# How Dell and Broadcom can help you make the transition to IPv6

### Provide a better user experience

Up to 13% better IPv6 performance and lower CPU utilization on read workloads vs. IPv4*

### Perform routine analysis more quickly

Up to 58% better IPv6 performance and lower CPU utilization on read workloads by enabling the Broadcom Offload feature**

### Reduce the time to search databases

Up to 77% better IPv4 performance and lower CPU utilization on read workloads by enabling the Broadcom Offload feature***

## IPv4 vs. IPv6: How we got here

As the internet grew and commercialized late last century, it became increasingly clear that Internet Protocol version 4 (IPv4) limitations would eventually present issues. Enter Internet Protocol version 6 (IPv6) in the 1990s. Despite the technology's age, its adoption has been slow in the US. Until very recently, many companies and other entities still primarily used IPv4, as shown by IPv6 adoption trackers such as one from Google.[1] Recently, however, the transition to IPv6 has been ramping, as some of those IPv4 limitations—such as the dwindling pool of available IPv4 addresses—are quickly becoming reality. One effort to encourage this transition includes a 2020 mandate from the U.S. Office of Management and Budget (OMB) requiring federal government agency devices be at least 80 percent IPv6-only by 2025.[2]

Regardless of mandates or address pools, there are a host of other reasons to choose IPv6 over IPv4. IPv6 includes features such as support for larger packets and multicasting, simpler header formats, smaller routing tables, and the elimination of the network address translation (NAT) process—all of which can increase performance over IPv4 in certain use cases. IPv6 also has built-in end-to-end encryption and name resolution protocol enhancements that contribute to better base security than IPv4.

Despite the advantages of using IPv6, some companies have resisted transitioning because it's not a small undertaking. In this paper, we explain why making the transition can be worth the investment and introduce a solution to help make the transition easier: Dell and Broadcom™ combine to have one of the first IPv6-only compliant end-to-end solutions. We also present the results of our testing, including performance advantages for IPv6 over IPv4 on read workloads and larger performance increases for IPv4 and IPv6 available by enabling the Offload feature in the Broadcom network interface cards (NIC).

* IPv6 without Offload vs. IPv4 without Offload
** IPv6 without Offload vs. IPv6 with Offload
***IPv4 without Offload vs. IPv4 with Offload

## Why organizations are shifting to IPv6

The transition to IPv6 has been a long, slow process that is complicated by the fact that IPv4 and IPv6 are not compatible, requiring companies to either choose just one or manage two networks via dual stack. For many companies, however, fully abandoning IPv4 is not an option. Doubling the number of networks you deploy means doubling the security concerns and hardware expenses. Additionally, applications built on IPv4 may need rebuilding or updating to work with IPv6. Despite these factors, we believe these complications to be worth the benefits you gain from taking advantage of the IPv6 landscape. Not only could organizations use the features we mention above, but as more companies and users move to IPv6, it will also be easier for others to follow. Thus, IPv6 will grow more valuable over time. In time, we hope this can lead to companies being able to shed their IPv4 network, leaving the single, more efficient IPv6 network in place.

Below, we detail some of the reasons to transition to IPv6 including current issues with IPv4, specific industries or government agencies with particular IPv6 requirements, and the benefits of IPv6.

### The shortage of IPv4 addresses

One of the most important reasons for the push to transition to IPv6 is the limited number of possible IPv4 addresses. The IPv4 address space is a 32-bit field, meaning there are a total of $2^{32}$, or roughly 4.3 billion, possible IPv4 addresses. As of November 2019, this address space was officially depleted, meaning there are no new IPv4 addresses to obtain.[3] This has created a premium on IPv4 addresses, allowing companies to sell or lease their addresses, increasing the costs of buying an IPv4 address. Amazon Web Services (AWS), for example, is adding a charge to every IPv4 address on its platform, citing a 300 percent cost increase of IPv4 addresses over the past 5 years.[4] Several workarounds for this problem exist, such as NAT, which allows organizations to map several private addresses within a local network to a single public address before transferring information to the internet. However, NAT comes with its own share of problems that can affect the performance and reliability of network applications. By adding an extra layer of translation and processing, NAT can introduce latency, errors, or packet loss.[5]

IPv6, by utilizing a 128-bit address field, increases available IP addresses to roughly $2^{128}$, or ~$3.4*10^{28}$, essentially solving the address limitation for the foreseeable future. Companies that work with the Internet of Things (IoT), virtual reality (VR), self-driving vehicles, telecom, and other technologies requiring many IP addresses could avoid the IPv4 address market and limited address availability by moving to IPv6.

## The federal government gets involved

US federal government agencies also find themselves impacted by the Office of Management and Budget (OMB) mandate, which claims that "full transition to IPv6 is the only viable option to ensure future growth and innovation in Internet technology and services."[6] The latest version of the mandate states that running dual stack IPv4 and IPv6 networks, as previous versions of the mandate dictated, is too complex and no longer necessary. Instead, this new mandate requires IPv6-only networking, outlining four actions agencies must take:

- Create an IPv6 project team.
- Create and publish an agency-wide policy that states their intentions to phase out all IPv4 use and make all federal IT systems IPv6 enabled by the end of 2023.
- Identify and test at least one IPv6 pilot by the end of 2021.
- Develop a plan by 2021 for implementing IPv6-only networking, with milestones including at least 50 percent of IP-enabled assets, transitioned to IPv6-only by the end of 2023, and 80 percent on IPv6-only networks by the end of 2025.[7]

This means that by the end of the first quarter of 2024, federal agencies should already have half of their systems converted to IPv6-only, and the rest fully transitioned in just two more years.

## Telecom and ISP industries are leading the way

With the development of the 5G cellular network and its need for high speeds and low latencies, much of the telecom industry has already converted to IPv6. As more and more devices connect over cellular networks with 5G, the increased address pool of IPv6 provides another benefit for internet service providers (ISPs) and cellular network companies. Additionally, the built-in quality of service (QoS) field in the IPv6 header allows ISPs to prioritize voice traffic over other traffic less vulnerable to latency such as http, SSH, and more. According to Akamai, a content delivery network and cloud computing company, ISPs and telecoms such as Comcast Cable, Verizon Business, AT&T, and T-Mobile have all reached IPv6 adoption above 70 percent—T-Mobile has as much as 92.7 percent IPv6 adoption.[8]

As these network providers continue to expand the IPv6 backbone, and mobile app developers continue to embrace IPv6 advantages, the rest of the world's industries will lag behind if they continue to rely on IPv4.

## How IPv6 helps your business

Even if you don't need a wide range of IP addresses or aren't part of the government mandate or the telecom industry, there are still benefits you can see from transitioning from IPv4 to IPv6. First, IPv6 could increase performance in several ways, mostly by increasing network efficiency. Second, IPv6 can offer some additional security benefits over IPv4. While IPv4 has had more security upgrades and patches to existing networks—simply by virtue of existing longer—the features that come with IPv6 offer stronger baseline security, which we examine in the following pages. Increased use and investment in IPv6 security enhancements should quickly close any existing gap between it and current IPv4 network security.

IPv6 can also improve network performance over IPv4 by using a more simplified header that takes less time and fewer resources to process. All IP packets include headers that contain the necessary information for proper route allocation and delivery. Much like the parts of a physical address tell postal workers the house, street, city, and country the letter originated from and is destined for, the IP header includes such information as the IP addresses of the source and destination devices, the version indicator, the total length of the packet, and other important information. Instead of requiring routers to perform a header checksum to ensure data integrity of every packet, IPv6 relies on Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and other existing protocol checksums. Additionally, this header allocates optional and non-essential header fields to header extensions, leaving only the most vital information in the header for processing. This practice increases the efficiency of processing data packets across the network by getting the metaphorical "letter" to the right "building" quickly and letting the extension headers then direct it more specifically, to the equivalent of the correct suite, floor, or person.

IPv6 also eliminates the need for NATs as each device on an IPv6 network can have its own unique IP address. If devices aren't having to share IP addresses, then the routers do not need to translate the network addresses to send the packet to the correct device, thus eliminating a step in the data transmission process. Finally, while the Maximum Transmission Unit (MTU) technology limits packet sizes in many networks to 1,500 bytes, IPv6 networks are capable of much larger packet sizes, up to 4 GB. IPv4 packets, on the other hand, are limited to much smaller theoretical maximums of 64K bytes.[9]

Another way that IPv6 increases network efficiency is with its ability to multicast instead of broadcast. On IPv4, data transmission is broadcast: When a packet leaves a source device, the information is sent to every host connected to the network. Every single host, in turn, checks the packet to see if the data is meant for it. IPv4 is like a kid yelling out, "Mom!" at a crowded playground, causing every mother in the vicinity to stop and check if their child is the one in need. IPv6, on the other hand, uses multicast, which is the ability to transmit a packet only to the device or devices for which it is intended. Now, instead of a child yelling, "Mom!" into a crowd, IPv6 taps their parent on the shoulder and talks directly to them, allowing the other moms to focus on their own children.

IPv4 routing tables, the list of networks and other links a router consults to determine where a packet should go, are quite large—and they continue to grow. The larger the table, the longer it takes to search and find the relevant data. While IPv4 addresses and networks can be difficult to aggregate or simplify, the structure of the IPv6 address allows for just that. The IPv6 address contains three parts: the network or site prefix, the subnet ID, and the host or interface ID.[10] Routers use a site prefix to route the packet through the internet. Network creators and ISPs can also use site prefixes to create packet groups that aggregate packets going in the same general direction. Aggregating the IPs like this allows the internet and routers to act similarly to public transportation. A person boards the train at a specific stop, rides it to another stop, then leaves the train and follows the maps to the exit nearest the actual destination they have in mind. With IPv6, packets behave similarly, where several packets with completely different final destinations can all "exit" the internet at the same stop and travel more granularly from there. This behavior allows for smaller, more efficient routing tables, speeding up the routing process and lowering the overhead on router hardware.

As we mentioned above, comparing IPv4 and IPv6 security isn't completely straightforward, as professionals have invested more time and effort into the older IPv4 network. However, IPv6 has the potential to be more secure than IPv4 due to at least two built-in advantages. First, IPv6 has end-to-end encryption and authentication built in via default Internet Protocol Security (IPsec) inclusion. While enabling IPsec on IPv6 networks may not be mandatory in some places, implementing it provides better security.[11] The second way IPv6 enhances security is simply by being much larger. With its more numerous IP addresses, IPv6 networks are nearly impossible to brute-force scan.[12]

One more benefit of IPv6 is the ability for users to implement stateful or stateless configurations. While IPv4 network devices rely heavily on dynamic host configuration protocol (DHCP) devices to assign IP addresses, IPv6 networks can use the stateless address auto configuration (SLAAC) technology to let devices generate their own IPs without manual application or the use of a third-party device such as DHCP. If users prefer to use a stateful network, IPv6's version of DHCP is available.

### The NIST and transitioning to IPv6

According to the National Institute of Standards and Technology (NIST), "the IPv6 protocol suite offers a vastly greater address space than IPv4 and supports significant new capabilities necessary to enable modern network environments."[13] To facilitate this transition, the NIST National Cyber-security Center of Excellence (NCCoE) "is planning a project to provide guidance and reference architecture that address operational, security, and privacy issues associated with the evolution to IPv6-only network infrastructures."[14]

The project aims to provide enterprise organizations attempting to transition to IPv6 with guidance and tools that will "ensure that evolving enterprise IT environments to be IPv6-only can be accomplished in a secure and robust manner."[15]

Regardless of the size of your network, the NIST Cybersecurity Practice Guide will provide best-practices and documentation to ensure that your transition to IPv6 is secure. For more information about the transition resources available from the NIST, visit https://www.nccoe.nist.gov/projects/IPv6-transition.

## How Dell servers and Broadcom NICs can help on your journey to IPv6

The National Institute of Standards and Technology (NIST) is part of the U.S. Department of Commerce and serves to support and promote technology and science innovation and investment.[16] In addition to mandating that government agencies switch to IPv6-only networks, the OMB has mandated the NIST to create a set of standards and tools to support the transition. The resulting U.S. Government IPv6 (USGv6) Program develops, tests, and maintains IPv6 standards to help companies and government agencies ensure successful IPv6 transitions and deployments. Vendors can use the IPv6 tests and certifications that NIST developed to certify that their products meet the requirements and standards defined by NIST's program.[17] In response to the 2020 OMB mandate for all government agencies, NIST revised their USGv6 program to include several objectives including updating specifications to add new and remove old technologies and streamlining their testing program based on previous experience.[18] With a set of agreed upon standards, definitions, and requirements, NIST and the USGv6-r1 provide OEMs a way to ensure their customers that their products are ready for IPv6 implementation.

Dell Technologies™ is the first company to offer a full USGv6-r1 certified server and storage stack for IPv6-only networking.[19] Dell's certifications include:

- Dell™ PowerEdge™ servers – first in the industry to be fully IPv6 Ready Logo 5.1.2 compliant with RedHat 8.4 and Windows 2019 and 2022.
- Dell PowerEdge iDRAC9 with FW version 5.10.0.00 – first baseboard management controller (BMC) validated by USGv6-r1 as IPv6-only compliant.
- Unity-XT storage array – first storage product validated by USGv6-r1 as IPV6-only compliant meeting the requirements in the IPv6-Only Functional v1.1 (36277) profile.[20]
- Additionally, PowerStore, PowerEdge with VMware 8.0.1, and PowerEdge with SUSE SLES15 SP4 are on the USGv6-r1 registry.

Additionally, Dell servers and storage leverage Broadcom NICs and Adapters to ensure network performance and security for IPv6 customers. Broadcom NICs such as the Broadcom BCM957508-P2200G dual-port 100GbE NIC include several IPv6 offloads that can boost network performance.[21] These offloads allow the NIC to directly handle some of the computational needs of the network rather than use the OS stack, which can provide lower latencies and lower CPU utilization dedicated to network traffic.[22] Broadcom BCM957508-P2200G dual-port 100GbE NICs also offer features such as NVME over Fabrics (NVMe-oF) capabilities that allow NVMe storage traffic to travel through network instead of directly through PCIe channels. NVMe-oF allows users to connect storage via Ethernet (TCP), Fibre, and RDMA.[23] This rerouting of storage network allows for extremely low latency to get the most out of NVMe-based storage.

To show how Dell and Broadcom can provide one great hardware stack option for your IPv6 needs, we conducted some testing to highlight the performance you can expect with IPv6-connected Dell PowerEdge R660 servers to a PowerStore 1200T storage array using 100GbE Broadcom 57508 NICs.

## Broadcom NICs

Broadcom NICs can serve most networking needs because they offer speeds ranging from 1G to 200G. According to Broadcom, their network cards feature:

- "Low power adapters and controllers with outstanding thermal performance
- Low latency and high throughput RoCEv2 [for] ground-breaking performance for machine learning, HPC and storage applications
- Broadsafe™ embedded security [for] Silicon Root of Trust and attestation delivering industry's most secure Ethernet controller
- Modern architecture [that] delivers industry's lowest latency and lowest CPU utilization for real-world network conditions
- TruFlow™ engine [to accelerate] virtual switch processing, reduces server CPU usage
- TruManage™ [for] end-user manageability needs to allow fine-tuning of networks for maximum performance
- On-chip tunneling protocol processing for Geneve, VXLAN, and NVGRE [that] provides up to a 5x throughput increase
- Acceleration engines for SDN and NFV [to] enable leading-edge service provider solutions"[24]

## Measuring performance

The goal of our performance testing was to show the benefits of the USGv6-r1 IPv6-only certified Dell PowerEdge server and Broadcom NIC solution. This included investigating the performance differences between IPv4 and IPv6 in a real-world environment. Most North American users continue to rely on IPv4, which typically requires NAT or packet fragmentation support from a network router.[25] In a typical scenario, a routing device need only read an IPv4 packet to determine its destination and send it on its way. In the case of IPv4 using NAT or requiring packet fragmentation, the routing device must modify the packet before it can send it along, which requires overhead. We wanted to quantify the impact of this overhead on network performance.

Other than the Layer 3 protocol, every aspect of the test scenario was the same. We used Linux standard tools and NVMe/TCP and/or NFS transport protocols for this test. We did not attempt to enable the best speed of each of these protocols; rather, we used those protocols to drive the tests to compare any differences in speed based on the journey of that data provided by Layer 3 (IP).

We configured two Dell PowerEdge servers as SUSE Linux Enterprise Server 15 SP4 hosts, sending data of diverse sizes to a Dell PowerStore storage array using the transport protocols we identified earlier. The data traversed multiple switches we configured to provide Border Gateway Protocol (BGP) routing and packet fragmentation within a heterogeneous multi-hop network.

We configured the host networks using an MTU of 9,000, with a 1,500 MTU on the switches emulating the core network (which forced packet fragmentation). The edge switches used BGP routing to communicate with the core network.

### Comparing IPv6 and IPv4 performance without the Broadcom Offload feature

First, we tested the relative performance of IPv6 and IPv4 on a write workload with the Broadcom Offload feature off. Table 1 presents the results. In terms of both performance (IOPS and throughput in MB per second) and CPU utilization, we observed approximate parity between the two IP versions at both block sizes we tested.

Table 1: IPv6 vs. IPv4 performance on a write workload with Offload off. Higher IOPS and MB/sec and lower CPU utilization are better. Source: Principled Technologies.

| Write workload, Offload off | | | | |
|---|---|---|---|---|
| **IP version** | **Block size** | **IOPS** | **MB/sec** | **Percentage CPU utilization** |
| IPv4 | 256K | 8,696.1 | 2,174.01 | 4.9 |
| IPv6 | 256K | 8,752.1 | 2,188.02 | 4.9 |
| **IPv6 % improvement** | | **0.64%** | **0.64%** | **0.00%** |
| IPv4 | 64K | 34,862.7 | 2,178.92 | 6.6 |
| IPv6 | 64K | 34,972.1 | 2,185.76 | 6.5 |
| **IPv6 % improvement** | | **0.31%** | **0.31%** | **1.51%** |

Next, we ran the same test using a read workload. As Table 2 shows, in contrast to the comparable performance we observed on the write workload, IPv6 had a performance advantage over IPv4 on the read workload. At the larger block size of 256K, IPv6 delivered 13.83 percent greater performance. At the smaller block size of 64 K, IPv6 delivered 9.83 percent greater performance. These results indicate that users in a real-world setting would enjoy better performance by using IPv6. We also observed a CPU utilization improvement for IPv6.

Table 2: IPv6 vs. IPv4 performance on a read workload with Offload off. Higher IOPS and MB/sec and lower CPU utilization are better. Source: Principled Technologies.

| Read workload, Offload off | | | | |
|---|---|---|---|---|
| **IP version** | **Block size** | **IOPS** | **MB/sec** | **Percentage CPU utilization** |
| IPv4 | 256K | 19,987.8 | 4,996.95 | 14.1 |
| IPv6 | 256K | 22,752.4 | 5,688.09 | 13.6 |
| **IPv6 % improvement** | | **13.83%** | **13.83%** | **3.54%** |
| IPv4 | 64K | 73,194.1 | 4,574.63 | 13.9 |
| IPv6 | 64K | 80,392.4 | 5,024.53 | 12.5 |
| **IPv6 % improvement** | | **9.83%** | **9.83%** | **10.07%** |

## Measuring the impact of the Broadcom Offload feature on IPv6 performance

A secondary component of our testing was investigating the capabilities of the Broadcom IP Offload feature. IP Offloading is a feature Broadcom has implemented in its NIC (Layer 2) to process IP (Layer 3) data to offload the processing of this data from the OS/CPU, leaving those clock cycles to process user data rather than managing the flow control of the protocol. We refer to this feature as Offload.

Table 3 presents IPv6 performance on a write workload with Offload off and with Offload on. While performance was comparable under both conditions, CPU utilization was lower with Offload on.

Table 3: IPv6 performance on a write workload with Offload off and Offload on. Higher IOPS and MB/sec and lower CPU utilization are better. Source: Principled Technologies.

| Write workload, IPv6 | | | | |
|---|---|---|---|---|
| | **Block size** | **IOPS** | **MB/sec** | **Percentage CPU utilization** |
| Offload off | 256K | 8,752.1 | 2,188.02 | 4.9 |
| Offload on | 256K | 8,615.5 | 2,153.88 | 2 |
| **Offload on % improvement** | | **-1.56%** | **-1.56%** | **59.18%** |
| Offload off | 64K | 34,972.1 | 2,185.76 | 6.5 |
| Offload on | 64K | 34,895.6 | 2,180.97 | 3.4 |
| **Offload on % improvement** | | **-0.21%** | **-0.21%** | **47.69%** |

Table 4 presents IPv6 performance on a read workload with Offload off and with Offload on. In contrast to the approximate parity we saw on the write workload, performance improved greatly with the use of Offload, particularly at the larger block size, where IPv6 delivered 58.15 percent greater performance than with Offload off. At the 64K block size, enabling Offload improved performance by 25.43 percent. Figure 1 illustrates these advantages. As we saw with the write workload, CPU utilization was lower with Offload on.

Table 4: IPv6 performance on a read workload with Offload off and Offload on. Higher IOPS and MB/sec and lower CPU utilization are better. Source: Principled Technologies.

| Read workload IPv6 | | | | |
|---|---|---|---|---|
| | **Block size** | **IOPS** | **MB/sec** | **Percentage CPU utilization** |
| Offload off | 256K | 22,752.4 | 5,688.09 | 13.6% |
| Offload on | 256K | 35,983.4 | 8,995.86 | 8.3% |
| **Offload on % improvement** | | **58.15%** | **58.15%** | **38.97%** |
| Offload off | 64K | 80,392.4 | 5,024.53 | 12.5% |
| Offload on | 64K | 100,840.9 | 6,302.55 | 7.3% |
| **Offload on % improvement** | | **25.43%** | **25.43%** | **41.60%** |

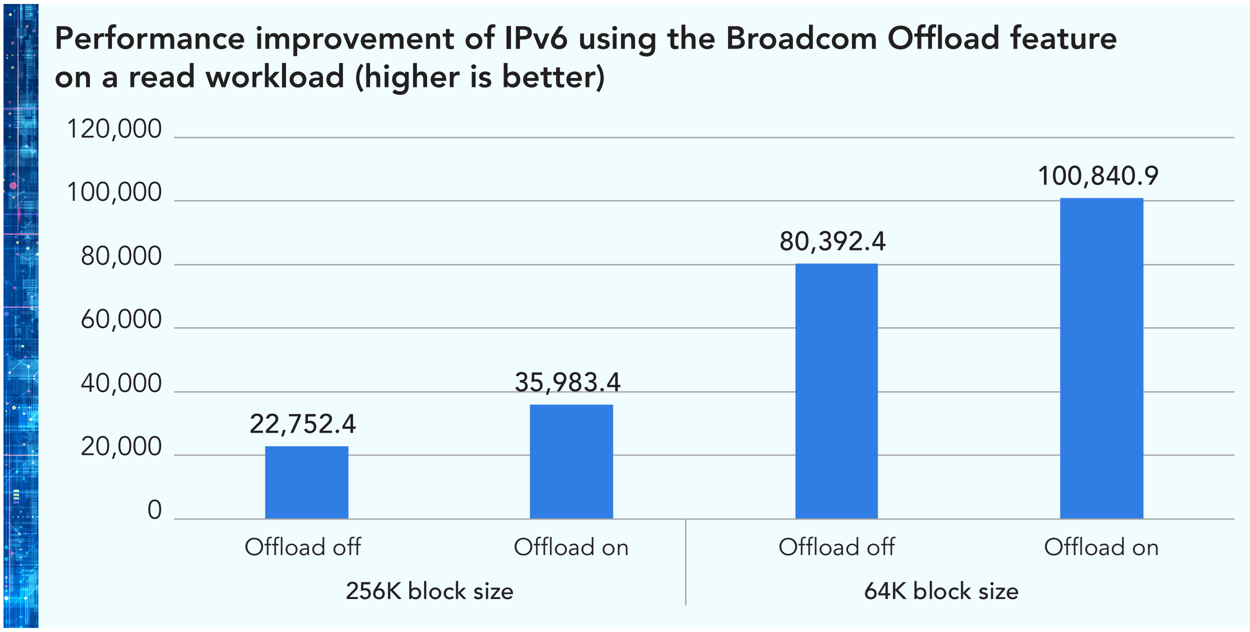


Figure 1: Performance improvement of IPv6 using Offload feature on a read workload. Higher is better.
Source: Principled Technologies.

## Measuring the impact of the Broadcom Offload feature on IPv4 performance

Table 5 presents IPv4 performance on a write workload with Offload off and with Offload on. As we saw with IPv6, performance was comparable under both conditions and CPU utilization improved with Offload on.

Table 5: IPv4 performance on a write workload with Offload off and Offload on. Higher IOPS and MB/sec and lower CPU utilization are better. Source: Principled Technologies.

| Write workload IPv4 | | | | |
|---|---|---|---|---|
| | **Block size** | **IOPS** | **MB/sec** | **Percentage CPU utilization** |
| Offload off | 256K | 8,696.1 | 2,174.01 | 4.9 |
| Offload on | 256K | 8,596.0 | 2,148.99 | 2.0 |
| **Offload on % improvement** | | **-1.15%** | **-1.15%** | **59.18%** |
| Offload off | 64K | 34,862.7 | 2,178.92 | 6.6 |
| Offload on | 64K | 34,727.8 | 2,170.49 | 3.6 |
| **Offload on % improvement** | | **-0.38%** | **-0.38%** | **45.45%** |

As Table 6 shows, the impact of enabling Offload on IPv4 read performance followed the same pattern we saw with IPv6. Using Offload dramatically improved read performance, by 77.90 percent at the larger block size and by 38.73 percent at the 64K block size. Figure 2 highlights these performance improvements. Once again, using Offload improved CPU utilization.

Table 6: IPv4 performance on a read workload with Offload off and Offload on. Higher IOPS and MB/sec and lower CPU utilization are better. Source: Principled Technologies.

| Read workload IPv4 | | | | |
|---|---|---|---|---|
| | **Block size** | **IOPS** | **MB/sec** | **Percentage CPU utilization** |
| Offload off | 256K | 19,987.8 | 4,996.95 | 14.1 |
| Offload on | 256K | 35,559.3 | 8,889.81 | 8.1 |
| **Offload on % improvement** | | **77.90%** | **77.90%** | **42.55%** |
| Offload off | 64K | 73,194.1 | 4,574.63 | 13.9 |
| Offload on | 64K | 101,545.6 | 6,346.60 | 7.4 |
| **Offload on % improvement** | | **38.73%** | **38.73%** | **46.76%** |

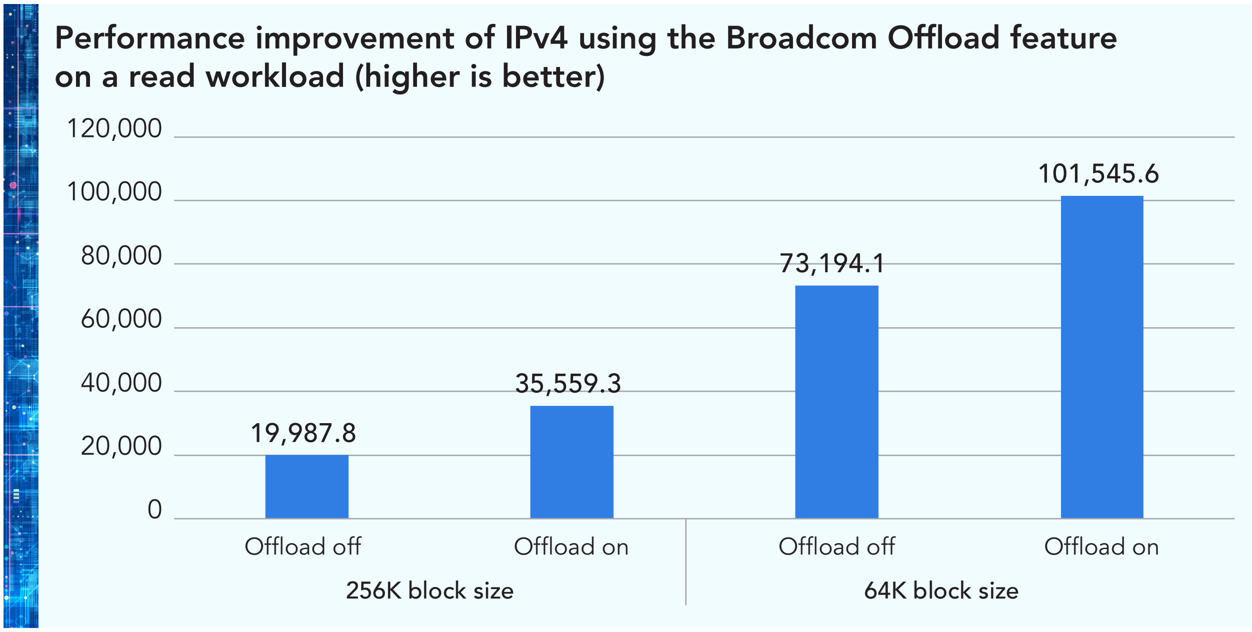


Figure 2: Performance improvement of IPv4 using the Broadcom Offload feature on a read workload. Higher is better. Source: Principled Technologies.

While we have discussed many advantages to making the shift to IPv6, our test results demonstrate that companies who opt not to do so immediately could reap performance benefits on read workloads—and CPU utilization benefits on both read and write workloads—by using the Dell-Broadcom solution we tested and enabling the Broadcom NIC Offload feature.

## Conclusion

With IPv4 address pools rapidly disappearing and a federal mandate for government agency devices to begin shifting to IPv6-only and telecom 5G with IoT and edge devices, it's clear that IPv6 is the future. Transitioning from IPv4 to IPv6 can be a challenge, so organizations may be interested to learn that switching to IPv6 has the potential to improve performance. In our testing without the Broadcom Offload feature, IPv6 delivered comparable performance to IPv4 on write workloads and better performance on read workloads while also reducing CPU utilization. When we enabled the Broadcom Offload feature on both IPv6 and IPv4, read workload performance increased dramatically and CPU utilization on both read and write workloads improved. Whether your organization is transitioning to IPv6 right away or choosing to delay the shift, this feature can boost performance on read workloads, which can improve the experience for users, reduce backup windows, and allow databases to load more quickly.

---

1. Google, "IPv6," accessed December 21, 2023, https://www.google.com/intl/en/ipv6/statistics.html.
2. Russell T. Vought, "MEMORANDUM FOR HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES," accessed December 21, 2023, https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf.
3. RipeNCC, "What is IPv4 Run Out?" accessed December 21, 2023, https://www.ripe.net/manage-ips-and-asns/ipv4/ipv4-run-out.
4. Jeff Barr, "New – AWS Public IPv4 Address Charge + Public IP Insights," accessed December 21, 2023, https://aws.amazon.com/blogs/aws/new-aws-public-ipv4-address-charge-public-ip-insights/.

5. Linkedin, "How can NAT affect the performance and reliability of network applications?" accessed December 21, 2023, https://www.linkedin.com/advice/0/how-can-nat-affect-performance-reliability-network.
6. Russell T. Vought, "MEMORANDUM FOR HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES," accessed December 21, 2023, https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf.
7. Russell T. Vought, "MEMORANDUM FOR HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES."
8. Akamai, "IPv6 Adoption Visualization," accessed December 21, 2023, https://www.akamai.com/internet-station/cyber-attacks/state-of-the-internet-report/ipv6-adoption-visualization.
9. Network Academy, "IPv4 vs IPv6 - Understanding the differences," accessed December 21, 2023, https://www.networkacademy.io/ccna/IPv6/IPv4-vs-IPv6#:~:text=Both%20the%20IPv4%20Total-Length,to%2065%2C355%20byte-long%20packets.
10. Diego Asturias, "Migration to IPv6: Benefits and Techniques," accessed December 21, 2023, https://www.rapidseedbox.com/blog/IPv6-migration.
11. Diego Asturias, "Migration to IPv6: Benefits and Techniques."
12. Diego Asturias, "Migration to IPv6: Benefits and Techniques."
13. "IPv6 Transition," accessed December 21, 2023, https://www.nccoe.nist.gov/projects/IPv6-transition.
14. "IPv6 Transition."
15. "IPv6 Transition."
16. NIST, "About NIST," accessed December 21, 2023, https://www.nist.gov/about-nist.
17. NIST, "USGv6," accessed December 21, 2023, https://www.nist.gov/programs-projects/usgv6-program/usgv6.
18. NIST, "USGv6 Revision 1," accessed December 21, 2023, https://www.nist.gov/programs-projects/usgv6-program/usgv6-revision-1.
19. George Dilger, "Industry First IPv6-only Support on Servers and Storage," accessed December 21, 2023, https://www.dell.com/en-us/blog/industry-first-ipv6-only-support-on-servers-and-storage/.
20. George Dilger, "Industry First IPv6-only Support on Servers and Storage."
21. Broadcom, "Broadcom Ethernet NICs PCIe NIC Ethernet Adapters Specification Sheet," accessed December 21, 2023, https://docs.broadcom.com/doc/PCIe-NIC-Ethernet-Adapters-Specification-Sheet.
22. Eden Kim and Fred Zhang, "Optimizing NVMe® over Fabrics (NVMe-oFTM)," accessed December 21, 2023, https://www.snia.org/sites/default/files/education/snia-optimizing-nvme-over-fabrics-nvme-of.pdf.
23. Juan Mulford, "What is NVMe-oF?," accessed December 21, 2023, https://www.storagereview.com/review/nvme-nvme-of-background-overview.
24. "Ethernet Network Adapters," accessed December 21, 2023, https://www.broadcom.com/products/ethernet-connectivity/network-adapters.
25. Google, "IPv6," accessed December 21, 2023, https://www.google.com/intl/en/ipv6/statistics.html.

**Read the science behind this report at https://facts.pt/HrC2sZI ▶**

Principled Technologies®

Facts matter.®



This project was commissioned by Dell Technologies.