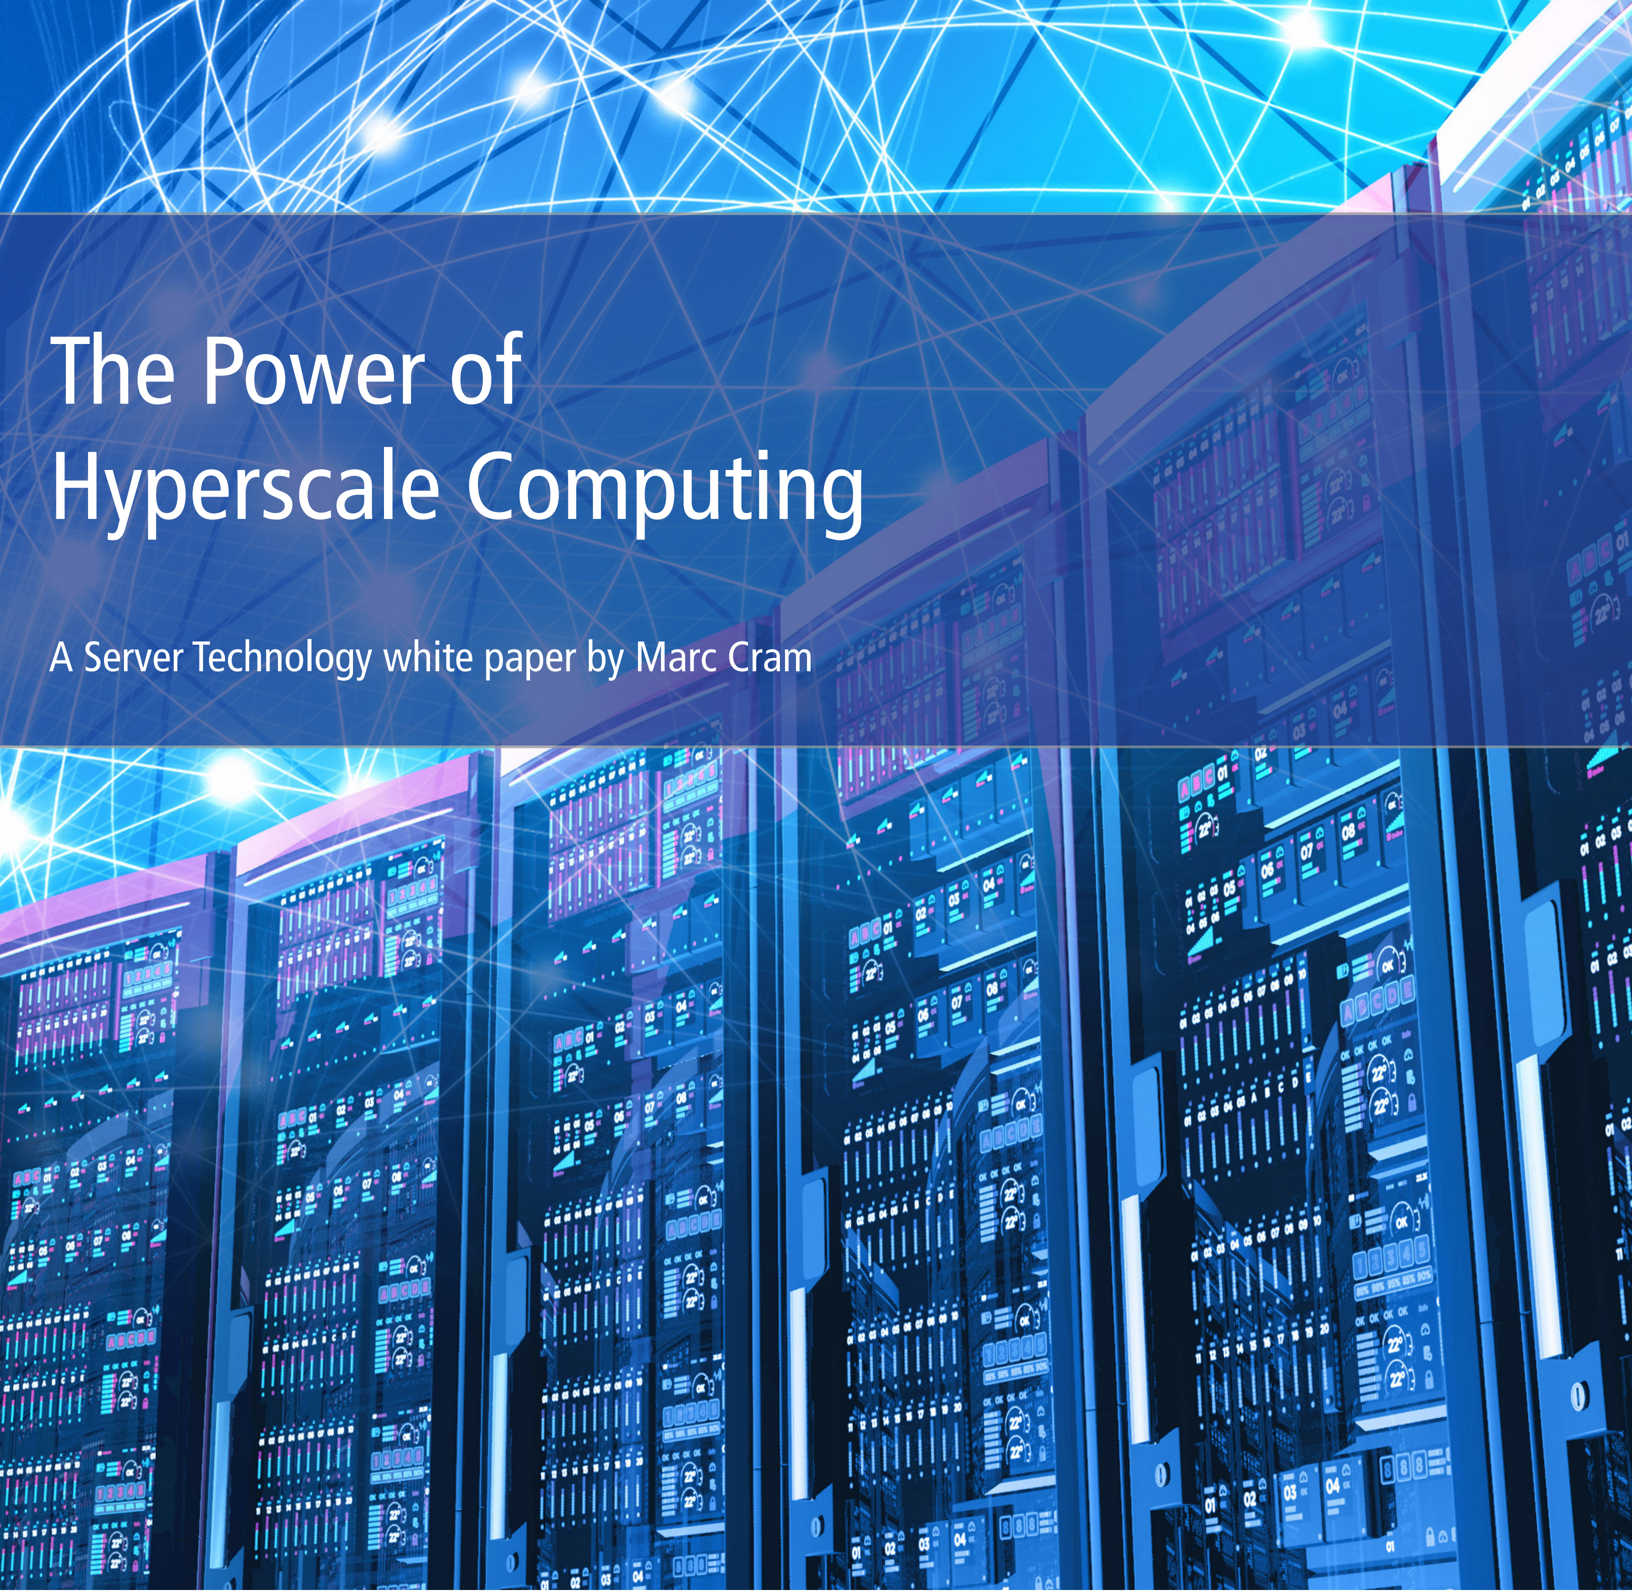

The Power of Hyperscale Computing

A Server Technology white paper by Marc Cram

Stay Powered. Be Supported. Get Ahead.
www.servertech.com

Server Technology
Quality Rack Power Solutions

EXECUTIVE SUMMARY

The growth of cloud IT adoption continues unabated. Today's landscape of cloud providers is dominated by a small handful of companies based in the United States and China that deploy company owned hyper-scale data centers. For the foreseeable future, these few will account for much of the growth to come in the cloud IT market, as many businesses decide to give up owning their own data centers and migrate their applications to cloud platforms. Figure 1 (below) shows the projected growth in the number of hyperscale data centers through 2020, along with the percentage of servers expected to reside in those data centers. By 2020, 47% of all servers sold are expected to go to hyperscale customers. For legacy hardware provid-ers, such as HPE and Dell, this trend portends a narrower customer base having both a higher technical acumen and a closer attention to the bottom line. HPE and Dell will see their sales to hyperscale customers increasingly captured by "no name" white box vendors who cater to giving their customers exactly what they want and no more. One server configuration no longer fits all. Open Computing is increasingly talked about as a path to lower costs and higher efficiencies.

New topologies in networking yield higher bandwidth for both "north-south" and "east-west" traffic in the hyperscale data center. Facebook has open sourced the specifications for their commodity-based switch designs and shown to the world the many benefits of software defined networking (SDN) and network function virtualization (NFV).

For the hyperscale data center operators, innovations in power delivery and new approaches to cooling enable the rack power density within the hyperscale data center to increase while improving the efficiency of the overall infrastructure, resulting in extremely low PUE.

This white paper discusses what comprises a typical hyperscale data center and the role that power infrastructure plays in enabling these data centers to operate as reliably and efficiently as possible.

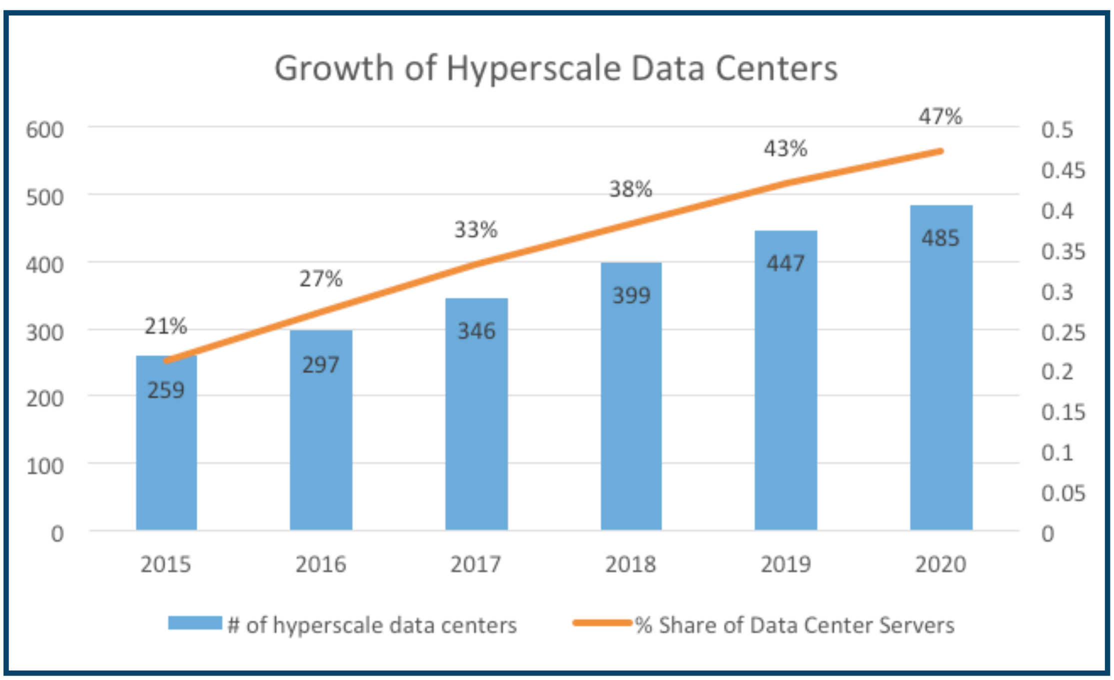

Figure 1- Source: Cisco Global Cloud Index, 2015-2020, Synergy Research

HYPERSCALE DEFINED

• • •

Large data centers, full of thousands of low cost, minimalist servers supporting a distributed computing environment. Usually intended for big data (like Hadoop) or cloud computing.

The compute and storage scale independently of one another, similar to Intel's Rack Scale Architecture (disaggregation).

WHAT IS HYPERSCALE?

Techopedia says that hyperscale computing "refers to the facilities and provisioning required in distributed computing environments to efficiently scale from a few servers to thousands of servers. Hyperscale computing is usually used in environments such as big data and cloud computing. Also, it is generally connected to platforms like Apache Hadoop.

"The structural design of hyperscale computing is often different from conventional computing. In the hyperscale design, high-grade computing constructs, such as those usually found in blade systems, are typically abandoned. Hyperscale favors stripped-down product design that is extremely cost effective. This minimal level of investment in hardware makes it easier

to fund the system's software requirements."[1]

Firms that operate hyperscale data centers tend to have strong engineering teams capable of specifying every aspect of the data center, from IT hardware, to software stack, to networking protocol, to HVAC and power infrastructure, and these teams are focused on achieving maximum performance, lowest cost, and peak power efficiency.

"According to a recent DOE study, most of the servers that have been responsible for the 3 percent annual increase in shipments have been going into hyperscale data centers. The cloud giants have created a science out of maximizing server utilization and data center efficiency, contributing in a big way to the slow-down of the industry's overall energy use, while data center providers have made improvements in efficiency of their facilities infrastructure, the power and cooling equipment that supports their clients' IT gear."[2]

[1] https://www.techopedia.com/definition/28869/hyperscale-computing

[2] http://www.datacenterknowledge.com/archives/2016/06/27/heres-how-much-energy-all-us-data-centers-consume/

WHO ARE THE HYPERSCALERS?

The most powerful names on the web today are the hyperscalers that began in the 1980s and 1990s during the internet boom. Microsoft turned its success with MS-DOS for the IBM PC into the over $60B behemoth it is today, migrating from packaged software (remember binders with floppy disks?) to SaaS (Office 365) and PaaS (Azure) with a global data center footprint. Amazon began in 1994 as an online bookseller. Today they are a $90B+ ecommerce juggernaut, as well as the world's largest cloud service provider. Facebook is the world's largest social media based enterprise, with numerous large data centers spread around the globe. Alibaba is China's home grown answer to Amazon's online sales presence, and is expanding their data center footprint beyond China's borders to provide competition to Amazon and others. Google began as a research project in the 1990s and quickly rose to become the largest search provider on the internet. They too rely on massive data centers in multiple geographies to deliver the quick performance everyone has come to expect when entering a query into the search bar of their web browser.

Other large and rapidly growing hyperscale data center occupants are Apple, Tencent, Baidu and EBay.

Most of the hyperscale data center builders own their own infrastructure. Originally populated with hardware made by Compaq (now Hewlett Packard Enterprise), Dell (now Dell Technologies, an amalgam of privately held Dell and EMC) and Cisco, the hyperscale data center operators are moving to lower cost, purpose built, custom compute and networking solutions coming from the ODMs of Taiwan and China, such as Quanta, Inventec, Wiwynn, Foxconn, Supermicro and Inspur.

Figure 2- https://www.google.com/about/data centers/gallery/#/all/2

HYPERSCALE DATA CENTER OPERATORS

• • •

Here are a few of the major hyperscale data center operators today

Google
Baidu
Facebook
Microsoft
Apple
Alibaba
Amazon
Softlayer
Tencent
eBay

The compute and storage scale independently of one another, similar to Intel's Rack Scale Architecture (disaggregation).

HYPERSCALE ENABLERS

• • •

Hyperscale data center operators require the partnership of a supply base that can support high volume, quick turn production of hardware – servers, storage, networking, racks, and so forth.

The major providers in this space include:

- Supermicro
- ZT Systems
- Hyve
- Quanta
- Foxconn
- Wiwynn
- Inspur
- Dell
- HPE
- Lenovo

Figure 3-
Inspur Cloud Rack

THE GROWTH OF HYPERSCALE DATA CENTERS

The idea of hyperscale data centers came from the notion that if software applications are robust enough, they can be gracefully moved from one compute instance to another without concern for the underlying hardware they are running on. The introduction of hypervisors and virtual machines as a hardware abstraction layer allowed software applications to be easily moved from one physical machine to another. Born at the height of the internet boom of 1999, VMware was the first commercially successful implementation of hypervisors and virtual machines on x86-based server hardware. The beauty of VMware is that "it enables a virtual machine guest to be paused, then move or copy that guest to another host and there resume execution exactly at the point of suspension."[3]

Once the software for hypervisors and virtual machines was in place, the evolution of data center operational goals began in earnest. It became the industry mantra to strive for "five nines of uptime" (99.999%), or roughly 5 minutes of downtime per year. Virtualization meant that the failure of any given server became a non-issue. Having a data center full of virtualized servers enabled a form of redundancy that did not require any particular server to be up and running.

Google, having begun in 1996, saw the opportunity that hypervisor and virtual machine (VM) technology (such as Xen and KVM, or Kernel Virtualization Machine) brought to their rapidly evolving application base of search, email, storage, and advertising. Once their applications were written in a way that supports VM adoption, Google began designing and building their own data center hardware in the form of custom servers. Today, it is estimated that Google runs over 1 million servers that are spread throughout the world in numerous data centers. Meanwhile, virtual machines are being displaced by container technology, another form of application abstraction through software such as Docker or Kubernetes. "Containers allow for the isolation of resources used by multiple applications running on a single machine."[4] Using containers means that a server can load a single instance of an operating system, and spin up a containerized application in faster time than a hypervisor could load up an entire operating system + application as would be found in the virtual machine environment. Container technology allows the data center operator to run more computational load concurrently on a given server than could be done with VM technology alone.

[3] https://en.wikipedia.org/wiki/VMware
[4] https://www.theregister.co.uk/2014/05/23/google_containerization_two_billion/

THE IT ARCHITECTURE INSIDE THE HYPERSCALE DATA CENTER

Along with changing the philosophy on compute infrastructure, the hyperscale customers have changed the way that they address networking. Traffic in a data center used to go mostly “north-south”, entering the data center on a main switch, going to an end of row aggregator switch, then on to a top of rack switch. In this arrangement, east-west traffic going between servers inside the data center was difficult to sustain for any significant period.

Facebook and other data center operators have taken the path of network function virtualization (NFV) deployed on open source hardware (Wedge, 6-Pack) and software (FBOSS). Wedge is commodity switching hardware of their own design that Facebook released to the world through Open Compute. Wedge is modular, and replaces most name brand top of rack switches. Facebook also refined the switching topology of their infrastructure to incorporate a “leaf and spine” approach such that any given server can quickly gain access to other servers within the same data center. The impact for Facebook has been to enable them to sustain higher amounts of website traffic, support long term storage of users’ pictures and posts, and deliver improved site performance on all key metrics.

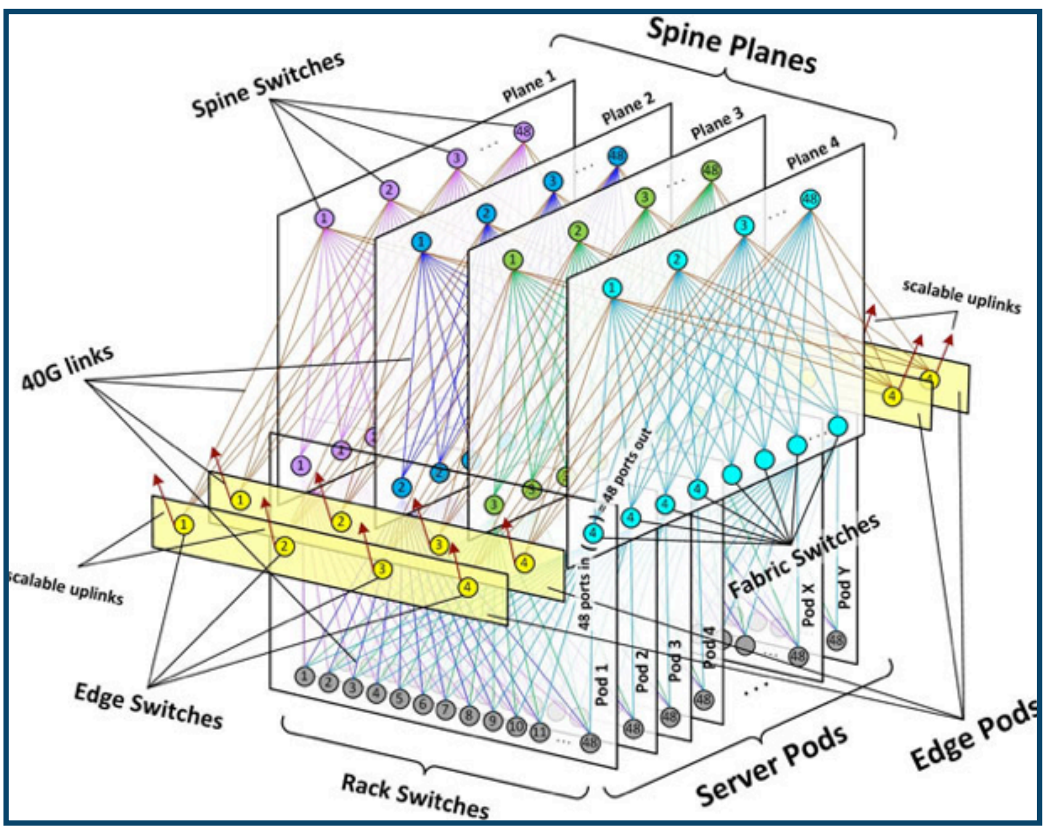

Figure 4- Fakebook’s data center fabric design

THE IMPACT OF PHOTONICS AND DISAGGREGATION ON HYPERSCALE

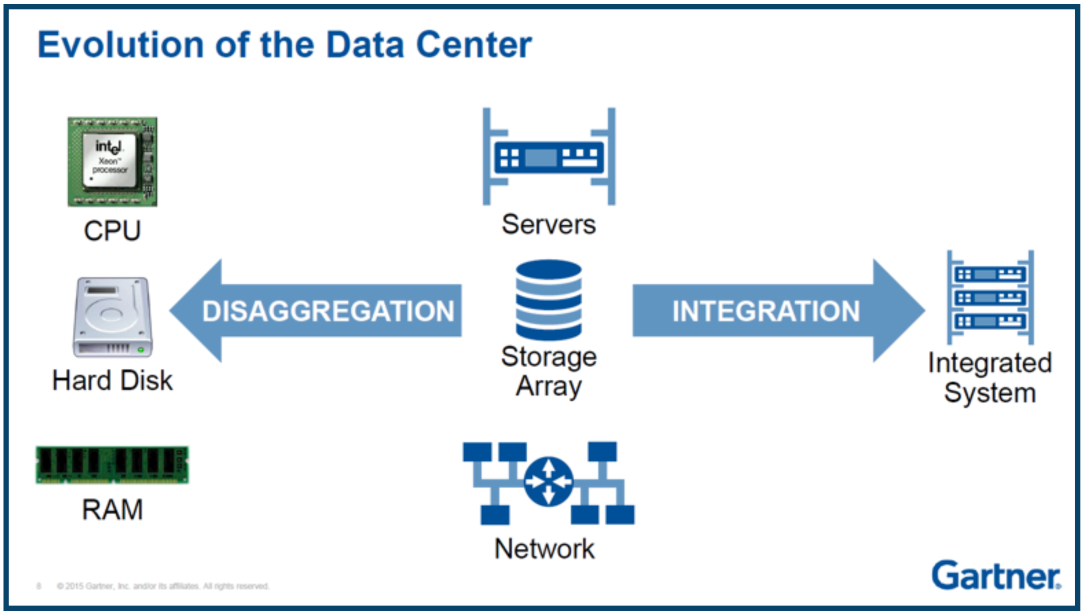

Figure 5- Next Generation Data Center Showdown - HCIS, Gartner 2016

The Facebook data center relies heavily on the concept of disaggregation. "Rack Scale is a disaggregated architecture where every individual component – processors, switches, hard drives and memory cards - can be replaced independently, and the whole rack behaves as one large virtualized server with shared cooling and power.

"The components are connected to each other with a high-speed optical fabric. Intel believes that Rack Scale will be the big thing after next - following and subsuming microserver architectures such as HPE's Moonshot. To this end, it has been working hard to make the advanced semiconductor technology called silicon photonics - the centerpiece of Rack Scale - more accessible."[5]

The use of optics in the data center enable lower cost per transmitted bit, lower power per bit, and better performance over a distance than does a comparable length of copper cabling.

100 Gb/sec and higher data transfer speeds necessitate the use of optics in the data center, whether through fiber or free air optics. The choice of which type of fiber comes down to a decision between multimode and single mode fiber. For reliable transmission over the longest distance, the use of single mode fiber is recommended.

For short distance communications of 1-3 meters in length, copper dominates. If going 10-30m, usually between racks in a row, then multi-mode optical fibers can be used. Going beyond the row, single-mode fiber is the rule.

[5]http://www.datacenterdynamics.com/content-tracks/servers-storage/ericsson-to-sell-intels-hyperscale-kit-to-network-operators/93484.fullarticle

Intel began shipping 100Gb/sec optical transceivers in August 2016. Silicon photonics is the term used by Intel for when a laser is integrated into a silicon chip. The Intel process for fabricating on-die lasers can be built at scale across a wafer, enabling them to produce a range of devices that can have native support for optical interconnections. The market for integrated photonics is expected to exceed the market for discrete optics beginning in the year 2018.

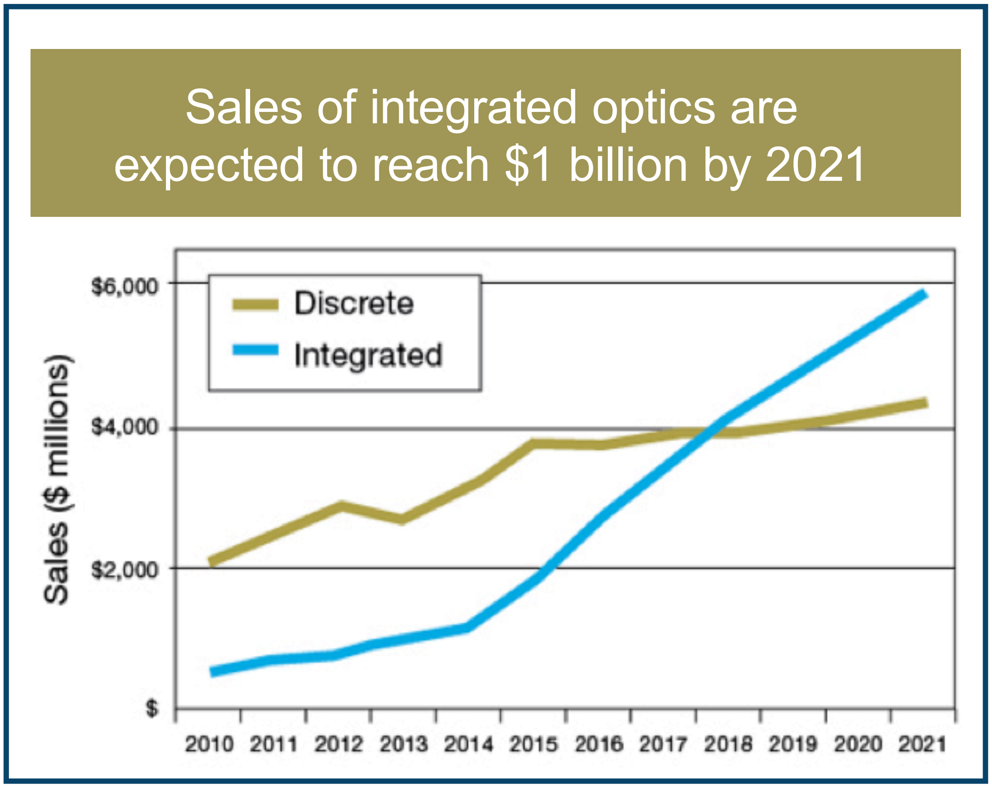

Figure 6- http://www.fibre-systems.com/feature/tipping-point-silicon-photonics

Future data rates are expected to go even higher than they are today – 400 Gb/sec in 2-3 (2020) years, according to Intel. Intel wants to skip the intermediate transition to 200 Gb/sec, and sees that integrated photonics is their path to supporting the bandwidth needs and enabling the disaggregated architecture espoused by Facebook and Open Compute.

Disaggregation takes the datacenter rack in the opposite direction of the hyperconverged systems that are available from the likes of Simplivity (now HPE) and Nutanix. Hyperconverged systems are tightly coupled computer, storage, and networking devices that have a deep software stack that allows newly added systems to be discovered and added to the logical stack already in the datacenter (Figure 5) and perform as if the systems are all one logical drive, for example. Rack systems as envisioned by Intel offer entire shelves and racks full of storage or compute devices, joined to other resources through a fabric interconnect like the Facebook leaf and spine architecture (see Figure 4).

Disaggregation works best in an environment where failure of individual components can be tolerated due to the robustness of the software systems running on the underlying hardware. Machine virtualization and container software enable a high degree of hardware fault tolerance, whether on the storage or the compute side. However, even the best software must run on a server somewhere, and that server is going to require electrical power.

"Form ever follows function" – Lewis Sullivan

THE IMPACT OF POWER DENSITY AND COOLING ON HYPERSCALE DESIGN

During the past twenty years, the data center industry has seen rack power density go up commensurately with compute and storage densities. More servers and more hard drives are placed into a single rack today than ever before, in a scale out approach (lots more servers) rather than a scale up (towards mainframe). Whereas a typical IT rack used to consume one to three kilowatts, today it is not uncommon to find a load of twenty to forty kilowatts in the cabinet. The US National Renewable Energy Laboratory (see Figure 7) reports that 30 kilowatt racks are not uncommon today.[6]

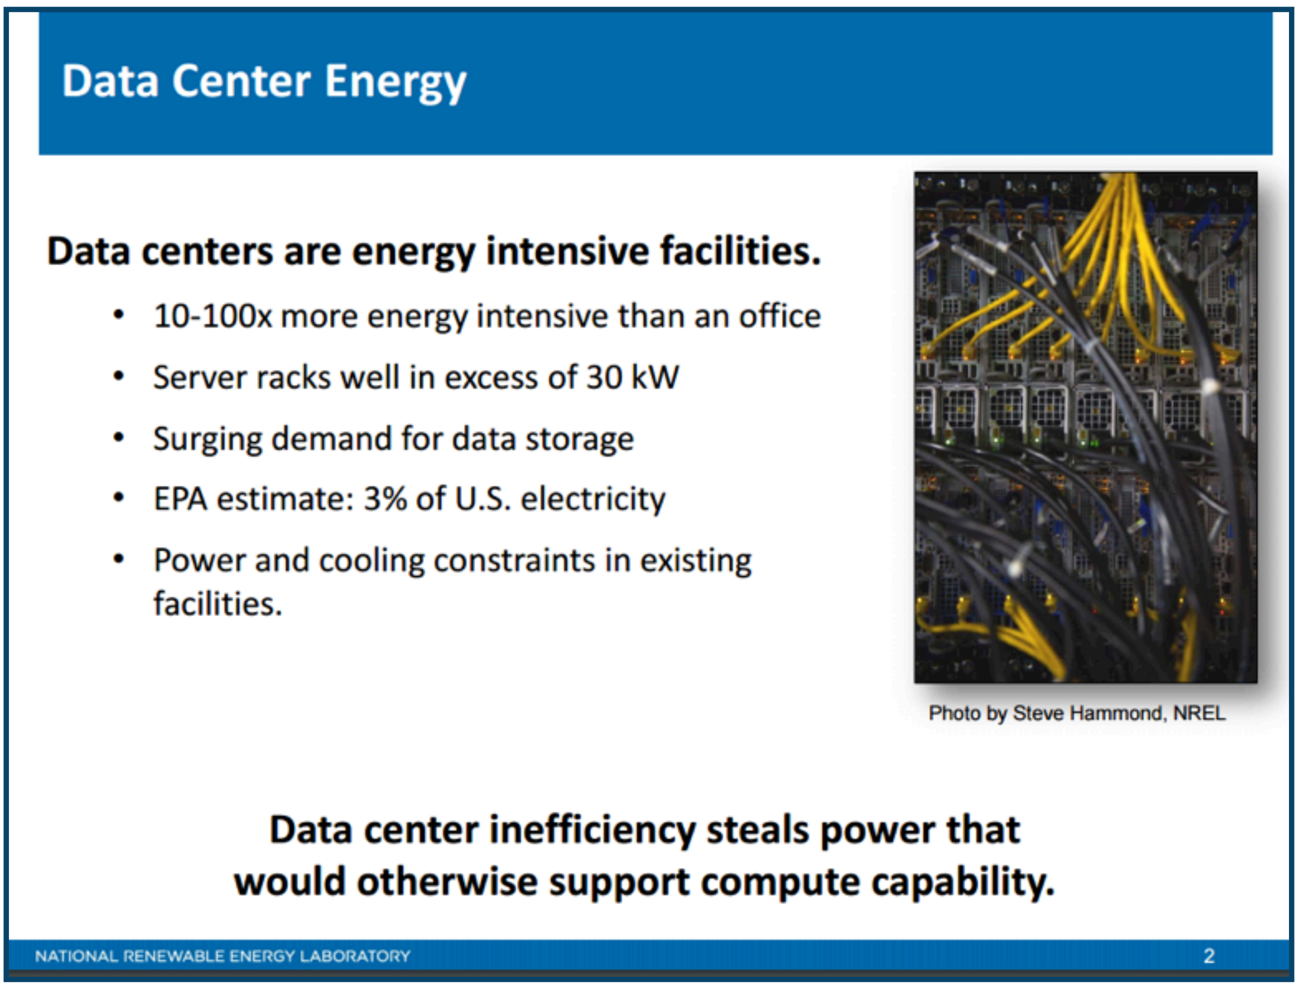

Figure 7- http://www.nrel.gov/docs/fy13osti/58902.pdf

[6] http://www.nrel.gov/docs/fy13osti/58902.pdf

Google used 5.7 terawatt hours of energy in 2015 according to Joe Kava, who heads the company's global infrastructure. "Data centers make up the vast majority of that," said Kava, which is why Google has also taken a leadership position in procuring renewable energy for its cloud campuses.[7]

Early data center designs incorporated large centralized battery banks, or distributed UPS systems located along the walls of the datacenter or at the end of a row of IT cabinets. Recent thinking by Microsoft and other hyperscale data center operators led to the elimination of the larger UPS systems in modern data center construction. Instead, the UPS functionality is being distributed across the datacenter, within the individual IT rack. The Power Shelf approach of Open Compute provides for a small lithium ion battery stack intended to provide only enough ride-through time to gracefully move VMs and containers to new locations before hardware servers shut down.
New servers based on the latest chip technologies from Intel, AMD, nVidia, and ARM deliver more compute operations per second and operations per watt than ever before. At the same time, the cost of data center real estate has gone up in most markets, causing the data center operator to look for ways of having taller IT cabinets to maximize the utilization of the square footage of the data center. Having a tall column full of IT gear generating heat leads to cooling challenges. The data center architect must early on commit to either air or water as the cooling medium of choice. In some facilities, both are deployed.

Efficiently air cooling a facility requires adherence to a number of design tenets:

- Alternate the direction of equipment to create naturally formed hot aisles and cold aisles
- Use blanking panels in empty rack spaces to ensure no leakage of cold air directly into the hot aisle
- Practice either hot aisle or cold aisle containment
- Natural convection can be used to move air through a well-designed facility – cold air falls to the floor, warm air rises to the ceiling, so overhead chilled air service to the cold aisles and either raised ceilings or chimneys atop the back of the racks can be used to push/pull air through the IT systems

[7] http://datacenterfrontier.com/google-used-5-7-terawatt-hours-of-energy-in-2015/

Recent air cooled data center designs from Microsoft feature large fan walls to move sufficient amounts of air.

Figure 8- A look at server for a Cray CS400 supercomputing cluster, which is cooled with an Asetek system that brings warm water to the chip. (Photo: Rich Miller)


Figure 9- Fan wall inside Microsoft's fifth-generation data center in Quincy, Washington (Source: Microsoft video)[8]

[8]http://www.datacenterknowledge.com/archives/2016/09/27/latest-microsoft-data-center-design-gets-close-to-unity-pue/

Adiabatic cooling relies on the process of reducing heat through a change in air pressure cause by volume expansion.[9] Adiabatic processes have been used in data centers to facilitate "free cooling" methods that make efficient use of water and electricity.
Liquid cooling is well suited to those applications where the IT cabinet power and thermal density exceed the cooling capacity of air flowing at a reasonable velocity (up to a few hundred CFM). Liquid cooling comes in many forms – chilled doors, chilled shelves, direct spray cooling onto a chip, and so forth. "In the HPC world, everything will move to liquid cooling," said Paul Arts, technical director of Eurotech. "In our vision, this is the only way to get to exascale. We think this is the start of a new generation of HPC, with enormous power. We are just at the beginning of the revolution."[10]

Immersion is the most recent innovation on the liquid cooling front. By placing the IT hardware directly into a liquid cooling medium, the maximum thermal transfer rate can take place, and hence the power density can be maximized. However, immersion comes with several drawbacks:

- Mess – any maintenance taking place requires the hardware to be removed from the liquid medium.; the potential for spillage is high
- Specialized hardware – spinning hard drives (old school magnetic disks) must be designed from the onset to operate in this environment, meaning they must be sealed and pressurized. Otherwise they may elect to go with solid-state disks.
- Higher up-front costs – datacenters for immersion cooling are rare, and retrofitting an existing facility can be costly
- Weight – not suitable for raised floor environments
- Specialized containers - Baths or vats large enough to hold entire racks of IT hardware along with the cooling liquids
- Filtration systems – circulating liquids always require some means of being cleaned

Ultimately, the decision on cooling will be driven by the parameters under which the datacenter will be operated: frequency of hard changes/repairs, number of available skilled personnel, how much time is allowable for repairs to be made, the level of power efficiency desired, and so forth.

[9] http://whatis.techtarget.com/definition/adiabatic-cooling
[10] http://www.datacenterknowledge.com/archives/2014/12/09/rise-direct-liquid-cooling-data-centers-likely-inevitable/

THE ROLE OF ENVIRONMENTAL MONITORING IN HYPERSCALE DATA CENTERS

Environmental monitoring is standard procedure for most hyperscale data centers. This is due to their reliance on high ambient operating temperatures and high power densities within the IT racks for achieving maximum energy efficiencies. The most efficient operating point for the datacenters is to run off free air cooling or adiabatic cooling, with the set point for the cooling systems adjusted to within just a few degrees of the highest possible inlet air temperature that the IT systems can tolerate. This minimizes the amount of time the cooling systems operate, and thus minimizes the amount of energy expended on cooling. As such, knowing the air temperature at every IT rack becomes crucial. This level of data granularity supports thermal modeling of the data center as well as localizing the cooling system response to the bare minimum, such as opening a louver or floor vent close to the rack that is in danger of experiencing a thermal excursion. In this example, the cooling system response must be timely to prevent thermal runaway from occurring. Numerous methods exist for environmental monitoring to be accomplished in the datacenter. There are wireless sensors such as those from RFCode, wired sensors from the likes of Johnson controls that are good for room level and system level (CRAC, CRAH) monitoring, and then there are rack level monitoring solutions incorporated into most intelligent IT rack PDUs, such as those available from Server Technology. By residing within the IT racks, right at the source of most heat generation within the datacenter, the sensors supported by rack PDUs provide an economical means of providing real-time information that is more responsive than sensors mounted at the ceiling or within the ductwork of the data center air handling systems.

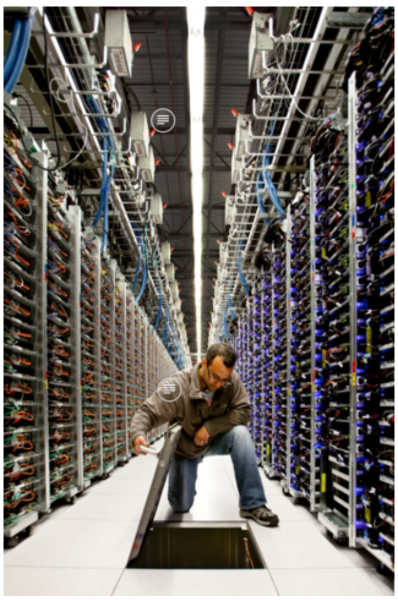

Figure 10 - https://www.google.com/about/data centers/gallery/#/all/26

POWERING THE FUTURE NEEDS OF HYPERSCALE DATA CENTERS

Numerous papers are available to the casual investigator that document what was happening to the growing power demands of the data center industry during the early 2000's. Had the industry not been helped by Moore's law (putting more functionality onto individual server CPUs) and server virtualization, the growth in power demand for data centers would have likely grown to exceed 10% of the USA's total power production capacity. Instead, the IT industry took a leadership role in policing itself to make IT systems more power efficient. The power draw per server chip (CPU) was capped by Intel and AMD at certain levels, and with each successive generation of CPU they improved the instructions per cycle and reduced the watts per instruction while remaining within a given power envelope.

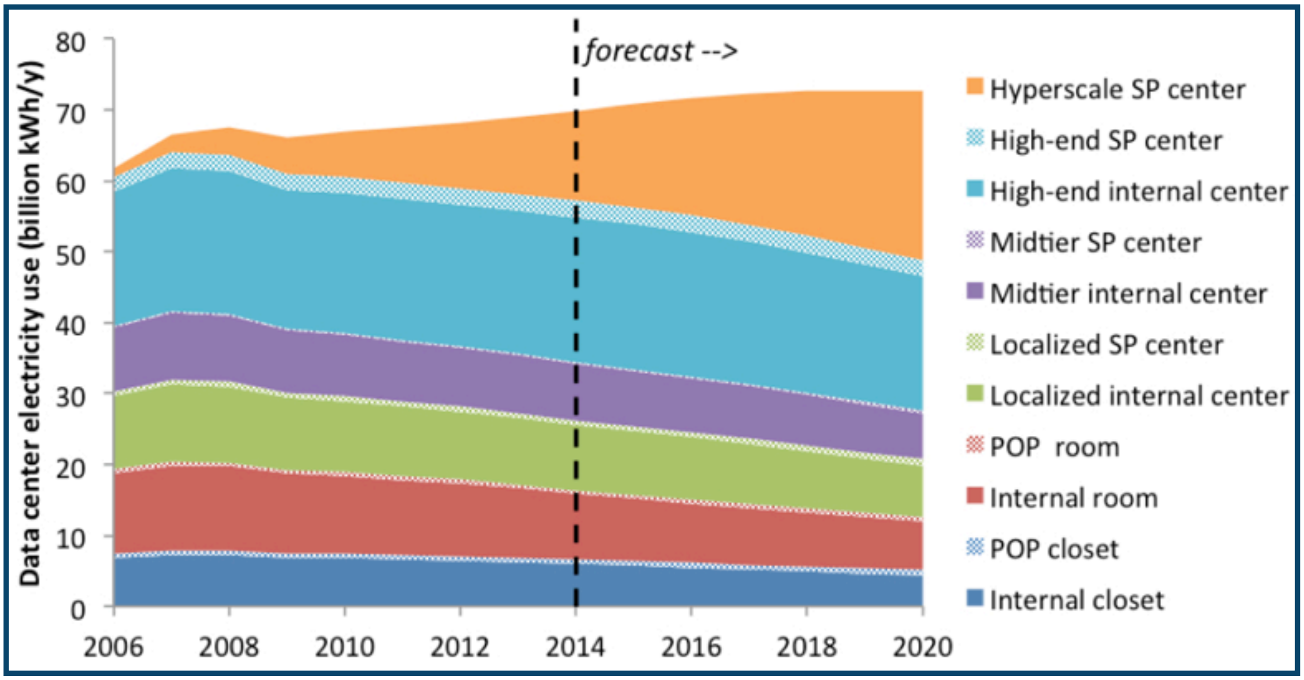

Figure 11- US DOE, Lawrence Berkeley National Labs; SP = Service Provider[11]

[11]http://www.datacenterknowledge.com/archives/2016/06/27/heres-how-much-energy-all-us-data-centers-consume/

CHOOSING THE RIGHT VOLTAGE

• • •

For a given power (such as 20kW to the rack), the higher the voltage delivered, the lower the amperage needed. Lower amperages can run on smaller diameter copper wiring than higher amperages. I^2R losses are also reduced with lower amperages.

Start at the highest voltage your utility or your onsite generation permits you to run, and specify your server power supplies (or power shelf) accordingly.

- 480VAC, 3ϕ
- 415VAC, 3ϕ
- 400VAC, 3ϕ
- 380VDC

Looking ahead, many hyperscale data center operators are making the commitment to power their facilities with renewable energy. First achieved through the purchase of renewable energy certificates and offsets (buying green energy from the local electrical utilities, such as hydroelectric), today many hyperscale data centers are being powered from energy generated on site by fuel cells or solar arrays. Apple and Facebook already have multiple locations running from solar. Microsoft has a biogas facility in Cheyenne
Wyoming, purchases wind power from Pilot Hill Wind, and generates solar energy at numerous facilities globally. Amazon is building a 253 MW wind farm in Texas, and Google has committed that it will run all its data centers with renewable energy during 2017. "We are the largest corporate purchaser of renewable energy in the world" says Joe Kava, Google's Sr VP of technical infrastructure.[12]

There are numerous proposals across the IT industry regarding what type of power infrastructure is best suited to the hyperscale data center. Early on, Facebook chose to go with 480V/277V AC power to the rack and 12VDC power to the IT loads within the rack. In Google's recent contributions to Open Compute, they have proposed 48V DC to power the servers, with direct conversion point of load power supplies running from 48V to the working voltages and amperages needed on the motherboard to run the CPU, the memory, and so forth. In Google's implementation, the power going to the rack from the power mains is 480/277V AC (or 400V/230V AC) with a 3phase rectifier on the power shelf converting AC power to DC power within the rack. Microsoft's most recent Open Compute contribution has 480/277VAC to the rack and 277V into the server, with dual 3phase power supplies in the server enclosure and a basic (non-intelligent) PDU that supports blind mating to the servers in the rack. See Figure 11.

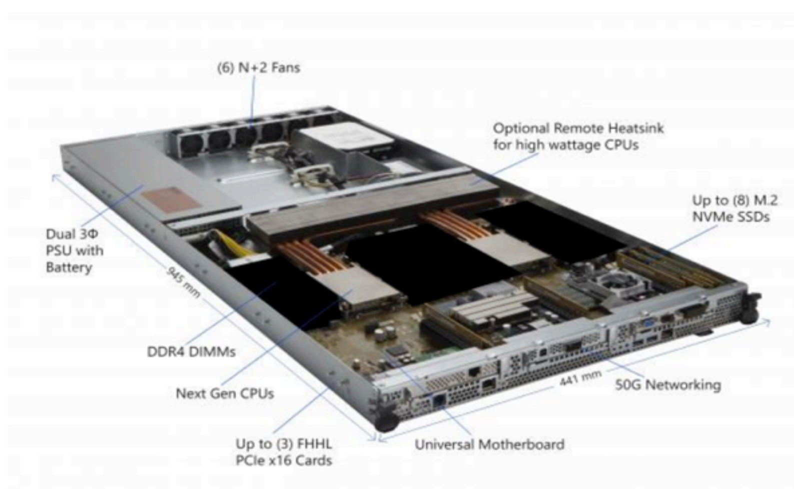

Figure 12- Microsoft Open Compute design

[12]https://www.nytimes.com/2016/12/06/technology/google-says-it-will-run-entirely-on-renewable-energy-in-2017.html?_r=0

Choosing which way to power the IT equipment in the hyperscale data center requires a series of questions to be answered followed by some decisions that are tempered by operating goals that are to be met by the data center. Some of these questions can be seen in Appendix B.
For those companies seeking to own the power generation equipment to run their data center, generating DC power on site and delivering that to the rack can make good sense. In that circumstance, bringing 380V DC power to the rack and running a DC to DC converter to drop to 48V is an option. Or they may even choose to bring 48V to the rack depending on distance between the power source and the IT rack.

Figure 13- Widows Creek Power Station, Alabama; photo TVA via Flickr

PDUS - THE FINAL STEP IN POWERING HYPERSCALE

Server Technology makes power distribution products that work at all voltage and amperage ranges needed to support today's hyperscale data centers. No matter the rack vendor or the IT rack load (whether 10kW or 100kW per cabinet), our PDUs work in ambient temperatures up to 60oC. And we have products that support both AC and -48V DC applications.

For hyperscale deployments, there are numerous PDU choices to be made that are driven by the following factors:

- Cost
- Feature set
- Reliability
- Simplicity versus flexibility

Basic PDUs offer simplicity, reliability, and support for high operating temperatures. They are well suited to applications where it is OK to lose a locked up or failed server that may have been running a dozen VMs or several dozen containerized applications. They lack remote power consumption monitoring and reporting capability, necessitating either upstream or downstream gear to support gathering information for PUE calculations.

PRO2 Smart and Smart POPS PDUs integrate inline power monitoring and remote monitoring capabilities through SNMP. These are great for those racks where the loads are neither static nor equally distributed, and there is a diversity of hardware within the rack that would require multiple protocols / software interfaces to extract power information from the respective devices. Environmental monitoring capability is an integral feature of Smart and Smart POPS PDUs, with support for multiple temperature and humidity probes that can be deployed inside and outside the rack. Smart and Smart POPS PDUs enable alarms to be set for user selectable predefined power and temperature thresholds.

PRO2 Switched and Switched POPS PDUs offer the most features and flexibility. They have the same capabilities as their Smart brethren, and give the data center operator the chance to reset power to the individual IT device via a single interface, whether that is a top of rack switch, a load balancer, storage array, or compute node. Furthermore, in the event of a main power source failure, they support load shedding by shutting off non-critical devices in a predefined prioritized order. Load shedding helps ensure a controlled migration of IT workloads so that the hyperscale data center can provide uninterrupted operation to the user community.

WHY SERVER TECHNOLOGY

Server Technology's power strategy experts are trusted to provide rack PDU solutions for demanding data centers worldwide ranging from small technology startups to Fortune 100 powerhouses. Because power is all we do, you will find us in the best cloud and colocation providers, forward thinking labs and telecommunications operations. Server Technology customers consistently rank us as providing the highest quality PDUs, the best customer support, and most valuable innovation. Let us show you – we have over 12,000 PDU configurations to fit every need, and over 80% of our PDUs are shipped within 10 days. Only with Server Technology will customers Stay Powered, Be Supported, and Get Ahead.
Interested in learning more about how Server Technology can help you manage and distribute power in your application? Visit us online at www.servertech.com

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

Germany Central Europe, Eastern Europe & Russia
10th + 11th Floor
Westhafen Tower
Westhafenplatz 1
60327 Frankfurt
Germany
+49 697 1045 6205 Tel
+49 697 1045 6450 Fax
salesint@servertech.com

Hong Kong APAC
Level 43, AIA Tower
183 Electric Road, North
Point, Hong Kong
+852 3975 1828 Tel
+852 3975 1800 Fax
salesint@servertech.com

India & Middle-East
Regus Municipal No. 3, Old Madras Road, Benniganahalli village,
Krishnarajpuram Hobli, RMZ Infinity, 1st floor,
Tower D, Bengaluru,
Karnataka 560016
India
+91 99022 44534 Tel
salesint@servertech.com

Server Technology®

Stay Powered Be Supported Get Ahead

REFERENCES

http://events.linuxfoundation.org/sites/events/files/eeus13_wikelius_masters.pdf

https://www.sdxcentral.com/sdn/definitions/what-is-a-data-center-fabric/

http://archive.data centerdynamics.com/focus/archive/2015/02/facebook-makes-modular-data-center-network-fabric

http://www.infoworld.com/article/3040038/data-center/what-hyperscale-storage-really-means.html

http://www.nrel.gov/docs/fy13osti/58902.pdf

http://www.techrepublic.com/article/9-iot-global-trends-for-2017/?ftag=TREe09998f&bhid=27430422563959956465665315772694

http://www.infoworld.com/article/3137968/data-center/inside-a-hyperscale-data-center-how-different-is-it.

html - tk.rss_data center

https://www.theregister.co.uk/2017/02/14/the_government_is_coming_for_your_code/

http://www.forbes.com/sites/patrickmoorhead/2013/08/21/intels-disaggregated-server-rack-is-a-mainframe-replacement/#3f35ee0b1f8f

http://www.intel.com/content/www/us/en/architecture-and-technology/rack-scale-design-overview.html

http://www.fibre-systems.com/feature/tipping-point-silicon-photonics

http://www.penguincomputing.com/solutions/fully-integrated-open-racks/

http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white_paper_c11-680202.pdf

https://en.inspur.com

APPENDIX A - DEFINING HYPERSCALE

The concept of building a hyperscale architecture is muddied by many tangential terms. In particular, we see customers confused about hyperconverged, hyperscale (or Web-scale), converged, software-defined, and commodity-based infrastructure.

Let's take a moment to clarify definitions on these ingredient terms:

- Software-defined: Infrastructure where the functionality is completely decoupled from the underlying hardware and is both extensible and programmatic. Read this post for our elaboration on software-defined storage in particular.
- Commodity-based: Infrastructure built atop commodity or industry-standard infrastructure, usually an x86 rack-mount or blade server. As we've written in the past, don't conflate commodity with cheap.
- Converged: A scale-out architecture where server, storage, network, and virtualization/containerization components are tied together as a pretested, pre-integrated solution. Components are still distinct in this architecture.
- Hyperconverged: A scale-out architecture that takes converged infrastructure one step further by combining software-defined components atop commodity hardware, packaged as a single solution -- often a single appliance. Components are no longer distinct.
- Hyperscale: A scale-out architecture that is also software-defined and commodity-based, but where the server, storage, network, and virtualization/containerization resources remain separate. Each component is distinct and can be independently scaled.

In summary, think of hyperconverged infrastructure as the modern, logical extreme of converged systems, whereas hyperscale is the modern, logical extreme of how we've been building data centers for 30 years. Both make sense for specific environments, as shown below.

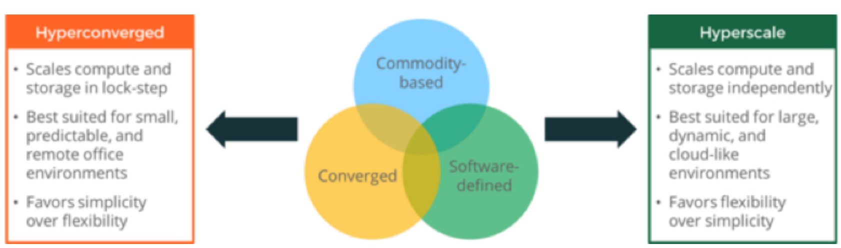

http://www.infoworld.com/article/3040038/data-center/what-hyperscale-storage-really-means.html

Based on our experience, there are some easy rules of thumb to determine which architecture is right for you.

- Choose hyperscale when… your organization has 5,000 employees or more, more than 500 terabytes of data, more than 500 applications, or more than 1,000 VMs.
- Choose hyperconverged when… you're below these watermark numbers, have five or fewer staff managing your virtual infrastructure, or you're in a remote or branch office.

http://www.infoworld.com/article/3040038/data-center/what-hyperscale-storage-really-means.html

APPENDIX B - QUESTIONS

Questions to answer when starting your next datacenter design:

- Do you have a target PUE (or another "Green" metric)?
- Costs to achieve your PUE or efficiency goal?
- Have you compared cooling methodology versus average climate at the data center location?
- What type of power is available from the utility – renewable? AC or DC?
- What are the available water supplies / chilled water sources – fresh water? Gray water?
- What renewable energy source – geothermal, water, wind, solar, biogas, and does it provide AC or DC, and cost per kWH?
- What is the typical solar exposure at the proposed facility location (how many sunny days, what angle to the sun, what solar intensity)?
- Are there tax abatements / incentives for hardware, energy and jobs created?
- Will the IT loads be virtualized / containerized?
- Will the failure of a given piece of IT equipment cause any adverse impact to the performance of the data center or can the IT load be moved with without concern? Failure in place threshold?
- Would you use a standard Basic PDU versus Bus bar?
- Would your application work better with a Managed (Switched) PDU in the Hyperscale environment?
- What is the availability of local skilled labor?