Infrastructure Virtualization Leads the Way in Reducing the Carbon Cost of Growth

An IDC White Paper, Sponsored by VMware

Authors: Stephen Belanger, Brad Casemore, Shannon Kalvar, Eric Sheppard

IDC ANALYZE THE FUTURE

Sponsored by: VMware

Authors:
Stephen Belanger
Brad Casemore
Shannon Kalvar
Eric Sheppard

August 2018

Infrastructure Virtualization Leads the Way in Reducing the Carbon Cost of Growth

IDC OPINION

Digital transformation (DX) enables organizations to compete more effectively in the digital economy, using rich insights from data collected via newer mechanisms such as connected devices. DX is a multifaceted transformational journey and forms the underpinning on which firms can out-innovate, outthink, and outpace their competition. Firms must embrace change and transform themselves digitally and become the disruptor in their industry so that they are not disrupted by their competitors. This requires them to transform in three key areas: IT, workforce, and security. Legacy systems and operating models hinder the ability of IT organizations to support DX initiatives. IT transformation (ITX) initiatives enable IT organizations to reduce the drag created by such systems and models and take a "people, process, and technology" approach to accelerate the digital readiness of their firm.

ITX enables IT to reduce the capital and operational costs of legacy IT and shift those dollars toward technology, personnel, and process transformation by:

- **Modernizing IT infrastructure:** Refreshing the server, storage, and networking infrastructure that supports key systems of record and trust (It also includes embracing hyperconverged systems, newer software-defined networking, data protection, and security technologies.)
- **Automating IT service delivery:** Deploying a hybrid cloud that includes private and public cloud resources (It also includes deploying newer IT-centric systems of engagement and insight.)
- **Transforming IT operations:** Moving to an IT-as-a-service model with internal and external resources (It involves moving IT to an "information drives process" paradigm along with the rest of the firm.)

VMware's infrastructure virtualization solution — which encompasses compute (server), storage, networking, and management capabilities — forms the underpinning of modern data center infrastructure. It enables firms to gain data center–wide and IT-wide efficiencies as well as establish metrics to track and ultimately curb carbon emissions resulting from IT infrastructure growth.

IT transformation (ITX) initiatives enable IT organizations to reduce the drag created by such systems and models and take a "people, process, and technology" approach to accelerate the digital readiness of their firm.

A reduction in the infrastructure footprint means reduced facilities-related costs pertaining to power consumption and cooling. In other words, the more densely virtualized the compute, storage, and networking infrastructure layers, the lower their carbon impact (measured in terms of carbon dioxide, or CO_2, emissions). Reducing carbon footprint has a direct effect on the environment, lowering release of greenhouse gases from power generation and aiding corporate social responsibility initiatives. IT can now be an important contributor to organizationwide sustainability initiatives.

Methodology

This is the third iteration of a VMware-sponsored IDC initiative to measure the impact of IT infrastructure growth in terms of carbon emissions and the influence infrastructure virtualization has on emissions.

For this initiative, like the previous two initiatives, 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 ESXi and vSphere. Where possible, we used IDC published data as the basis for secondary calculations.

IDC started with published data from IDC's Worldwide Quarterly Trackers, which cover areas such as enterprise servers, enterprise storage systems, and server virtualization. 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 2017.

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 2017.

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.

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 CO2) emissions avoided.

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

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

At the most basic level, infrastructure virtualization enables the reduction of the number of physical servers, storage arrays, and networking devices deployed in any IT environment.

On the storage side, IT can reduce spend on power-hungry storage arrays by deploying server-based flash along with a software-defined hyperconverged storage stack such as VMware vSAN.

Situation Overview

At the most basic level, infrastructure virtualization enables the reduction of the number of physical servers, storage arrays, and networking devices deployed in any IT environment. The software (i.e., the hypervisor) virtualizes the hardware layers, presenting a unique hardware instance (virtual machines) to every operating system instance for compute, a built-in software-defined data persistent layer for these virtual machines (using server-based storage), and a networking subsystem (for east–west and north–south networking traffic). This enables IT to consolidate multiple server operating system instances as well as their associated storage and networking layers and the workloads hosted on them on to 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 applications while sharply reducing spend on new servers.

Reducing the number of physical servers, storage arrays, and networking devices in use leads to an associated savings in data center floor space and power consumed by IT infrastructure installations that have been avoided. In addition, there is often a reduction in cooling system power consumption, although power consumption associated with servers is the bulk of the overall calculation.

Within a few years, the role of a hypervisor has expanded beyond virtualizing compute to providing core infrastructure functions for data persistence (storage) and networking (switching, routing, and security) functions. On the storage side, IT can reduce spend on power-hungry storage arrays by deploying server-based flash along with a software-defined hyperconverged storage stack such as VMware vSAN. On the networking side, IT can reduce spend on physical network devices and appliances 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 over 14 million in 2017.
- **Installed base of servers worldwide avoided annually because of the cumulative and retroactive effects of server virtualization grew from 107,000 in 2003 to over 50 million in 2017.** 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 devices) avoided due to hyperconverged storage and software-defined networking is shown in Figure 1. 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.

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

- The years measured for this document capture storage 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.

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.

- 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 compute overhead through better server utilization by performing network segmentation (microsegmentation) in the hypervisor rather than requiring physical server segmentation. Second, its ability to switch and route east–west (VM-to-VM) network traffic within the hypervisor (rather than having to traverse the physical data center network) reduces the number of physical networking devices that are required. Third, it minimizes the use of network security appliances because NSX performs distributed firewalling within the hypervisor. NSX typically reduces the number of physical and virtual hosts enterprises need to deploy, and that has a beneficial impact on data center footprint, but that effect is mitigated on an aggregate level as a result of NSX's smaller installed base relative to vSphere.

vSAN — storage virtualization and hyperconvergence

The more pervasive the vSAN deployment, the greater the avoidance of external networked storage arrays. For example, in just a few short years (2014–2017), servers deployed with VMware vSAN were equivalent to almost 4.5% of all midrange networked storage system units shipped during the same period. The storage virtualization portion of these two new markets is further along in its life cycle.

FIGURE 1

Worldwide Infrastructure Avoided (Servers, Storage, Networking) Due to the Use of VMware Server Virtualization Software, 2003–2017

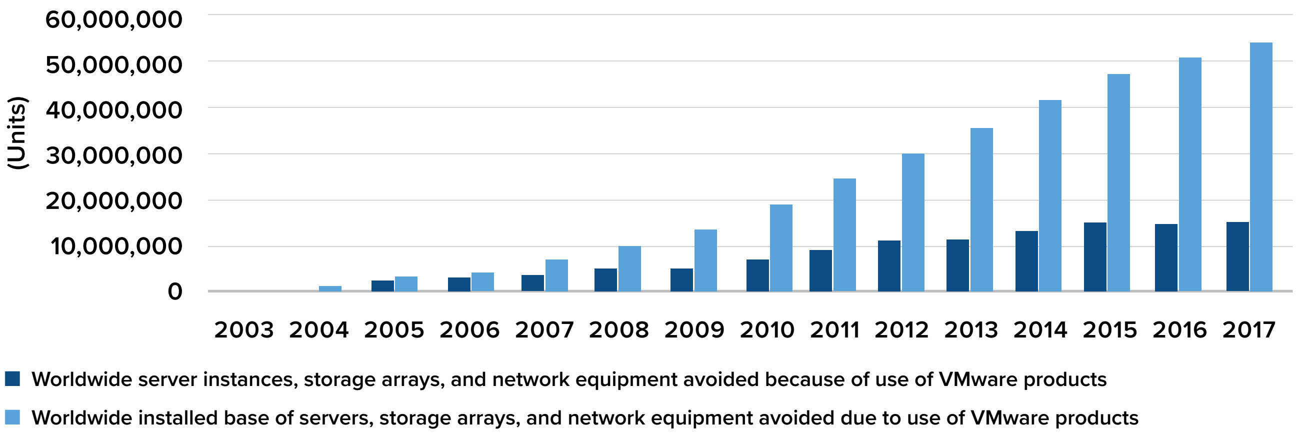

Note: There are minor changes to historical data from the prior iteration to incorporate IDC's most up-to-date published tracker data and other modeling refinements.
Source: IDC, 2018

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 (ESXi and vSphere) that primarily virtualizes compute and storage (vSAN) and secondarily networking (NSX). IDC calculated the power consumption avoided using the following formula:

$$P_a = N * S * H$$

$$P_c = \sum_{2013} P_a$$

Where

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

Figure 2 illustrates overall power consumption avoided using VMware virtualization products. 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 solid state storage).

FIGURE 2

Worldwide Power Consumption Avoidance Associated with the Use of VMware Products, 2003–2017

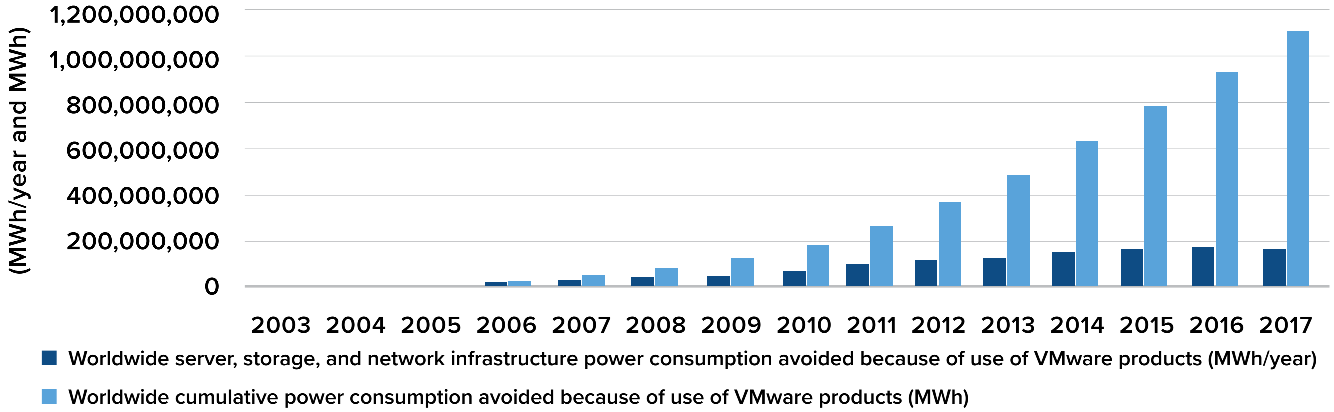

Note: There are minor changes to historical data from the prior iteration to incorporate IDC's most up-to-date published tracker data and other modeling refinements.
Source: IDC, 2018

Figure 3 illustrates the CO_2 emissions reduction related to the use of VMware virtualization products. IDC calculated CO_2 emissions using the following formula:

$$V_a = X * P_a$$

$$V_c = \sum_{2013} V_a$$

Where

- Va is the annual CO_2 emissions reduction (metric tons/year)
- X is the average amount of CO_2 produced by each server or storage array (metric tons)
- Pa is the annual power consumption avoided (MWh/year)
- Vc is the cumulative CO_2 emissions reduction (metric tons)

FIGURE 3

Worldwide CO_2 Emissions Reduction Associated with the Use of VMware Products, 2003–2017

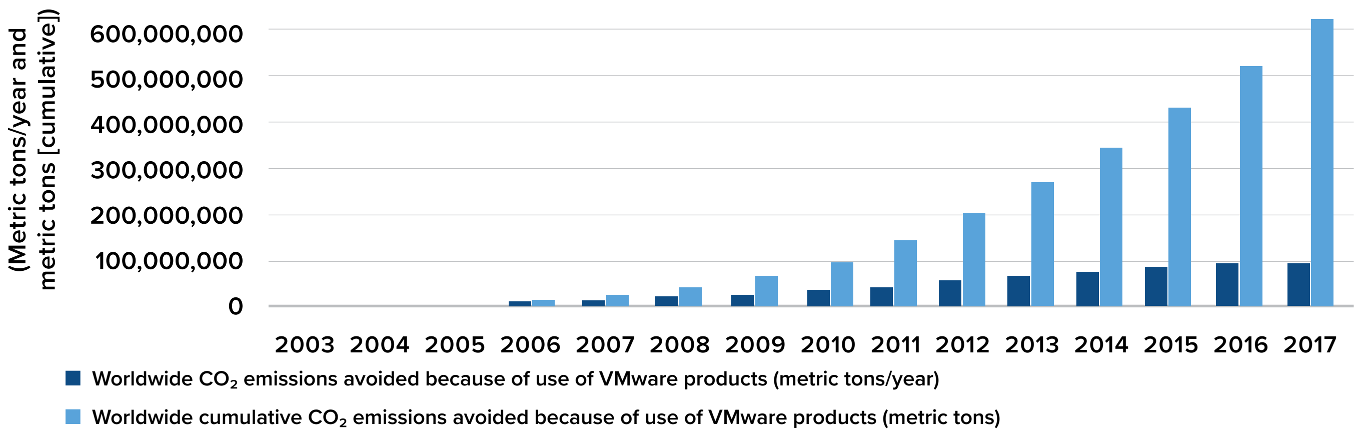

Note: There are minor changes to historical data from the prior iteration to incorporate IDC's most up-to-date published tracker data and other modeling refinements.
Source: IDC, 2018

As noted in Figure 3, CO_2 emissions avoided related to the use of VMware virtualization products grew from 112,651 metric tons/year in 2003 to over 84 million metric tons/year in 2017. In terms of cumulative emissions, avoided emissions grew from 112,651 metric tons in 2003 to over 539 million metric tons in 2017.

The CO_2 avoidance in 2017 alone is the equivalent of removing over 18 million cars from the road and the avoidance of having driven over 270 million miles in 2003 and 200 billion miles in 2017. In cumulative terms, 539 million tons is the equivalent of having eliminated over 1 trillion automobile miles being driven over the past 15 years.

Introducing CCOG — Carbon Cost of (Infrastructure) Growth

This document introduces the concept of CCOG — carbon cost of (infrastructure) growth. Future iterations of this document will seek to measure the impact of IT infrastructure in terms of carbon emissions, thereby providing a metric for organizations to measure the impact of their corporate sustainability initiatives.

Global growth is driven by demographic, geographic, economic, logistical, social, and technical factors. The "cost of growth" is a measurement of how much input in terms of capital, energy, and human time is required to achieve a certain amount of growth from the current state and then sustain that growth over time. "Carbon cost of growth" refers to carbon emissions required to instantiate and sustain a specific amount of growth in a region.

From IDC's perspective, the cost of growth must be approached as an IT infrastructure system play. Furthermore, it must be measured in terms of carbon saved because of a reduced carbon load (emissions). This approach takes the cost of each addition into the overall system and includes:

- The cost of the addition itself
- The cost to integrate the addition into the existing structure and second/third order impacts
- The cost to maintain the now more complex system over time

The "cost" of each addition may take various forms, including energy, environmental impact, manpower, material, measurement, method, or milieu. For energy or environmental impact, carbon emissions is a measurable metric.

Ideally, the carbon cost of growth should be negative over time (i.e., a lower volume of carbon emissions is required to sustain growth over a period).

Ideally, the carbon cost of growth should be negative over time (i.e., a lower volume of carbon emissions is required to sustain growth over a period). This would introduce a drag into world carbon emissions, eventually causing net-new growth to become carbon neutral. Since the Industrial Revolution, economic growth has come at a cost of increased carbon emissions. Achieving negative carbon cost of growth requires a radical rethinking of materials, energy production, energy consumption, logistics, and manufacturing approaches.

Infrastructure platforms are a pivotal factor in calculating the carbon cost of growth, specifically through "organic" productivity growth and more crucially via mortality reduction and logistical coordination of computing, storage, and networking resources using virtualization. This is not just in the IT field but also across the industrial spectrum. Accounting, communications, industrial robots, marketing, product development, sales, and even unsuspecting consumers rely upon compute to perform basic functions, connect with their peers, and create/explore new markets. This suggests, then, that there might be a "compute requirement of growth," or a ratio at which access to compute is tied to overall economic growth.

Infrastructure growth is a crucial factor in the energy cost of growth, consuming electricity at a potentially calculatable rate. Companies such as VMware focus on solutions that decouple the usage of an asset from the physical attributes of the asset itself (using virtualization) and then monitor/configure the resulting virtual asset. This enables greater and more efficient use of existing compute, storage, and networking assets; greater control over useful configurations; and a net reduction in overall demand for new assets.

The extent to which infrastructure growth can be offset through virtualization — a means to increase compute asset utilization and enable centralized pay-as-you-go-consumption — is otherwise known as cloud computing these days. IDC plans to measure the contribution of virtualization in reducing the carbon cost of growth in data centers and what it takes for enterprises to ultimately make it negative.

Conclusion

IDC's analysis illustrates that virtualization software has a significant, tangible, and net positive effect on lowering the amount of compute, storage, and networking infrastructure that is required. The more pervasive the virtualization, the higher the cumulative savings and therefore the organization's contribution to worldwide carbon reduction and sustainability initiatives.

IDC estimates that the cumulative infrastructure avoidance–related savings can quickly add up when firms proportionately complement their vSphere installations with vSAN and NSX.

IDC estimates that the cumulative infrastructure avoidance–related savings can quickly add up when firms proportionately complement their vSphere installations with vSAN and NSX. In the big picture, this enables companies to aggressively pursue green IT and 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 calculate 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 2017 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 ESXi, 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 2017 being a modeled effort. Servers running competitive hypervisors from Microsoft, Red Hat, Citrix, and others were specifically excluded from this analysis.
- IDC's Worldwide Quarterly Enterprise Storage Tracker provided the basis for total storage shipments and was used to calculate the offset of external storage capacity because of vSAN.

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 2008–2015 (modeled for 2017), 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 27% in 2017. 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 2017 (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" — from 20 hours to 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 36% of new server shipments in 2017, while Western Europe accounted for 17% of new server shipments. Asia/Pacific (excluding Japan) accounted for 34% of new server shipments in 2017. 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 power consumption equivalency per household is as follows:

- Annual CO_2 pounds avoided via the use of VMware virtualization divided by CO_2 pounds emission from electric utilization per year per household. Average CO_2 pounds emission per household ranged from 14,020 CO_2 pounds to 13,334 CO_2 pounds as households in general continued to become more efficient.
- 185.6 billion pounds CO_2 (annual emissions avoided in 2017) divided by 13,334 pounds CO_2 emission from electric utilization per year per household equals nearly 14 million households' emission for one year in 2017. Equivalency per household was calculated for each year starting in 2003, and the cumulative household equivalency was over 86,000,000 in 2017.
- 86,000,000 divided by 126,224,000 households total (per U.S. Census data, 2017) equals 68%.

12. The calculation for annual CO_2 avoidance per household is as follows:

- Annual power consumption avoided via the use of VMware virtualization divided by average MWh usage per year per household. Average MWh usage per year per household ranged from 11.320MWh to 10.776MWh as households in general continued to become more efficient.
- 150,292,559 MWh per year (in 2017) avoided via the use of VMware virtualization divided by 10.776MWh per year per household equals equivalent emissions of 13.3 million CO_2 pounds for *one* year in 2017.
- 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.

IDC Global Headquarters

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

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.