Executive Summary

Data Center Energy and Carbon Emission Reductions Through Compute, Storage, and Networking Virtualization

Sponsored by: VMware

Ashish Nadkarni
Eric Sheppard
Brad Casemore
September 2017

IDC OPINION

Digital transformation (DX) – a technology-based business strategy – is essential for firms to thrive in a digital economy. DX initiatives serve as a conduit via which firms can improve their top line – by creating differentiation in the market, extending their competitive advantage, developing new business relationships, and enhancing customer experience. Firms also utilize DX for increasing operational efficiencies, which in turn leads to an improved bottom line. And finally, DX also helps firms reduce their carbon footprint by embracing green IT technologies. Server virtualization (aka software-defined compute) – a technology that has become a mainstay of most modern data centers – ushered in an era of unprecedented operational efficiency in the data center. Virtualization increased the overall utilization of the servers by removing the disparity between CPU, memory capacity, and disk utilization. It also helped IT mitigate rampant (pre-virtualization) data center expansion and reduce server sprawl by rightsizing the compute footprint.

VMware's server virtualization solution – the first commercially available software for x86-based servers – enabled IT to accelerate the deployment of Windows- and Linux-based applications while sharply reducing spend on new servers. Reduced servers mean reduced facilities-related costs pertaining to power consumption and cooling. In other words, the more densely virtualized the compute layer, the lower the data center carbon footprint (measured in terms of carbon dioxide, or CO_2, emissions). Thanks to advances in technology, IT can now further reduce its data center carbon footprint by extending the benefits of the server virtualization layer to storage and networking:

- On the storage side, IT can reduce spend on storage arrays by deploying server-based flash along with traditional spinning media and a software-defined hyperconverged infrastructure stack such as VMware vSAN.
- On the networking side, IT can reduce spend on physical network equipment by deploying software-defined networking such as VMware NSX.

There is no doubt that server virtualization has positively impacted the industry by enabling customers to achieve a cumulative reduction in server acquisition and life-cycle costs, data center space, and operational costs. It has streamlined infrastructure management. Crucially, it has enabled IT to reduce its data center carbon footprint, initially just with compute – and now increasingly so with storage and networking. Reducing carbon footprint has a direct effect on the environment, lowering release of greenhouse gases from power generation and aiding corporate social responsibility initiatives.

METHODOLOGY

IDC used a conservative and defensible approach to calculate the power consumption avoided and the associated carbon dioxide emissions that were avoided because of the industry's use of VMware's virtualization software (referred to as a hypervisor) such as ESX and vSphere. Where possible, we used IDC published data as the basis for secondary calculations.

IDC started with published data from IDC's Worldwide Server Tracker and IDC's Worldwide Quarterly Server Virtualization Tracker. This data was used to establish the number of physical servers deployed, including those running without a hypervisor and those deployed with a VMware hypervisor product, from 2008 to 2016.

IDC then extended that data model back to 2003 to establish the number of servers deployed with VMware hypervisors during the early years of x86 virtualization adoption. By multiplying the number of servers deployed with virtualization software and the average virtual machine (VM) density, IDC established the number of virtual machines that were put into service each year from 2003 to 2016.

The resulting new deployment data was then aggregated using installed base calculations to determine cumulative server counts for deployments that were assumed to be avoided. The count of servers not deployed and the installed base totals derived from the count of servers not deployed were then multiplied by average power consumption rates to determine megawatt hours (MWh) avoided, which were also converted into metric tons of carbon dioxide (MT CO_2) emissions avoided.

IDC extended these calculations to the net displacement (reduction) of storage and networking technologies should the server virtualization technology include storage (through hyperconverged infrastructure) and networking capabilities.

A detailed methodology is included in the Appendix section at the end of this document.

SITUATION OVERVIEW

At the most basic level, server virtualization enables the reduction of the number of physical servers deployed in any IT environment. The software (i.e., the hypervisor) virtualizes the hardware layer, presenting a unique hardware instance to every operating system instance. This enables IT to consolidate multiple server operating system instances and their associated workloads on a smaller number of physical servers. VMware's virtualization software — the first commercially available software for x86 servers — enabled IT to accelerate the deployment of Windows- and Linux-based applications while sharply reducing spend on new servers.

Reducing the number of physical servers in use leads to an associated savings in data center floor space and power consumed by server installations that have been avoided. In addition, there is often a reduction in cooling system power consumption, although power consumption associated with infrastructure was not considered as part of the calculation presented in this IDC Executive Summary.

Up until a couple of years ago, a hypervisor was limited to virtualizing compute — it essentially acted as a pass-through for storage and networking functions. With advances in technology, hypervisors like VMware vSphere can manage storage and networking too — enabling a reduction of the respective infrastructure. On the storage side, IT can reduce spend on power-hungry storage arrays by deploying server-based flash storage capacity along with traditional spinning media and a software-defined hyperconverged infrastructure stack such as VMware vSAN. On the networking side, IT can reduce spend on physical network equipment by deploying software-defined networking such as VMware NSX.

Figure 1 illustrates the "before and after" scenario with virtualization:

- **Annual server instances worldwide avoided because of the use of new server shipments virtualized and annual existing installed servers newly redeployed with hypervisors (combined).** This constitutes the entry of new virtualized instances or, in this document, "servers avoided" each year. Annual servers avoided because of the use of VMware virtualization products grew from a combined total of 107,000 in 2003 to 16.2 million in 2016.
- **Installed base of servers worldwide avoided annually because of the cumulative and retroactive effects of server virtualization** grew from 107,000 in 2003 to 53.5 million in 2016. The totals shown in Figure 1 serve as the basis for calculating the amount of power and CO_2 emissions avoided because of the use of VMware virtualization products.
- **A further annual reduction in the number of infrastructure devices (storage arrays and network equipment) avoided due to hyperconverged infrastructure and software-defined networking** (see Figure 2). IDC used virtual machine-to-storage and virtual machine-to-network port density ratios to calculate the number of such devices displaced. IDC used 2014 as the starting point for this add-on, which was considered the first year with material shipments of such products. Annual infrastructure equipment avoided because of the use of VMware virtualization products grew from a combined total of 595,000 units in 2014 to 1.07 million units in 2016.

Further considerations regarding reduction of additional storage and networking infrastructure are as follows:

- The years measured for this document capture storage (through hyperconverged infrastructure) and network virtualization while they were still nascent markets. Therefore, IDC has not separated the CO_2 impact analysis of storage and networking equipment displaced. IDC estimates it to be in the range of 10% of the total annual savings as a direct result of this equipment being displaced.
- In general, reduction in the number of physical storage and network ports improves the overall efficiency of, and therefore the energy consumption in, the data center. The degree to which physical network ports, storage networking ports, and storage units and capacity can be reduced depends on the extent to which vSAN and NSX are used in the environment:
  - **NSX — network virtualization:** NSX can have an impact in three ways. First, it can reduce networking-related compute overhead because of better server utilization. Second, east-west (VM-to-VM) network traffic, which is now switched within the hypervisor (versus traversing the physical data center network), reduces the number of infrastructure switches and load balancers required. Third, it minimizes the use of security appliances for east-west traffic, as NSX does firewalling within the hypervisor. NSX has a negligible impact on the data center footprint.
  - **vSAN — hyperconverged infrastructure:** The more pervasive the vSAN deployment, the greater the avoidance of external networked storage arrays. For example, in just three short years (2014–2016), servers deployed with VMware vSAN were equivalent to almost 4.5% of all midrange networked storage system units shipped during the same period. The hyperconverged infrastructure portion of these two new markets is further along in its life cycle.

FIGURE 1

Worldwide New Server Shipments, Installed Base of Servers, and Additional Infrastructure Avoided Because of the Use of VMware Server Virtualization Software, 2003-2016

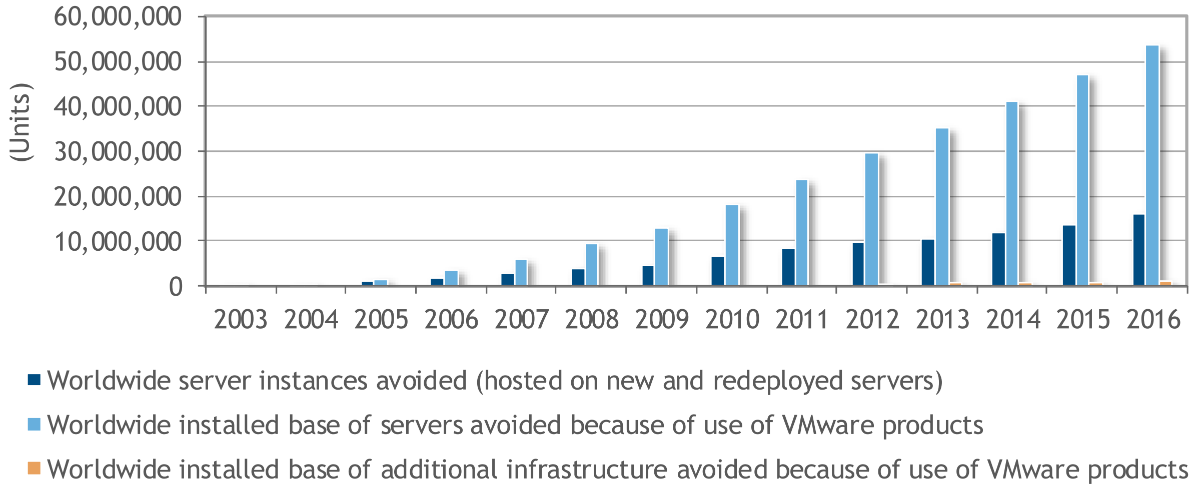

Source: IDC, 2017

FIGURE 2

Worldwide Additional Infrastructure Avoided Due to the Use of VMware Server Virtualization Software, 2012-2016

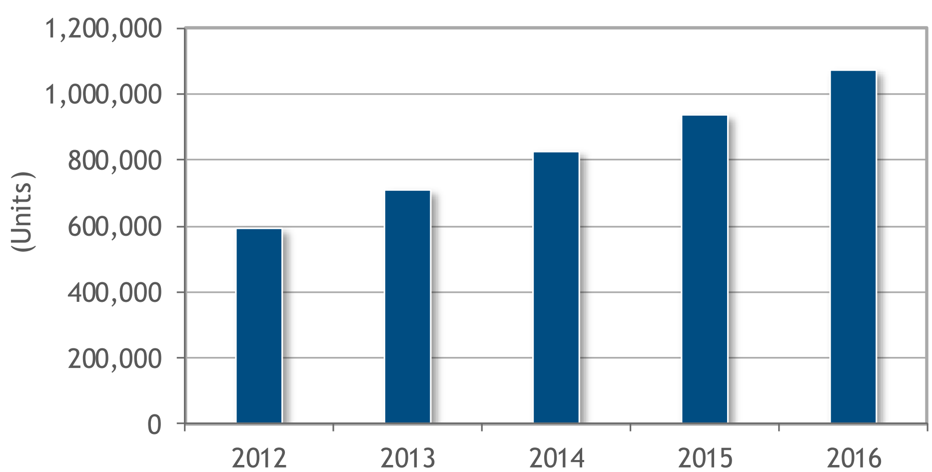

Source: IDC, 2017

Impact on Data Center CO_2 Emissions

The power consumption avoided is directly proportional to the avoidance of new physical server deployments and the corresponding workloads (and their associated storage and networking) being deployed on virtual machines hosted by a hypervisor running on existing physical servers.

Figure 1 presents the calculated number of servers that were avoided because of the use of a hypervisor from VMware (ESX, vSphere) that primarily virtualizes compute and secondarily, storage (vSAN hyperconverged infrastructure) and networking (NSX). IDC calculated the power consumption avoided using the following formula:

$$\boldsymbol{Pa = N * S * H}$$

$$Pc = \sum_{2013}^{2016} Pa$$

Where

- P_a is the annual power consumption avoided (MWh/year)
- N is the number of physical servers avoided
- S is the average power consumption of each server (MWh)
- H is the average number of hours of use per day per server
- P_c is cumulative power consumption avoided (MWh)

What is noteworthy is that the savings accrued by deploying virtualization software is orders of magnitude beyond savings via incremental improvements in server efficiency (i.e., by deploying more efficient processors or more efficient power supplies or moving from spinning media to a mix of spinning media with solid state storage).

Figure 3 shows worldwide power consumption avoidance associated with the use of VMware server virtualization products. Power avoidance grew from 191,000MWh in 2003 to over 135,000,000MWh in 2016. Cumulatively, the power consumption savings grew from 191,000MWh in 2003 to almost 739,000,000MWh in 2016. To put some perspective on that total, 739,000,000MWh of cumulative power consumption avoided is equivalent to the power consumed by 52% of total U.S. households for one calendar year (during 2016 per U.S. Census data).

Figure 4 illustrates the CO_2 emissions reduction related to the use of VMware server virtualization products. IDC calculated CO_2 emissions using the following formula:

$$\boldsymbol{Va = X * Pa}$$

$$Vc = \sum_{2013}^{2016} Va$$

Where

- V_a is the annual CO_2 emissions reduction (metric tons/year)
- X is the average amount of CO_2 produced by each server (metric tons)
- P_a is the annual power consumption avoided (MWh/year)
- V_c is the cumulative CO_2 emissions reduction (metric tons)

As noted in Figure 4, CO_2 emissions avoided related to the use of VMware server virtualization products grew from 107,000 metric tons/year in 2003 to over 76 million metric tons/year in 2016.

In terms of cumulative emissions, avoided emissions grew from 107,000 metric tons in 2003 to almost 415 million metric tons in 2016.

The CO_2 avoidance in 2016 alone is the equivalent of removing over 15 million cars from the road and the avoidance of having driven over 257 million miles in 2003 and 170 billion miles in 2016.
In cumulative terms, 415 million tons is the equivalent of having eliminated 900 billion automobile miles being driven over the past 15 years.

FIGURE 3

Worldwide Power Consumption Avoidance Associated with the Use of VMware Products, 2003–2016

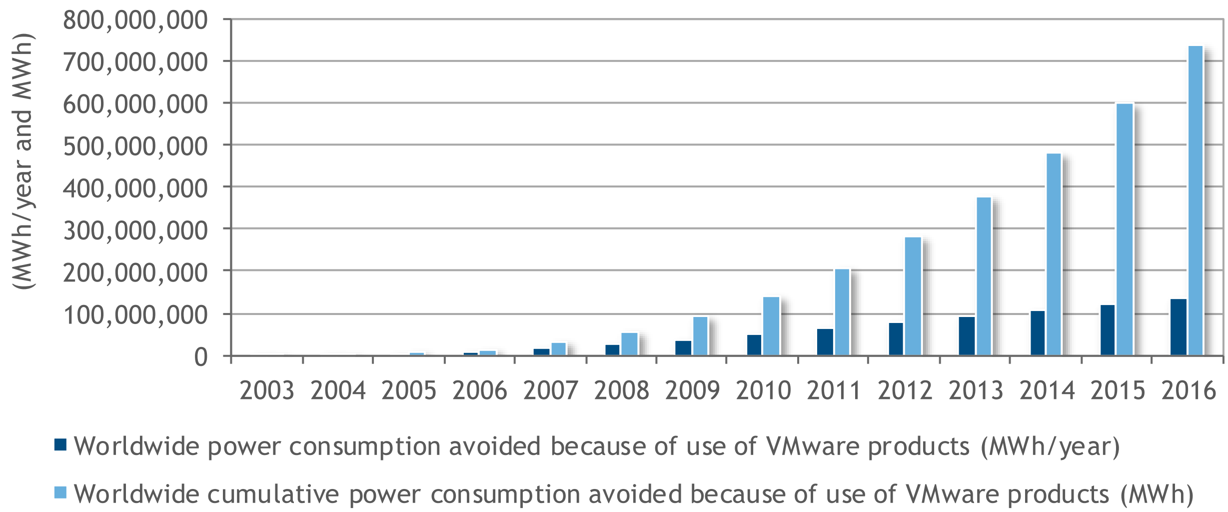

Source: IDC, 2017

FIGURE 4

Worldwide CO_2 Emissions Reduction Associated with the Use of VMware Products, 2003-2016

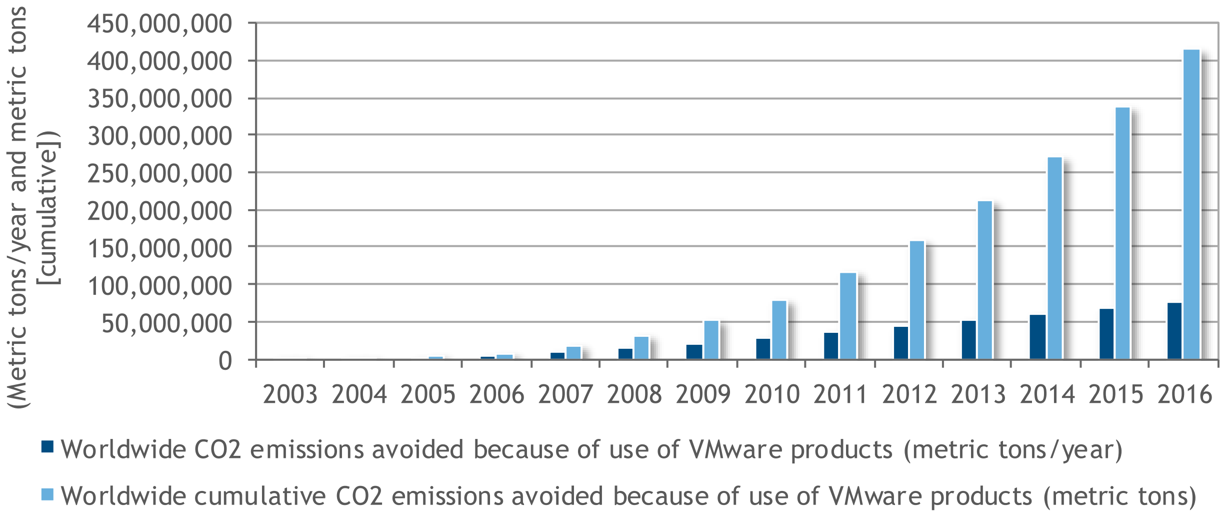

Source: IDC, 2017

CONCLUSION

IDC's analysis illustrates that virtualization software has a significant, tangible, and net positive effect on lowering the size of compute, storage (i.e., hyperconverged infrastructure), and networking infrastructure. The more pervasive the virtualization, the higher the cumulative savings.

IDC estimates that the cumulative infrastructure avoidance-related savings can quickly go north of 20-25% when firms proportionately complement their vSphere installations with vSAN and NSX. In the big picture, this enables companies to aggressively pursue green IT (i.e., do their part in reducing their data center carbon footprint).

The data presented in this IDC Executive Summary illustrates the reductions in power consumption and related CO_2 emissions directly attributed to VMware products in use worldwide. VMware is a market leader in server virtualization software, and its products have been and continue to be a major driving force in helping customers realize higher levels of operational efficiencies in their data centers – resulting in a positive benefit to the planet by reducing CO_2 emissions.

APPENDIX

Detailed Methodology

IDC used a conservative and defensible approach to calculating the power consumption avoided and the associated carbon dioxide emissions that were avoided because of the industry's use of VMware virtualization software. Where possible, we used published IDC data as the basis for the secondary calculations.

The methodology applied is as follows:

1. IDC based this model on syndicated (published) IDC data including the following dimensions:
   - IDC's Worldwide Server Tracker data provided the basis for total worldwide server shipments. IDC used data from calendar years 2003 through 2016 as a starting point for this model.
   - IDC's Worldwide Quarterly Server Virtualization Tracker (discontinued) was the basis for new server shipments virtualized using VMware virtualization. IDC included the following VMware virtualization products: VMware GSX, VMware ESX, VMware Server, and VMware vSphere. In addition, IDC included data for VMware vSAN and NSX to model the storage and networking impact. IDC's Worldwide Quarterly Server Virtualization Software Tracker provides data from 2008 through 2015, with 2016 being a modeled effort. Servers running competitive hypervisors from Microsoft, Red Hat, Citrix, and others were specifically excluded from this analysis.
2. To bridge earlier historical years that predate IDC's Worldwide Quarterly Server Virtualization Tracker, IDC applied assumptions to solve for a complete historical view for the use of VMware hypervisor products on new server shipments virtualized (new servers shipped with a hypervisor installed before shipment, during shipment, or immediately after delivery). Assumptions applied for the period from 2007 back to 2003 were as follows:
   - In 2003, the assumption applied is that VMware was the only viable x86 server virtualization technology in (relative) widespread use. IDC scaled the overall worldwide penetration of virtualization deployments on new server shipments from 18% in 2008 (reported in IDC's Worldwide Quarterly Server Virtualization Tracker) to 2% in 2003.
   - IDC scaled VMware's overall share of the worldwide total of new server shipments virtualized from 61% in 2008 to 100% in 2003. (In other words, VMware captured 100% of the 2% of new server shipments that were virtualized in 2003). In 2003, VMware was the only viable virtualization technology on x86 hardware.
   - IDC's working models behind the top-level conclusions produced here were built using detailed dimensions including product type (blade, density optimized, rack optimized, and tower) and socket count (1 socket, 2 sockets, and 4+ sockets).
3. IDC applied the same assumption to bridge from 2007 back to 2003 to develop a complete historical view for the use of VMware hypervisor products on existing installed servers.
4. Virtual machine density (VM density), as reported in IDC's Worldwide Quarterly Server Virtualization Tracker for 2008 to 2015 (modeled for 2016), was scaled back linearly to solve for historical data: The starting point for this model — 2003 — assumes two VMs/new server shipments virtualized.
5. Total instances were calculated by multiplying VM density for new server shipments virtualized and installed base deployments (individually) by their respective unit volume to come up with total instances placed into service each given year.
6. At this point in the model, IDC applied a "discount" to reduce instances from the installed base that may have been the result of virtualization sprawl.
   - The discount applied ranged from 6% of instances in 2003 to 25% in 2016. The presumption is that virtualization software, in conjunction with today's data center-oriented virtualization-friendly licensing, makes it easy (and affordable) to spin up more instances than would have happened if a hardware and software purchase was necessary for each individual instance to be created.
7. Installed base of total server instances in use each year was determined by an installed base calculation on the discounted instance total, using a mathematical formula that replicates the calculations for physical server installed base totals produced in IDC's Worldwide Server Tracker database. This calculation is done individually for new server shipment virtualized instances and

separately for instances aboard installed base servers that have had hypervisors installed and varied by form factor. Overall, the average life expectancy for servers included in this study was about 4.5 years.

- The presumption is that virtualized operating system/workload instances on a server newly deployed will have a life cycle much like that of an existing server installation (because the alternative would have been to install that instance on a dedicated server).
- Separately, IDC calculated the life expectancy/installed base for instances aboard installed servers that received a hypervisor through a redeployment midlife; instances running on these servers were assumed to have a considerably shorter life cycle than instances running on brand-new servers.
- These separate installed base calculations were combined to produce a total number of server instances that were avoided because of the use of VMware virtualization solutions.

8. The overall total installed base of servers in use were then multiplied by power consumption estimates for each server product type (blade, density optimized, rack optimized, and tower) and by socket count.
   - IDC used U.S. Department of Energy estimates to help shape the actual power consumed by servers, which tends to be roughly 70% of the rating of the power supply included with the server itself. (Source: Energy Technologies Area, Berkeley Lab)
   - Power consumption was assumed to be higher in early product years, with efficiency improving each year up to 2016. (Source: Energy Technologies Area, Berkeley Lab)
9. The results of step 8 were then multiplied by the number of hours of utilization per day those servers experienced.
   - Commercial servers commonly used in large-scale data centers and cloud environments, including blade, density-optimized, and rack-optimized form factors, were assumed to have a high level of "on time" — between 20 hours and 24 hours — because the result of shutting down servers during low-use periods can be detrimental to the balancing of cooling systems counteracting the heat exhausted from data center infrastructure.
   - Density-optimized form factors, commonly used in hyperscale data centers, were assumed to have the highest uptimes. However, density-optimized servers are not heavily used
     with VMware hypervisor products because most density-optimized servers tend to run bare metal workloads.
   - However, form factors more likely to have non–data center deployments were treated uniquely. Tower form factors are commonly used in small and medium-sized businesses (SMBs) and branch offices and were assumed to have a comparatively short daily "on time" of 12 hours per day. (Source: IDC estimates)
   - The resulting data produced watt hours of power consumed per day and per year.
10. IDC converted annual power consumption to megawatt hours annually and in turn converted MWh to equivalent CO_2 emissions associated with that power consumption.
    - CO_2 generation rate used is for the overall United States, or 1,238.516lb of CO_2/MWh of power generated. IDC recognizes that there are differences in global emissions factors and this would be an area for further study.
    - Our assumption is that emerging geographies have higher CO_2 emission rates and mature geographies (such as Western Europe) have lower CO_2 emission rates.
    - The United States accounted for about 34% of new server shipments in 2016, while Western Europe accounted for 19% of new server shipments. Asia/Pacific outside of Japan accounted

for 31% of new server shipments in 2016. IDC used the average U.S. CO_2 emission rate/MWh for the overall worldwide calculation, assuming higher emission rates (because of the use of inexpensive, high-emission fuels) in the fast-growing emerging market segments will more than offset the lower emission rates of Western Europe. Thus the U.S. average is a conservative conversion factor to use. (Source: U.S. Census Bureau)

11. The calculation for cumulative power consumption equivalency per household is as follows:
    - 915 billion pounds CO_2 (cumulative) divided by 14,020 pounds CO_2 emission from electric utilization per year per household equals 65,252,774 households' emission for *one* year.
    - 65,252,774 divided by 135,697,926 households total (per U.S. Census data, 2016) equals 48%.
12. The calculation for annual CO_2 avoidance per household is as follows:
    - Using MWh: 135,309,481MWh per year (in 2016) avoided via the use of VMware Virtualization divided by 11.320MWh per year per household (note that IDC divided kWh by 1,000 to get to MWh) equals equivalent emissions of 11,644,534 households for *one* year.
    - Alternatively, IDC could also arrive at this value using CO_2 released: 167,582,957,726 pounds of CO_2 released per year avoided via the use of VMware Virtualization (in 2016) divided by 14,020 pounds of CO_2 released per household per year equals equivalent emissions of 11,644,534 households for *one* year.
    - CO_2 pounds were converted into metric tons to be consistent with GHG Reporting Protocol standards (www.ghgprotocol.org/corporate-standard).

IDC worked with VMware to correlate these findings against internal historical data to arrive at a closer approximation of real-life savings.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com