Server Technology
A brand of legrand

stay powered. be supported. get ahead.

Density in the Data Center 2018 Update

Written by Marc Cram, CDCD

Executive Summary

The original version of this paper was released in 2014. Four years on in 2018, the original premise of the paper has withstood the test of time. Density in the data center continues to climb. Meanwhile, the solutions for addressing density have changed to meet the new requirements coming from enterprise and hyperscale data center operators.

Density in the datacenter refers to the electric power consumed per square foot of floor space (or unit volume of space), number of servers, and the cooling systems load. A number of factors in today's market are driving the datacenter operator to increase the density and the operational loading of the compute infrastructure. Among these factors are the desire to reduce both capital and operational expenditures, improve energy efficiency (lower PUE), avoidance of new build-out, reduction of latency times, and better management of network traffic flow.

Historically, the aggregate power load of the datacenter cabinet has ranged from 1-5 kW for a 42U rack. A full rack of single (or dual) corded 1U servers in the aptly nicknamed "pizza box" form factor were the original drivers for a high density (scale out) rack and high outlet count PDU. As the equipment manufacturers adopted a number of new form factors such as 2U, 4U, 5U and 10U enclosures incorporating multiple server chips and large banks of both storage and memory, the demands placed on the power supplies within those servers went up, requiring higher power delivery per cord and outlet, driving the widespread adoption of C19 outlets over the previously dominant C13 outlets.

According to the 2013 IHS report "The World Market for Rack Power Distribution Units," the market for 5-10kW

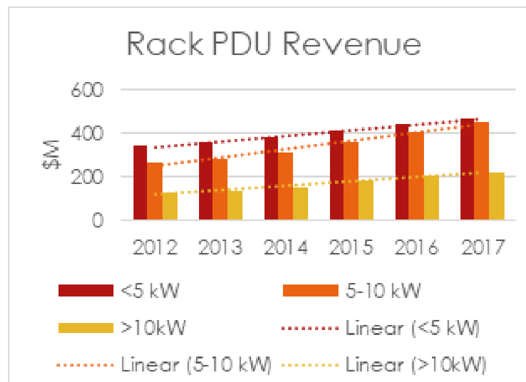

Source: IHS Electronics & Media, World Market for Rack Power Distribution Units 2013 Edition

and >10kW PDU will outgrow the low power (<5kW) segment at nearly twice the rate through 2017. The growth for the high-power segment will average over 11% while the low power PDU market is estimated to be just over 6%.

The demand for high-power, High Density Outlet Technology (HDOT) based PDUs from Server Technology during the years projected by IHS bear testimony to the accuracy of the IHS report. Figure 1 at right below shows the relative proliferation of unique PDU SKUs at various power levels and feature sets on the HDOT platform.

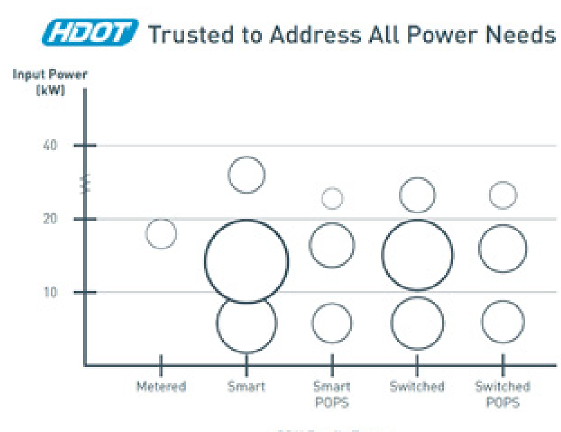

Figure 1: Server Technology Unique HDOT SKUs 2018

There are numerous challenges for both the Facilities and IT teams of a datacenter that are associated with implementing a power-dense data center. Amongst these are cost, heat, load balancing, environmental and power monitoring, increased network traffic, and configuration management of systems.

	Challenge	Impact
Facilities	**Reducing Capital Expenditure**	• Requires consolidation and integration of servers and systems along with virtualization • Reduces the amount of gear present on the data center floor
Facilities	**Higher Operating Temperatures**	• Unbalanced phases lead to power inefficiency • Unbalanced phases lead to larger electrical bills. • With heavier loads in the servers, the potential swings in power demand are greater, making the task of load balancing the power phases even more difficult
Facilities	**Phase Balancing**	• There is less time for IT and facilities to react in the event of a cooling system failure or thermal runaway of a server system. • Hot aisle or cold aisle containment are frequently needed. • HVAC systems need automation
Facilities	**Power Consumption**	• Higher voltages and currents are often demanded in order to support the desired IT load densities. • Today's dense IT loads often require 208V 3-phase 60A (21 kW) circuits, or even 415V 3-phase 60A (40+ kW)
Facilities	**Criticality of Power Availability**	• Less room for error, as more computations/ work are being done by a given asset. • Also less tolerance for a CRAC or CRAH to go down, as the speed with which thermal runaway can occur is faster.
IT	**Increasing Network Traffic**	• IOT, video, and sensors are all contributing to higher network loads. • Through consolidation and virtualization, the datacenter is able to reduce the number of switches and cables necessary to peform interconnect all of the working pieces of the datacenter
IT	**Configuration Management**	• Today's servers have more CPUS with more cores running more VMs than ever before. Making sure the configuration of the servers is correct becomes a difficult task requiring significant upfront efforts to develop automation, particularly at large scale installations
IT	**Complexity**	• In addition to the servers themselves, all of the other supporting systems are more complex: • virtualized storage on SAN or NAS, • software defined networking • the links to physical infrastructure such as UPS systems, CRACs and CRAHs over BACnet all require vigilent attention on the part of the IT team to ensure proper operations
IT	**Idle/ Unused VMs & Idle Machines**	• Unused or underutilized systems need to be identified and then either be put into a deep sleep state or taken offline altogether in order to reduce power consumption

Figure 2 – Density Challenges

Managing the Challenges of Density

There are a number of companies and industry organizations devoted to improving the state of the art of datacenter design and power utilization effectiveness. ASHRAE, IEEE, The Green Grid, and Emerge Alliance are among those organizations continuously seeking to foster new ideas and new technologies leading to more energy efficient datacenters. Most studies and papers from these groups have concluded that minimizing the number of power conversions and operating power supplies at the most efficient (highest) input voltage possible will result in significant savings on operational expenses (OPEX) for the datacenter. Modern green field datacenters are using 480V/277VAC, 415V/240VAC, and 380VDC as typical power going to their racks, with amperages ranging from 16A all the way up to 100A. A 1200mmx800mm cabinet fed with 415V/63A power has a usable capacity of 45.4 kW in a footprint just over 10.3 sq. ft. This configuration represents the practical upper limit of density for most datacenters today. This is a typical configuration for many modular datacenters that are readily available. Most applications that require higher density are running very specialized applications and are supported by liquid cooling technology.

Figure 2 shows a few of the challenges of density and their respective impacts.

By making use of the available power and configuration management tools provided by the various server manufacturers, the datacenter designer and operator can come to a consensus on the best approach to achieving the power and compute density needed for their application. These tools are readily available from manufacturers such as HPE, IBM, and DellEMC, and they can give the best idea of what level of compute density can be achieved for a given budget.

Another category of software tool that is available to support critical compute loads running in dense datacenters is Datacenter Infrastructure Management (DCIM). Typically incorporating support for SNMP, most DCIM tools provide a way of measuring power consumption, assessing available power capacity, profiling thermal conditions within the datacenter, exercising closed-loop control of HVAC systems, providing some degree of IT asset tracking, and determining the status of the various elements of UPS that may be distributed throughout the datacenter. DCIM tools help provide a “single pane of glass” for knowing what is going on in the datacenter. Depending on the DCIM tool selected, it may also provide hooks into any of a number of configuration management and virtualization management packages that enable the dense datacenter to rapidly configure new servers or move compute loads around to available underutilized resources within the datacenter.

Computational fluid dynamics software tools are used to determine how the cooling infrastructure is performing relative to the heat loads that are distributed within the datacenter. CFD will help ensure that hot aisle / cold aisle containment systems are performing optimally while identifying opportunities to redeploy assets to make the best use of available cooling capacity.

Operating the datacenter at the upper limits of the allowable inlet air temperature range helps achieve lower cooling costs and higher PUE, while requiring the temperature monitoring and control systems to operate reliably and respond quickly to changing conditions. The potential for thermal runaway is exacerbated in this case. The ability for modern PDUs to monitor both power consumption and temperature within the cabinet provides a “last line of defense” in the high temperature environment of the exhaust (hot aisle) side of the datacenter rack. The alerting, alarming, and reporting functions of modern PDUs provide critical information to HVAC systems, DCIM tools, and the IT and Facilities personnel responsible for overseeing the successful operation of the datacenter. The more individual loads within the racks, the more critical it becomes for the PDU to have granular measurement and reporting capabilities along with supporting the targeted hot aisle temperature.

Placing many servers in a single cabinet plugged into a 3-phase power source necessitates careful planning and monitoring as well. The datacenter can minimize wasted power and maximize power delivery to the IT gear when the loads on all three phases are as close to equal as possible. This requires the servers and the PDUs to report power consumption and provide an aggregate figure of consumption by phase in order to give the datacenter operator a chance at having a maximally efficient datacenter. Utilizing the alternating phase HDOT modules minimizes the potential for accidentally unbalanced loads across the phases.

Solutions

Providing power to a dense heterogeneous computing environment requires highly integrated and extremely reliable power distribution. High Density Outlet Technology(HDOT) Cx from Server Technology provides both. The PRO2 HDOT Cx is available in a "build your own" PDU that allows the user to configure a PDU meeting their exact specification. By selecting the appropriate input power (voltage, amperage, and phases), input cord, PDU orientation, outlet mix, and connectivity, the datacenter designer and IT specifier can get a PDU that is custom-tailored to their high compute density application.

HDOT Cx modules are available in our alternating phase configuration and feature our unique 2-in-1 outlet design that supports both C14 and C20 plugs, thus future-proofing the data center rack by allowing the PDU to be re-used and re-configured more easily than any other PDU on the market today.

Benefits

The PRO2 HDOT Cx delivers on the promise of modern PDU design capabilities. With its modular construction and its completely re-thought outlet layout, PRO2 HDOT Cx provides unique, 2-in-1 outlets that act as both locking C13 and C19 power outlets, all in a standard form factor that fits most datacenter cabinets. At just 2.2" wide, 2.5" deep, and 70" tall, the PRO2 HDOT Cx PDU is suitable for all server racks 42U and taller. PRO2 HDOT Cx is also capable of operating at full power load in a 60°C environment, allowing the datacenter to run with a warm ambient temperature. The benefits of the alternating phase modules are numerous – they simplify load balancing while minimizing cord lengths, and they make it easier to identify which asset is plugged into each outlet of the PDU.

The Smart, Smart POPS, Switched, and Switched POPS versions of PRO2 HDOT Cx allow for remote monitoring of power and temperature, provide alerting and reporting, thus enabling the datacenter manager to make the most use of the available datacenter compute infrastructure and available power. PRO2 Switched HDOT Cx provides individual outlet control, and POPS delivers power consumption data from the individual outlet. The PRO2 firmware features an open MIB and OID tree information that is freely available for download from the Server Technology website, and PRO2 firmware is supported by more DCIM tools than other PDUs on the market.

Why Server Technology for High Density Data Centers

Server Technology is the foremost provider to the data center industry of high density, high power intelligent power distribution products. With the best initial quality and longest mean time to failure (MTTF), Server Technology products provide uncompromising reliability and value for the data center.

Server Technology holds many of the key patents for providing power in a zero-U PDU configuration compatible with today's racks and dense compute infrastructure requirements.

Our HDOT-based products have been recognized multiple times by DCS and others in the data center industry for being the global leader of innovative, intelligent power distribution products. Our PDUs offer solutions for data centers, telecommunications operations and branch offices. Over 60,000 customers around the world rely on Server Technology's power distribution units (PDUs) and power management and measurement solutions to Stay Powered, Be Supported, and Get Ahead.

2018 WINNER DCS AWARDS

References:

www.ashrae.org
www.thegreengrid.org
http://h10032.www1.hp.com/ctg/Manual/c00064724.pdf "Optimizing facility operation in high density data center environments"
http://www.dell.com/html/us/products/rack_advisor_new/
http://www.zdnet.com/blog/btl/data-center-design-101/26158 by Larry Dignan, 2009
http://www.te.com/dm/datacenter/media/401947HK.pdf The Three Principles of Datacenter Design

Interested in learning more about how Server Technology can help you manage and distribute power in your datacenter?
Visit us online at: www.servertech.com/products

North America Headquarters
1040 Sandhill Road
Reno, Nevada 89521
1-775-284-2000 Tel
1-800-835-1515 Toll Free
1-775-284-2065 Fax
sales@servertech.com
www.servertech.com
www.servertechblog.com

U.K., Western Europe, Israel, & Africa
Fountain Court
2 Victoria Square
Victoria Street
St. Albans
AL1 3TF
United Kingdom
+44 (0) 1727 884676 Tel
+44 (0) 1727 220815 Fax
salesint@servertech.com

Server Technology. A brand of legrand

stay powered. be supported. get ahead.

©2018 Server Technology, Inc. Version 08/10/2018. Sentry and Server Technology are registered trademarks of Server Technology Incorporated. Information is subject to change without notice. Printed in USA. Server Technology offers a wide range of products for North America and Global markets; for more information visit our website at www.servertech.com