power. Copyright© 2023 OmniOn Power Inc. All rights reserved.