DISTRIBUTED DC CONCEPT FOR HIGH DENSITY DATA CENTERS APPLICATIONS

Kfir L. Godrich
Director of Technology Development
EYP Mission Critical Facilities, Inc.

ABSTRACT

High density Data Center topology is one of the biggest challenges of our industry. With power densities rising from 150 W/sf to over 1,000 W/sf, the Data Center infrastructure is definitely changing – the Data Center compound is about 1.2% in the US load. In addition, the efficiency factor becomes more important in every decision, the reason being the fact that operational costs of the Data Centers have a 30-50% compound of energy cost!

On the servers side, regular servers of blades are manufactured more and more without the traditional power supply input stage, directly with –48VDC bus or soon we might be able to see racks with direct feed of DC power at higher voltages. The efficiency is improved by about 15%-30% at the local server level, with promises to save even more for the end-to-end system. The end-to-end approach is the key that owners and operators are starting to implement. This "drawback" is the simultaneous increase in the servers per cabinet density, therefore increasing the density per rack.

We will present a different design direction – the $(DC)^2$ TM (pronounce "DC square"). This unique approach suits multiple Data Centers spaces with backup from an array of tri-dimensional (3D) optional high density sources/storage devices. Relevant cooling options as well as other aspects will also be discussed. The design includes references to the architectural design of the Data Center, such that the Operational Availability is improved at the whole level application and not for several sub-systems only. It is important to emphasize that different proposed storage dimensions are based on different technologies adding another best practice besides the more traditional redundancy levels implemented, as dual path or fault tolerance. We'll present design examples with modules up to 4MW distributed with up to 600VDC systems including considerations and discussions regarding equipment and protection.

INTRODUCTION

The opening of the current High-Density (HD) trend was made possible by the latest technological breakthrough in silicon "nanometer" improvements by a factor of about 30% - thus creating the time for the multi-core technology adopted by AMD, Intel and others. In conjunction with the capabilities boost in the processors, there is a change in the configuration of the servers. The memory of the new servers is upgraded with buffered DIMMS — taking a 30% stake out of the needed power.

Without a doubt, energy efficiency is one of today's key issues [6]. Here we will try to examine the HD power environment from the resolution of processors, servers and racks, to the range of end-to-end AC and DC sub-optimal solutions with the goal of improving power efficiency. We will look at additional performances and how their efficiencies are impacted by the various solutions: including computational performance, cost performance and space performance.

DATA CENTER TRENDS

The data center market changed in the last two years when, besides the traditional availability/reliability requirement, additional inputs entered the game and created an interesting roadmap for the next several years – including: HD, energy efficiency, "green" energy and additional performance related variables.

Everything starts from the basic Moore's law that generally claims you can squeeze twice as many transistors onto an integrated circuit every 24 months. See hereafter a curve describing this development over the years.

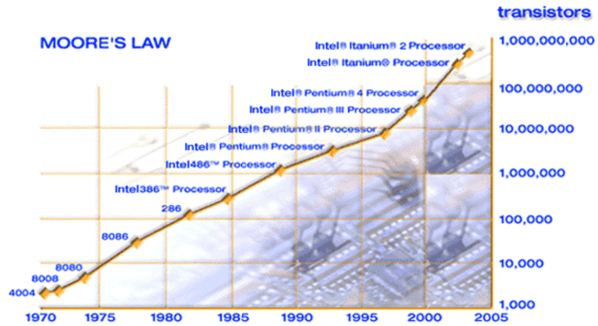

Figure 1. Moore's Low development over the years. (Source – Intel)

This can be seen in Fig.1 describing the situation until the last year in the CPUs world. When referring to the performance of the CPUs power in the various applications of the market, a rule of thumb describing the utilization of the CPUs is that in average, 80% of the time the applications use 20% of the CPUs potential. A lot is happening in trying to improve this misperformance and one of the most popular is the "virtualization" trend.

The main idea here (see [1]) is that in order to virtualize a set of diverse concrete resources you have to access them through a uniform interface that, from the user's perspective, enables them to behave as one "virtual" resource that can be shared by diverse users with varying degrees of dynamic behavior. The practical way of implementing that is to create a "virtualization" pod in the data center, usually containing high density blade servers. The applications running in the regular pods will use protocols that will create maximum loading of the CPUs in the virtualization pod (by means of CPU utilization – power and cooling in this case will be also at the highest utilization mode, close to nameplate). Virtualization creates a "single system illusion" for certain "users" while there are no illusions about the fact that this process will create relatively high dynamic behaviour from the load point of view in the data center's environment.

In [2] Ray Kurzweil describes the pace of the metrics changes in his "futuristic" approach by looking at Gene's law, that his basic trend line follows that of Moore's Law in that DSP power dissipation per MIPS halves every 18 months (see Fig. 2).

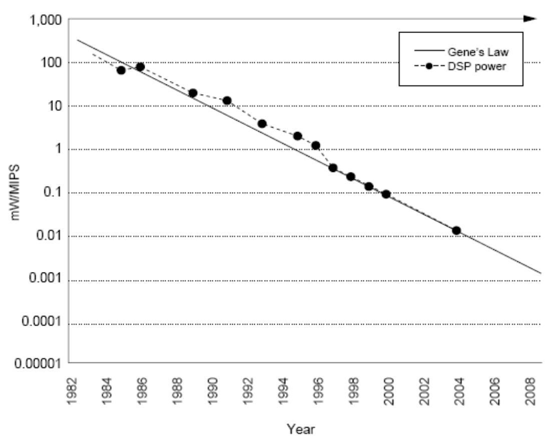

Figure 2. Gene's Low development over the years. (Source – Singularity)

So we are going forward and increasing the density of the transistors on the CPUs and also grouping CPUs in dual/quad/multicore platforms in order to increase performance (see Fig. 3).

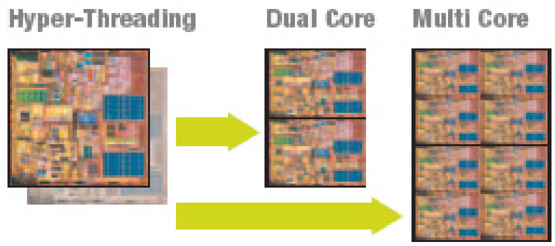

Figure 3. Multiple Core Configurations.

The metrics are changing due to over-performance as well as under-performance (see also [3]). Looking at the last 10 years in the processing world Fig. 4 summarizes the achievements as follows:

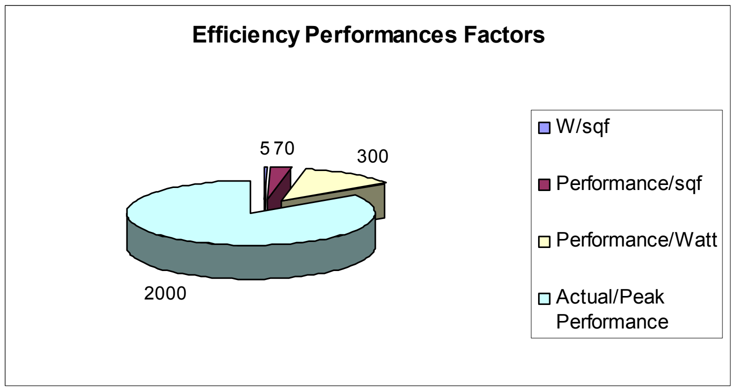

Figure 4. Efficiency Performance Factors Over the Past 10 Years.

As can be seen from that figure, a lot has been done to improve the processing performance, very little has been done to improve the power/space efficiency performance.

One may ask, what is the limitation in this pace, if any? Well, it looks like this limitation has to be the thermal limitation. In today's technologies we are about 50W away from the 200W air cooling limitation of chips. While this will occur, the next step will be to move over direct coolants technologies at the chip level, direction which some of the semiconductors manufacturers started to seriously consider and implement. As a note we should mention that this pace has nothing to do with current solutions at the rack level, as water cooling for example, that is still using air inside the rack for the direct cooling of the servers/blades.

Evolving from the micro to the macro, the rack level solutions available today are impressive for new "green fields" as well as in existing data centers looking for blading of creating high density pods.

This range of solutions, under the category of supplemental cooling has a variety of possibilities most of them fitted for the hot aisle – cold aisle arrangements, pressurizing the cold aisle from the top of the rack, the ceiling between the racks of in-row solutions.

As mention before other then air cooling solutions are being considered including water, CO_2 and other coolants that offer reduced energy consumption. Rejecting heat given off by servers to water (such as in a water-cooled cabinet) versus air (such as in a traditional air cooled data center) is more efficient, simply by the fact that water, per unit volume, has greater ability than air to remove heat. However, there are other issues related to reliability and functionality of the rack that may present challenges in implementing water cooled cabinets into a High-Performance Computing (HPC) environment. The limitations of the water cooling solutions are not only psychological (quite a lot of data center operators are not ready yet to make this step after for years they are trying to keep water out of the white floor space) but also an economical barrier – water cooling solutions are an exponential increase in the infrastructure costs as well as the cost of the racks (up to $33,000 per rack, not installed). A High Density row of 5 water cooled racks of blades can get a sticker of $1M, not including installation and the infrastructure share.

One technology that includes rack level power and cooling approach we find interesting is the Rackable solution. This solution doubles the density of blades potential for a rack to around 190 blades, by using both sides of the rack. This solution has a -48VDC option that makes the total efficiency even higher. We believe that if the price of the -48VDC power supply option from the servers manufacturers will go down to the prices of the AC inputs, which it should this solution will be even more popular.

There are several things we learn from implementing efficiency approaches to High Density and HPC environment:

a. There is no such a thing as one solution fits all.
b. The geography of the KWh changes the relevancy between various solutions – $0.05 a KWh in Washington State versus $0.29 a KWh in Manhattan is quite a wide spectrum.
c. Local solutions are low impact only, end-to-end solutions are the right way to go.

A positive example of HPC efficiency success is the IBM BlueGene/L. This 64 racks machine has an impressive 360TF/s capacity running over 2,500 SF raised floor with 1.5MW of power, 20KW/rack – about 500 W/SF. This amazing machine runs air cooled with individual DC fans next to each rack over regular raised floor (see Fig. 5 hereafter).

Figure 5. BlueGene/L install example. (Source - IBM)

The HPC platforms targeted for the year 2010 will perform around 1 PF/s with power densities of up to 1.2KW/SF – a real engineering challenge.

The latest solution in the cooling technologies for Data Centers is introduced by HP and is called DSC (Dynamic Smart Cooling) - it is beyond the scope of this paper to get into the details of that approach, but it is considered a breakthrough in the energy efficiency optimization (see also [6]).

THE $(DC)^2$ TM CONCEPT

A parallel approach to the DSC and related to end-to-end efficiency is the $(DC)^{2\,TM}$ Concept as presented in figure 6 and described in [4, 5]. This concept has a DC data center configuration as the main approach. The idea was to rectify MV directly to the distribution voltage window of 500-600VDC in the Data Centre and to use DC/DC converters at the row level to supply the High Density racks with -48VDC. When compared to the regular AC high available data centers, the MTBF is improved by a 25% factor and the probability of failure for a period of 5 years is improved by about 17%.

Looking in the AC data center distribution versus the DC data center distribution efficiencies, our analysis yielded the followings efficiency ranges along with their variations (computations in the AC case included the chain between the UPS input down to the servers inputs, while in the DC case from the RC inputs down to the servers inputs):

a. $82.0\% < \eta_{AC} < 88.0\%$
b. $84.9\% < \eta_{DC} < 91.2\%$

While implementing the end-to-end approach with the related cooling solution, the efficiency calculation yields the followings***:***

a. AC system - 59.5%
b. $(DC)^{2\,TM}$ concept – 75.2%

There is a 21% improvement in the total energy consumption in favour of the $(DC)^{2\,TM}$ concept.

About the same improvement (20%) is seen in the space efficiency when compared between the concepts. Fist High Density data centers are deployed these days up to power capacities of 50MW each – results to follow.

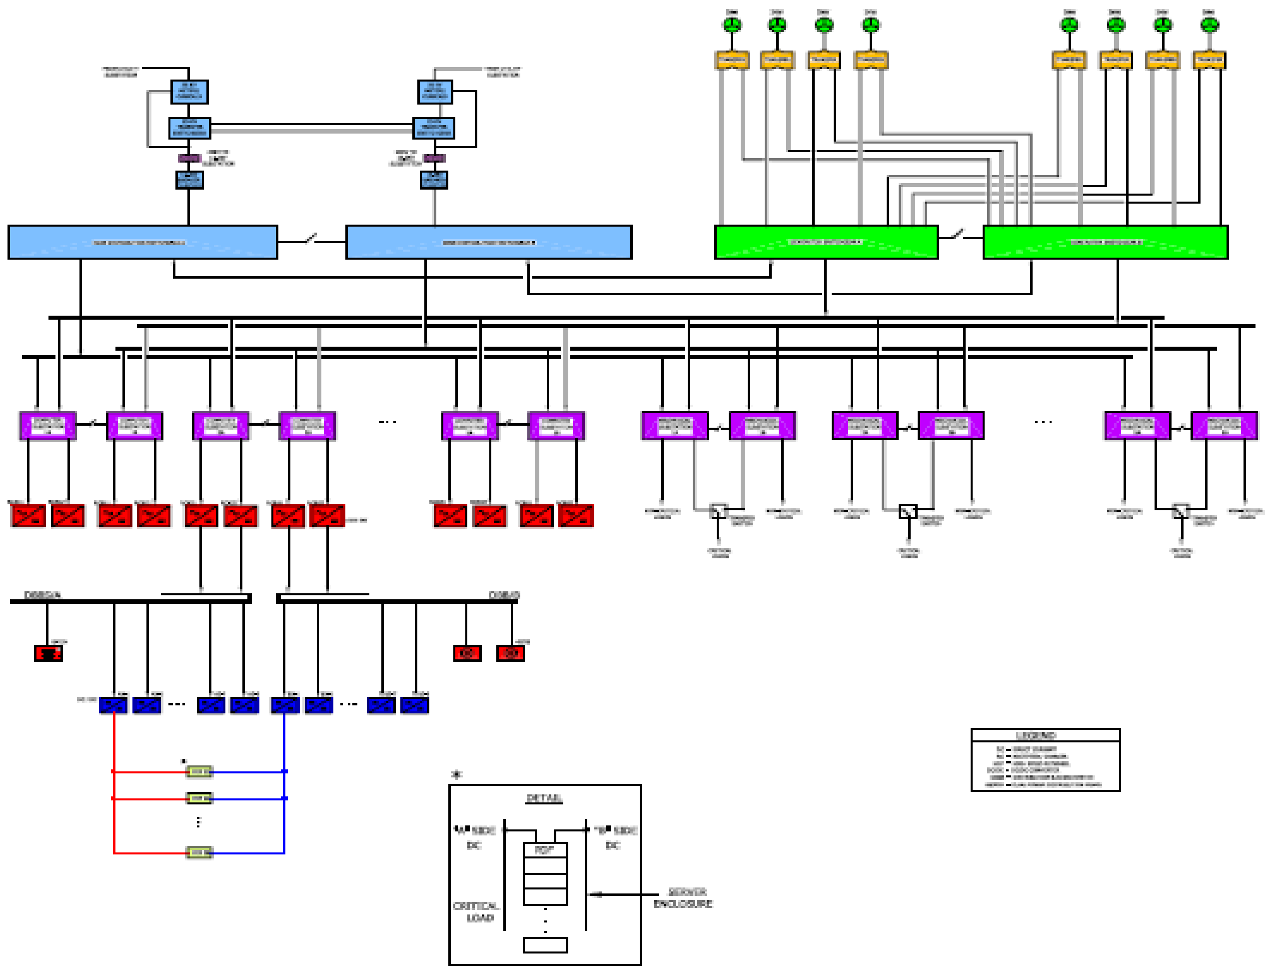

Figure 6. DC^{2}TM Data Center Concept.

CONCLUSIONS

We presented herein a review of various concepts for improving performances in data center environments to optimize operation costs. The DC^{2}TM is a distributed DC solution that is currently a good alternative to AC Data Centers with improved reliability and energy efficiency performance. We envision in the future improved DC distributed solution with higher voltage and fewer components, once some issues of safety, protection will be solved (not to mention the psychological factor that we're dealing today while spreading the DC concepts). Taking into account that a large part of these costs are hidden in the energy bill, this being between 30-50% of the total operations cost, a careful attention should be given to the energy efficiency as one of the most important performances of a data center. This approach must be applied end-to-end. We can summarize that in our opinion the future of data center and HPC is going to high density, efficient and "green" as the alternative energy solution will improve their performance.

REFERENCES

[1] Wunder, D.: How to get started, landmarks on the way, presented at Marist On Demand Business conference, IBM 2005

[2] Kurzweil, R.,: The Singularity Is Near, Viking, 2006

[3] Gross P.: Needed: New Metrics, Energy User News, 08/2002

[4] Gross P., Godrich K. L.: Total DC Integrated Data Centers, Intelec, 2005

[5] ***(DC)*** $^{2\,TM}$ **-** US Patent pending 2005

[6] The Economist, Going Green, March 2007