



DPU POWER EFFICIENCY

White Paper

Table of Contents

Increasing Data Center Power Efficiency with the NVIDIA BlueField DPU	4
The Data Center Revolution	4
Data Centers Consume Increasing Amounts of Power	5
Strategies for Reducing Data Center Power Costs	6
Increasing Server Efficiency with Domain-Specific Processors.....	8
Testing Hardware-Accelerated DPU Offloads to Reduce Server Power Consumption..	8
5G User Plane Function Network Offload Combined with CPU Microsleeps.....	9
OVS Network Offload using BlueField-2 DPU.....	10
IPsec Security Offload to BlueField DPU Saves up to 247W Per Server	12
Networking Offload with VMware vSphere and Redis Key Value Store.....	14
Network Offload Power Savings from Watts to Dollars	16
Higher Energy Prices Mean Larger Savings from DPU Offloads	17
Additional Cost Savings from DPU Offloads and Power Efficiency	18
TCO Savings from DPU Offloads	21
DPU Offloads Are the Next Wave for Improving Data Center Efficiency	22

List of Figures

Figure 1:	Data centers forecast to consume from 3% (best case) to 13% (worst case) of global electricity demand by 2030 (Source: Huawei 2015 paper by Anders S.G. Andrae.)	5
Figure 2:	Data Center PUE has improved dramatically since 2007 but has plateaued in recent years. (Source: Uptime Institute 2021 survey.)	7
Figure 3:	CPU power savings from CPU efficiencies and network offloads for a telco UPF workload at idle, 50% load, and 100% load.	9
Figure 4:	The BlueField DPU OVS offload reduced power consumption by 127 Watts or 29% compared to running OVS in the kernel on the CPU.	11
Figures 5 and 6:	Testing of IPsec encryption and decryption for both server and client showed that offloading to the BlueField-2 hardware provides substantial power savings over running IPsec entirely in software.	13
Figure 7:	Testing of VMware vSphere Distributed Services Engine with a Redis workload on a 25Gb/s network showed offloading networking to the DPU saved 12 cores while delivering faster Redis performance.	15
Figure 8:	Average power prices in Italy, France and Germany doubled (or more) during 2022 due to the war in Ukraine. (Source: Axios chart based on Rystad Energy data.)	18
Figure 9:	A typical 2007 breakdown of data center energy usage with only 50% of electricity used to power IT equipment and a PUE of 2.0 (Source: US EPA Energy Star Report to Congress on data center efficiency, 2 August 2007).	19
Figure 10:	Typical 2014 US data center energy use breakdown with 57% of power used for IT equipment and only 43% used for cooling, power provisioning, lighting, and other purposes, leading to an improved PUE of 1.75. (Source: "United States Data Center Energy Usage Report" by Sehabi, et.al. of Lawrence Berkeley National Laboratory, published December 2016.)	19

List of Tables

Table 1.	CPU power Savings from CPU efficiencies (micro-sleep and frequency scaling) and networking offloads for a telco UPF workload, as published by Ericsson.	10
Table 2.	Power savings from OVS offload to the BlueField DPU for a wireless telco workload.	11
Table 3.	Power savings from offloading IPsec encryption from the CPU to a BlueField DPU.	14
Table 4.	TCO calculation for offloading VMware ESX networking to a BlueField DPU, for a Redis workload initially running on 10,000 servers and supporting 14 million Redis transactions per second per server.	16
Table 5.	Higher power costs mean larger savings from DPU offloads. Estimated cost savings by using DPU offloads to reduce power consumption by 200W per server for 10,000 servers over 3 years, in various countries and US states.	17
Table 6.	The decline in typical data center PUE over time and the total power savings at different PUE levels resulting from a 100 Watt power savings on a server	20
Table 7.	TCO calculation from offloading IPsec encryption/decryption to a BlueField DPU, for a large data center with 10,000 servers.	21

Increasing Data Center Power Efficiency with the NVIDIA BlueField DPU

Power efficiency is increasingly important in data centers due to rising costs and constraints on power capacity. One of the best ways to improve efficiency is to use a Data Processing Unit (DPU) or SmartNIC to offload and accelerate networking, security, storage or other infrastructure functions and control-plane applications, which reduces server power consumption up to 30%. The amount of power savings increases as server load increases and can easily save \$5.0 million in electricity costs for a large data center with 10,000 servers over the 3-year lifespan of the servers, plus additional savings in cooling, power delivery, rack space, and server capital costs.

Learn why power efficiency is increasingly an important metric for IT and the six main strategies for improving it in the data center. See multiple examples of how the [BlueField DPU](#) acceleration and offloads cut power consumption and reduce both CapEx and OpEx to deliver a lower TCO.


The Data Center Revolution

Power efficiency was not initially a focus for data centers. Other criteria such as maximizing compute density, accelerating time to market, and achieving high availability levels with redundant systems were the highest priorities. Now that most data centers can be brought online rapidly and offer high levels of availability and compute density, improving power consumption and reducing associated power costs have become top goals both for optimizing existing data centers and designing new ones.

Data Centers Consume Increasing Amounts of Power

It's currently estimated that data centers consume just over 1% of worldwide electricity production, 1.8% of US electricity production,¹ and 2.7% of European electricity production.² This is projected to rise to 8% (likely estimate) or as much as 13% (worst case projection) of worldwide electricity production by 2030, according to a Huawei study published in 2015.³

Figure 1: Data centers are forecast to consume from 3% (best case) to 13% (worst case) of global electricity demand by 2030 (Source: Huawei 2015 paper by Anders S.G. Andrae.)



¹ ["United States Data Center Energy Usage Report"](#) by Arman Shehabi, Sarah Josephine Smith, et.al. of the Berkeley Lab, published June 2016.

² "Energy-efficient Cloud Computing Technologies and Policies for an Eco-Friendly Cloud Market" [European Commission Report](#) by Environment Agency Austria and Borderstep Institute, published 9 November 2020.

³ "On Global Electricity Usage of Communication Technology: Trends to 2030" by Anders S.G. Andrae and Tomas Edler of Huawei Technologies Sweden AB, published 30 April 2015

Recently, electricity prices have been soaring worldwide due to high demand, supply chain constraints, and geopolitical disruptions to oil and gas supplies. In addition, many data centers face hard limits on how much power can be supplied—that is, no additional power can be provisioned to a data center at any price, even if the owners are willing to pay. This creates a powerful incentive to increase efficiency to allow more applications, tenants, and productivity from the same data center without needing to build or lease another one.

Cloud service providers are also squeezed between constant competitive pressure to reduce hourly rental costs for Software, Platforms, and Infrastructure-as-a-Service (SaaS, PaaS, IaaS) and the rising costs of electricity in many locations. And organizations of all types—service providers, enterprises, and government agencies—face pressure to increase efficiency to combat climate change, often by purchasing more expensive “green” power rather than power generated from burning fossil fuels.

With the pressure to save money, maximize utilization of each data center, and to fight climate change, every data center operator is on a quest to improve server power efficiency.

Strategies for Reducing Data Center Power Costs

Data center operators typically pursue several strategies to reduce power consumption and costs:

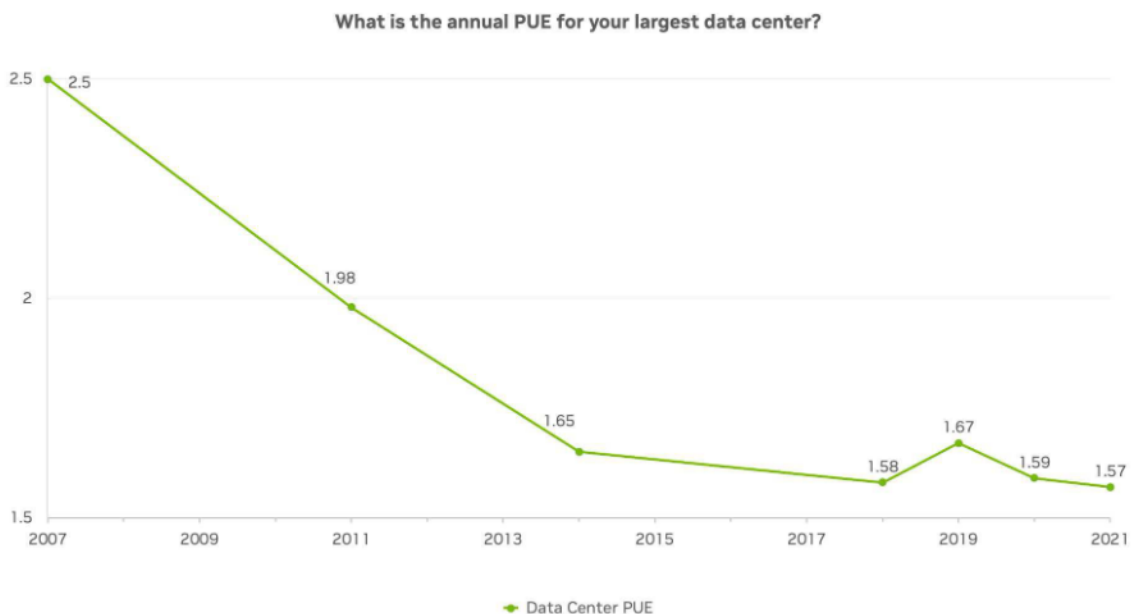
1. Build new data centers in locations that have less expensive and/or more plentiful electricity supplies, which reduces power costs but not power consumption.
2. Improve Power Usage Effectiveness (PUE), which means maximizing the percentage of power delivered to the data center that is used for actual computing equipment. However, this does not itself measure nor improve the power efficiency of individual servers.
3. Improve cooling efficiency, since cooling can account for up to 30% of data center power usage. This can include liquid cooling of components, separating hot and cold aisles, allowing higher data center temperatures, and locating data centers in cooler climates that can utilize free air cooling and/or recycle excess data center-generated heat. This reduces the electricity needed for cooling (and decreases the PUE ratio).
4. Maximize the workload utilization on each server using virtualization, containers, and composable infrastructure.
5. Improve the power efficiency of individual servers.
6. Outsource selected IT workloads to the public cloud or to colocation centers that have lower power costs. However not all workloads can be outsourced, and public cloud power efficiencies are not always passed on to tenants.

The first three strategies reduce power consumption and/or costs but do not improve the efficiency of the servers. All data centers should strive for lower PUE ratios, which divide the amount of electricity going into the data center by the amount going to the

servers themselves. An inefficient data center might have a PUE of 2 or higher (meaning only half the power delivered to the data center actually powers computing equipment), and the most efficient hyperscaler data centers claim a PUE from 1.12 to 1.25.

PUE can be improved with more efficient power distribution equipment and cooling arrangements, such as using external air, raising the data center temperatures, and employing variable-speed fans. However, PUE does not account for server efficiency, and the average data center PUE has plateaued in the last few years because most data centers have already taken advantage of the easiest ways to improve cooling and power distribution efficiency.

Figure 2: Data Center PUE has improved dramatically since 2007 but has plateaued in recent years. (Source: [Uptime Institute 2021 survey](#).)



The 4th strategy—maximizing utilization of each server—reduces the number of servers needed and increases the efficiency of each server, but it might also increase the power consumption of each server. The 5th strategy seeks ways to reduce the power consumption of each server. Networking offloads help improve these metrics: higher server utilization (strategy #4) and better server efficiency (strategy #5).

Outsourcing applications from an on-premises data center to large cloud service providers saves electricity because large cloud providers build very power-efficient servers and data centers. In fact, most of the larger ones already use SmartNICs and DPUs to perform networking and storage offloads. However, not all workloads can be moved to the public cloud and enterprises are repatriating some workloads from the public cloud either because of data privacy and sovereignty concerns or because they can run those workloads less expensively on-prem. Also, it's not clear if all the power efficiency savings of large cloud service providers are passed on to their customers as cost savings.

Increasing Server Efficiency with Domain-Specific Processors

Today's data centers rely on software-defined infrastructures to provide flexibility, scalability, and ease of management. In typical servers, virtualization, networking, storage, security, management, and provisioning are all handled by VMs, containers, or agents running on the servers' main CPUs. Not only does this consume up to 30% of the processor cycles, but a CPU is not efficient in running these types of infrastructure workloads. General-purpose CPUs excel at general-purpose, single-threaded workloads and are typically optimized for performance, not power efficiency.

Using domain-specific accelerators improves performance, reduces power consumption, and frees up server CPU cores to run the types of applications they do best. For example, a GPU can take over parallelizable tasks, especially those that are math-intensive, graphics-oriented, or AI-focused, making the server faster and more power efficient—completing more compute tasks per Watt of power used than the regular CPU would.

Similarly, a DPU runs data center infrastructure tasks much more efficiently than a general-purpose CPU. DPUs have dedicated hardware engines for accelerating networking, data encryption/decryption, key management, storage virtualization, and other tasks. In addition, the CPU cores on a DPU tend to be more power efficient than general server CPUs and have direct access to the networking pipeline. So even when a networking task cannot be accelerated by DPU domain-specific silicon, the DPU cores can perform SDN, telemetry, deep packet inspection, or other networking tasks more efficiently than the server CPU.

Testing Hardware-Accelerated DPU Offloads to Reduce Server Power Consumption

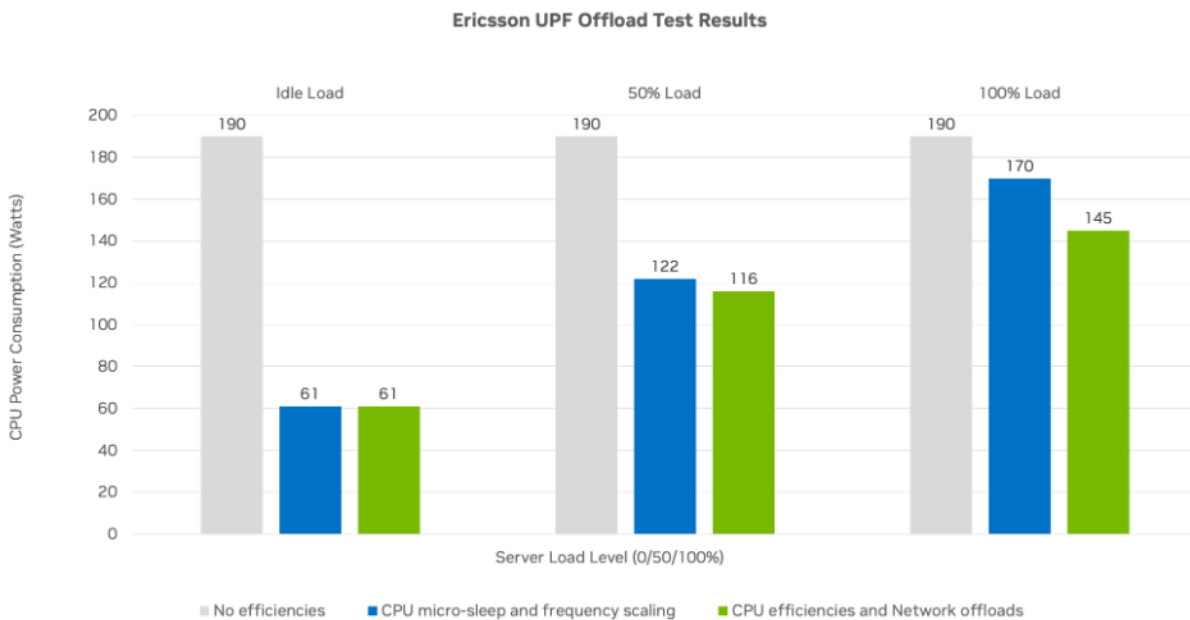
Significant efficiency gains are delivered through offloading networking and security tasks to the NVIDIA BlueField DPU. It includes domain-specific accelerators that excel at handling infrastructure processing, including networking, storage, and security. The DPU also features Arm CPU cores that are more energy-efficient than x86 CPUs and ideal for offloading the control plane of infrastructure applications and/or parts of the hypervisor or container management software.

NVIDIA worked with key partners to test the power savings from offloading various tasks to the SmartNIC or DPU in multiple tests.

5G User Plane Function Network Offload Combined with CPU Microsleeps

Ericsson, a leading telecommunications equipment manufacturer, tested a 5G User Plane Function (UPF) on servers with and without network offload to an NVIDIA ConnectX-6 Dx SmartNIC.⁴ They also tested the savings from enabling CPU microsleeps (CPU sleeps when not busy) and frequency scaling (CPU adjusts its frequency based on the workload level). They found that the CPU micro-sleeps and frequency scaling provided greater power savings when the workload was lower, and the ConnectX-6 network offloads provided the greatest power savings when the workload was heaviest.

Figure 3: CPU power savings from CPU efficiencies and network offloads for a telco UPF workload at idle, 50% load, and 100% load.



The study showed that the network offload provided up to 23% power savings (45 Watts per CPU) when the server was at 100% load. The CPU power management provided only limited power savings at 100% load. This is because the ConnectX SmartNIC offers increasing power savings and efficiency gains at higher workloads, when micro-sleep and frequency scaling are ineffective.

⁴ Ericsson power efficiency paper found at this [web page](#)

Table 1. CPU power Savings from CPU efficiencies (micro-sleep and frequency scaling) and networking offloads for a telco UPF workload, as published by Ericsson.


Telco UPF Server at 100% load	Power Use (Savings) per CPU	PowerCost for 10,000 servers over 3 years (at \$0.15/kWh)
No CPU efficiencies or offload	190W	\$7.49M
CPU micro-sleep & freq. scaling	170W (20W or 10.5% savings)	\$6.70M (\$790,000 savings)
Network offload over CPU effic.	145W (25W or 14.7% savings)	\$5.72M (\$980,000 savings)
CPU effic. and network offload	145W (45W or 23.7% savings)	\$5.72M (\$1.77M savings)

The NVIDIA BlueField-2 DPU embeds a ConnectX-6 Dx SmartNIC and would provide similar power savings as the ConnectX-6 Dx, or even greater power savings if used to offload additional control plane or management plane functions. This type of power savings from network offload is beneficial because many enterprise data centers have already enabled some form of CPU micro-sleep and frequency scaling. Still, very few enterprise data centers have already implemented hardware-accelerated network offloads.

OVS Network Offload using BlueField-2 DPU

A second example of power savings is illustrated by the work NVIDIA demonstrated with an Open vSwitch (OVS) networking offload for a major North American wireless carrier. OVS is a common open-source tool used for software-defined networking (SDN). Normally, OVS runs on the server’s x86 CPU as OS kernel software but can be offloaded to the networking accelerator on the BlueField DPU. In this test, the workload varied from 0% to 100%, and NVIDIA compared the power consumption between running OVS in the kernel (on the CPU) and offloading OVS to the DPU.

Figure 4: The BlueField DPU OVS offload reduced power consumption by 127 Watts or 29% compared to running OVS in the kernel on the CPU.



The DPU offload reduced power consumption up to 29% (127W) at 100% workload because the BlueField DPU is both faster and more power efficient than the x86 CPU at processing OVS SDN tasks.

Table 2. Power savings from OVS offload to the BlueField DPU for a wireless telco workload.⁵

OVS offload to BlueField-2 at 100% load	Power Use Per server	Power cost for 10,000 servers over 3 years (at \$0.15/kWh)
OVS in Kernel on the CPU	432 Watts	\$16.9M
OVS offloaded to BlueField DPU	305 Watts	\$11.9M
Power Savings	127 Watts (29%)	\$5.0M (savings)

⁵ Dell PowerEdge R740 server with 2 Intel Xeon Gold 6248 “Cascade Lake” CPUs @2.50GHz (40 physical cores, 80 HT cores) and Red Hat 8.3 KVM. BlueField-2 DPU card with 2x25GbE ports.

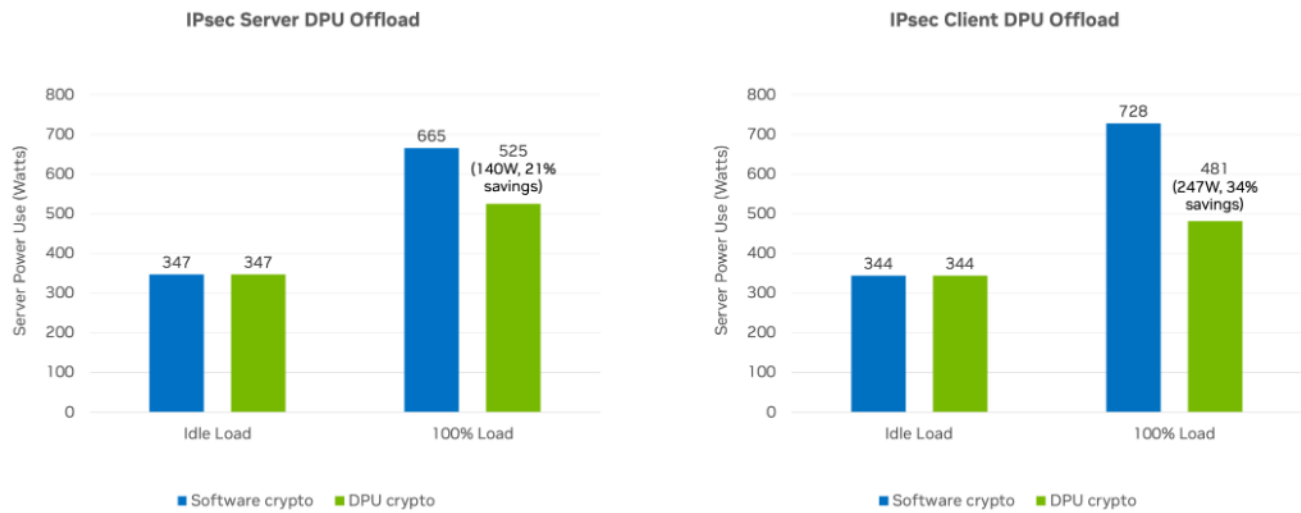
The server with DPU offloads delivered more than twice the networking throughput (49.3Gbps versus 19.8Gbps) while freeing up 18 virtual CPU cores that were required to run OVS in the kernel. In addition, at a 30% load, testing showed that offloading OVS to the DPU resulted in an average latency of 5X lower (5.5µs versus 31.5µs) than running OVS in the CPU. The power cost savings table above assumes the same number of servers would be deployed with or without a DPU. In reality, with the faster OVS network throughput, lower latency, and CPU cores saved, many fewer servers would be required with the DPU offloads, resulting in much larger cost savings than shown in the table.

IPsec Security Offload to BlueField DPU Saves up to 247W Per Server

In a third test, NVIDIA compared the power consumption while running network traffic encrypted with the popular IPsec algorithm. IPsec encrypts (and decrypts) network traffic at Layer 3 of the 7-layer OSI model - the IP protocol layer. Using encryption on every network link is increasingly a requirement in data centers because it is part of a zero-trust security stance, protecting data even if some other server or application in the data center is compromised by a cybersecurity adversary. IPsec is the most popular tool for encrypting non-web intra and inter-data center traffic between servers or between clients and servers, but at high network speeds, it places a heavy burden on the CPU. A SmartNIC or DPU with appropriate offloads can perform IPsec encryption and decryption faster and with less power consumption than a standard CPU.

Comparing the power consumption of IPsec-encrypted traffic using the CPU (IPsec in software) versus the DPU (IPsec accelerated in DPU hardware) showed a 21% power savings (140W per server) for the server and 34% power savings (247W per server) for the client.

Figures 5 and 6: Testing of IPsec encryption and decryption showed that offloading to the BlueField-2 hardware provides substantial power savings over running IPsec entirely in software. The power savings was significant on both the IPsec client and IPsec server.



As expected, the power savings was minimal when the server was at 0% load (no load = no encrypted traffic = nothing to offload) and very high at 100% load (lots of encryption work to offload). This means that maximum power savings from the DPU and server efficiency is achieved by running each server as close to 100% load as possible, which aligns with a general data center strategy for minimizing the number of servers needed and maximizing the benefits of offload using hardware acceleration.

Table 3. Power savings from offloading IPsec encryption from the CPU to a BlueField DPU.⁶

IPsec offload to BlueField-2 DPU at 100% load	Power Use (savings) Per server	Power cost for 10,000 servers over 3 years (at \$0.15/kWh)
IPsec Server, software crypto	665W	\$26.2M
IPsec Server DPU crypto offload	525W (140W, 21% savings)	\$20.7M (\$5.5M savings)
IPsec Client, software crypto	728W	\$28.7M
IPsec Client, DPU crypto offload	481W (247W, 34% savings)	\$19.0M (\$8.7M savings)


It’s worth noting that many implementations of IPsec that were known to be required in advance use a dedicated cryptographic accelerator card, or optimized CPU libraries. However, offloading cryptography to the DPU eliminates the need for a dedicated crypto card and still provides power efficiency benefits over using optimized CPU libraries. It’s common in today’s world for new encryption requirements (such as IPsec or TLS) to be imposed on existing applications and servers due to new cyber threats and/or stricter data privacy regulations. Deploying servers with NVIDIA BlueField DPUs enables customers to turn on network or storage encryption offloads if and when needed, without consuming large numbers of CPU cores.

Networking Offload with VMware vSphere and Redis Key Value Store

VMware and NVIDIA tested the performance of a Redis key value store on servers both with the VMware ESXi networking running on the server CPU and with ESXi networking offloaded to a BlueField DPU. With 36 Redis streams running, the server with the DPU offload achieved slightly faster performance (3.5% more transactions per second) and freed up 12 CPU cores to run Redis or other applications. At 36 Redis streams, the DPU offloads saved 18% of the 64 previously busy CPU cores.

⁶ IPsec server and IPsec client machines each with 2 Intel Xeon Platinum 8380 “Ice Lake” CPUs @2.30GHz (80 physical/160 HT cores total) running RHEL 8.3. DPU is BlueField-2 VPI card with 2x100GbE/EDR ports, cryptography enabled, and 16GB DRAM running Ubuntu 20.04.

Figure 7: Testing of VMware vSphere Distributed Services Engine with a Redis workload on a 25Gb/s network showed offloading networking to the DPU saved 12 cores while delivering faster Redis performance.⁷



Due to the CPU cores saved, the customer can deploy fewer servers to support the same workload. This results in both a CapEx savings from purchasing fewer servers and paying for less data center infrastructure, and an OpEx savings because the reduced number of servers consume less electricity (and related power distribution and cooling electricity). For a Redis-on-VMware deployment that initially required 10,000 servers, a TCO analysis shows the BlueField offload reduces the number of required servers by 15% and saves \$56.5M over a 3-year period.

For the power savings calculation, we make the simplistic assumption that the server with the DPU uses an extra 65W per server and thus costs an extra \$500 to operate for three years. In reality, it's quite possible the servers with a DPU would actually consume less power than the standard servers due to the highly efficient networking features of the BlueField DPU, resulting in even larger OpEx savings in the TCO.

⁷ Server with 2 Intel "Ice Lake" Xeon Platinum 8380 CPUs @2.30GHz with 80 physical cores (40 cores per socket) and 1TB of DRAM. BlueField-2 DPU card with 2x25GbE ports and 16GB DRAM.

Table 4. TCO calculation for offloading VMware ESX networking to a BlueField DPU, for a Redis workload initially running on 10,000 servers and supporting 14 million Redis transactions per second per server. ⁸

Redis on VMware vSphere DSE	Without DPU	With BlueField DPU Offload
Servers needed	10,000	8,500
Cost per server (HW+SW)	\$51,071 (no DPU)	\$53,911 (\$2,840 higher)
Total Server CapEx	\$510.7M	\$458.2M (\$52.5M savings)
OpEx per server, 3-years	\$5,500	\$6,000 (\$500 higher)
Total OpEx, 3 years	\$55.0M	\$51.0M (\$4M savings)
3-year TCO	\$565.7M	\$509.2M (\$56.5M / 10% savings)

Network Offload Power Savings from Watts to Dollars

We see that offloading networking tasks to a BlueField DPU reduces power consumption by as much as 34% or up to 247 Watts per server. We also see that the more highly-utilized a server is, the greater the power savings from hardware-based network offloads. The value of the power savings depends on the local cost of electricity and the PUE ratio. Additional CapEx and power savings may be realized if the DPU offload and acceleration results in fewer servers being needed to support the workload.

Assuming a power savings of 247W per server that runs 24x7x365, and a power cost of 15 cents per kWh (actual electricity costs in the United States vary from 8 cents to 23 cents per kWh, and from 12 cents to 38 cents per kWh in Western Europe), this means an operational cost savings of \$3.25 Million per year for a large data center with 10,000 servers $[0.247 \text{ kW} * (24 \text{ hours/day} * 365 \text{ days/year}) * \$0.15/\text{kWh} * 10,000 \text{ servers}]$. This not only provides a significant cost savings; it may allow for more servers to be installed in a power-constrained data center and improves the green credentials of both the data center and its tenants.

⁸ Cost per server with DPU reflects typical street price of a BlueField-2 DPU E-series Ethernet card with 2x25GbE network ports and 16GB DRAM, minus the typical street price of a SmartNIC with 2x25GbE ports, each configured in a server configured for and bundled with VMware vSphere DSE, including necessary software licenses.

Higher Energy Prices Mean Larger Savings from DPU Offloads

Power savings from DPU offloads are set to increase as energy prices rise. Many developed countries have significantly higher prices than the \$0.15/kWh used in our examples above. Here are typical electricity prices for various countries and some of the US States as of December 2021.


Table 5. Higher power costs mean larger savings from DPU offloads. Estimated cost savings by using DPU offloads to reduce power consumption by 200W per server for 10,000 servers over 3 years, in various countries and US states based on 2020/2021 power prices.

Country	Electricity cost per kWh (Dec 2021) ⁹	Value of 200W savings/server in 10,000 servers over 3 years (does not count cooling savings)
Baseline for this paper	\$0.150	\$7.89M
China	\$0.076	\$3.99M
USA Texas	\$0.084	\$4.4M
Mexico	\$0.085	\$4.47M
USA Oregon	\$0.088	\$4.63M
South Korea	\$0.100	\$5.26M
USA Florida	\$0.101	\$5.31M
Israel	\$0.157	\$8.25M
USA average	\$0.162	\$8.51M
France	\$0.177	\$9.30M
USA California	\$0.180	\$9.46M
Brazil	\$0.185	\$9.72M
Sweden	\$0.205	\$10.77M
Japan	\$0.214	\$11.25M
United Kingdom	\$0.298	\$15.66M
Spain	\$0.306	\$16.08M
Germany	\$0.320	\$16.82M

⁹ Electricity prices as of December 2021 according to the website GlobalPetrolPrices.com for countries and from the US Energy Information Agency (EIA) 2020 stats for US States (published Nov 2021). Does not include 2022 energy price increases due to conflict in Ukraine.

Due to pandemic-related shortages and the Russian invasion of Ukraine, energy prices have risen, especially in Europe. Prices have roughly doubled from the numbers shown in the chart above, which means the value of TCO savings from DPU offloads has also doubled. In addition, in some areas there are legal requirements or social expectations to use more “green” power generated from renewable sources, and this can also be expected to increase electricity costs over the next few years.

Figure 8: Average power prices in Italy, France and Germany have doubled (or more) from the beginning of 2022 due to the war in Ukraine. (Source: [Axios](#) chart based on Rystad Energy data.)



Additional Cost Savings from DPU Offloads and Power Efficiency

There are additional ancillary cost savings from reducing server power consumption because that also reduces the power distribution and cooling costs. Every Watt of power going into a data center requires power management hardware such as uninterruptible power supplies, generators (plus fuel delivery and storage), and power distribution units. Every Watt consumed also generates heat that must be extracted from the data center. Up to 40% of data center power consumption is dedicated to cooling. So, every Watt of electricity saved on the server side reduces both the amount of power needed to cool the data center and the amount of capital equipment needed to supply the power and remove the heat.

Figure 9: A typical breakdown of data center energy usage for a typical 2007 data center with only 50% of electricity used to power IT equipment and a PUE of 2.0 (Source: US EPA [Energy Star Report to Congress on data center efficiency](#), 2 August 2007).


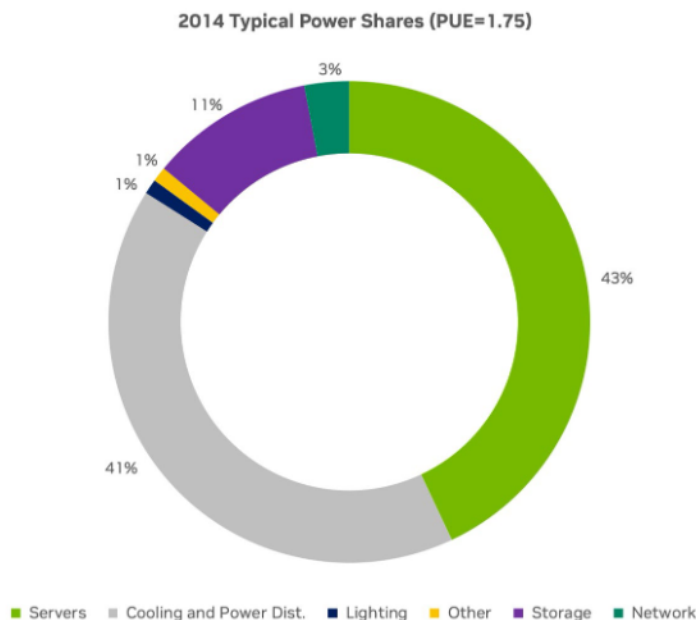


Figure 10: Typical 2014 US data center energy use breakdown with 57% of power used for IT equipment and only 43% used for cooling, power provisioning, lighting, and other purposes, leading to an improved PUE of 1.75. (Source: “United States Data Center Energy Usage Report” by Sehabi, et.al. of Lawrence Berkeley National Laboratory, published December 2016.)



The pie charts in Figures 9 and 10 show how reducing the percentage of power for cooling, power provisioning equipment, and other uses caused typical US data center PUEs to improve from 2.0 in 2007 to 1.75 in 2014. By 2020, the average data center PUE had improved (fallen) further to 1.57, which would mean a typical power consumption breakdown of 29% for Cooling/HVAC, 5% for Electrical equipment, and 2% for Lighting/Other, leaving 64% for IT load (servers, storage, and networking). With a PUE of 2.0, every Watt of power saved on servers or networking equipment equals two Watts of power saved overall. With a more efficient PUE of 1.5, every Watt saved at the server level still results in 1.5 Watts saved overall.

Table 6. The decline in typical data center PUE over time and the total power savings at different PUE levels resulting from a 100 Watt power savings on a server

Data Center PUE Rating	Power Saved at Server	Total Power Savings
PUE of 2.0 (typical in 2007)	100W per server	200W per server
PUE of 1.75 (typical in 2014)	100W per server	175W per server
PUE of 1.5 (typical in 2020)	100W per server	150W per server
PUE of 1.2 (2021 hyperscaler)	100W per server	120W per server

The higher the PUE, the more power savings at the server level are multiplied. A PUE of 2.0 means 2 Watts of power goes into the data center for every Watt that goes to a server, so saving 100 Watts at the server saves 200 Watts overall. This reflects the operational cost (and power losses) of operating the power delivery, fans and air cooling equipment plus miscellaneous OpEx overhead such as lighting and cameras. It's a testimony to the efficiency benefits of DPUs that the hyperscalers have largely implemented DPU offload technology on all of their newer servers, even though they typically pay lower prices for electricity and have a smaller multiplier on server-level power efficiency gains due to their lower PUE ratios. Enterprises, which are just starting to implement DPU offloads, stand to save much more power per server than hyperscalers due to enterprises' higher PUE ratios and higher average electricity costs.

It is important to keep in mind that PUE does not measure server efficiency or power consumed per delivered application. Once servers have been virtualized and the "low hanging fruit" of power delivery and cooling efficiency has been harvested to improve PUE to a good level (currently about 1.5 for enterprise data centers), the best place to seek additional power savings is by using offloads within the servers themselves.

There will typically also be additional savings from the ability to run more revenue-generating workloads as well on each server thanks to the CPU cycles freed up by the networking offload. Deploying DPU offloads in servers usually allows each server to perform more work (more connections, more virtual machines, more users, etc.). This results in a large CapEx savings because fewer servers are needed, as well as a significant OpEx savings because fewer servers consume less power, floor space, and other data center resources (cooling, power distribution, management).

TCO Savings from DPU Offloads

The total cost of ownership (TCO) savings from DPU offloads will equal...
 CapEx savings from deploying fewer servers to perform the desired work
minus... CapEx additional spending to deploy DPUs in servers
plus... OpEx savings from operating fewer servers
plus... OpEx savings from reduced per-server power consumption due to offloads
minus... additional OpEx (if any) from DPU power consumption.

The following example is based on real-world efficiencies from the DPU. Using our previous example of encrypting all traffic in and out of an IPsec client, we calculate the TCO with 1000 or 10,000 servers over a three-year period. The use of the DPU reduced power consumption on the server and freed up CPU cores, allowing for fewer servers to be deployed to handle the same workload. We calculate both the CapEx savings from needing fewer servers and the power savings on each server from using the DPU crypto offloads. Power costs assume \$0.15/kWh and a PUE of 1.5.

Table 7. TCO calculation from offloading IPsec encryption/decryption to a BlueField DPU, for a large data center with 10,000 servers.

Large Data Center TCO	Servers without DPU	Servers with DPU Offload
Servers needed	10,000	8,200 (18% reduction)
Cost per server	\$10,500 (no DPU)	\$12,000 (with DPU) ¹⁰
Total Server CapEx	\$105,000,000	\$98,400,000 (\$6.6M / 6.3% savings)
Power use per server	728W (0.728 kW)	481W (247W/34% reduction)
Total power use, 3 years	191,318,400 kWh	103,653,576 kWh (45.8% reduction)
Server power cost (\$0.15/kWh)	\$28,697,760	\$15,548,036 (\$13.1M savings)
Total power cost (PUE=1.5)	\$43,046,640	\$23,322,054 (\$19.7M OpEx savings)
3-year TCO (CapEx + OpEx)	\$148,046,640	\$121,722,054 (\$26.3M / 17.8% savings)

We see significant two-way savings from the offload and acceleration capabilities of the BlueField DPU. The offload frees up CPU cores allowing fewer servers to be deployed, reducing CapEx. The lower number of servers and lower per-server power consumption combine to reduce OpEx substantially. The result is a substantial savings of \$26M over three years in a large data center with 10,000 servers. The use of the DPU saves a significant 6.3% on the server CapEx but a huge 46% savings on OpEx due to lower power costs. In this example, this DPU delivers a server CapEx savings of \$6.6M but an OpEx power savings of \$19.7M over three years of operation, assuming a power cost of \$0.15/kWh and PUE of 1.5. (Higher electricity costs and/or a higher PUE ratio will result in even larger OpEx and TCO savings.)

¹⁰ Cost per server with DPU reflects typical street price of a BlueField-2 DPU VPI card with 2x100GbE/EDR network ports, crypto offloads enabled, and 16GB DRAM, minus the typical street price of a SmartNIC with 2x100GbE ports.

DPU Offloads Are the Next Wave for Improving Data Center Efficiency

Using hardware-accelerated networking offloads with an appropriate DPU can greatly reduce power use per server, resulting in more efficient servers, a more efficient data center, and significant cost savings from reduced electricity consumption and reduced cooling loads. In a world facing rising energy costs and increasing demand for green IT infrastructure, the use of DPUs will become increasingly popular to reduce TCO by decreasing both CapEx and OpEx in the data center.

To learn more about the NVIDIA BlueField DPU, visit <https://www.nvidia.com/en-us/networking/products/data-processing-unit/>.

To test the NVIDIA BlueField DPU with VMware vSphere online, apply here: <https://www.nvidia.com/en-gb/launchpad/infra-optimization/experience-vmware-project-monterey-early-access-on-bluefield-2-dpu/>.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation and/or its affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2022 NVIDIA Corporation & Affiliates. All rights reserved.