



Boost Performance for Virtual Workloads by Using GPUs and DPUs

Testing by Prowess Consulting demonstrated significant performance jumps for workloads running in VMware vSphere® virtual environments when they connected to Dell™ PowerEdge™ servers configured with hardware accelerators.

More and more small and medium-sized businesses (SMBs) are turning to hardware acceleration to provide a performance jump for many critical workloads. For example, applications that rely on AI and machine learning (ML) require higher levels of performance than many CPU-only configurations can provide. The same is often true for 3D modeling, rendering, and advanced simulations used in architecture, financial services, healthcare, and manufacturing deployments. Hardware accelerators not only improve performance for targeted workloads, they also free up valuable CPU cycles that can be allocated to other tasks.

Because so many organizations rely on the convenience, simplified management, and security benefits of VMware vSphere, Prowess Consulting decided to investigate the potential gains from hardware accelerators in a virtualized environment in this study commissioned by Dell Technologies.

We deployed two Dell™ PowerEdge™ R760 servers with VMware vSphere, one with a graphics processing unit (GPU) and a data processing unit (DPU) and one without hardware accelerators. We tested both the direct performance gains and the reduced latency benefits provided by offloading functions from the CPU, and the results were impressive. A video rendering workload was 3,681% faster, and an iPerf® file transfer benchmark test resulted in 36% gains for the accelerated configuration, compared to the CPU-only configuration. In addition, we recorded 30% better CPU utilization when accelerators were used.

Highlights

Compared to a non-accelerated deployment, adding GPUs and DPUs can significantly improve performance for demanding workloads in virtualized environments.

**3,681%
faster**
video
rendering

**36%
higher**
iPerf® file-transfer
performance

**30%
better**
CPU
utilization

Testing Hardware Accelerators for Performance and CPU Utilization

For the first test, we recorded Blender® benchmark scores on a single VM, both with and without GPU acceleration. As Figure 1 shows, the accelerated configuration generated scores 36.81x (or 3,681%) higher than the non-accelerated configuration.

Blender® Benchmark Rendering Scores in Samples Per Minute—Single VM (higher is better)

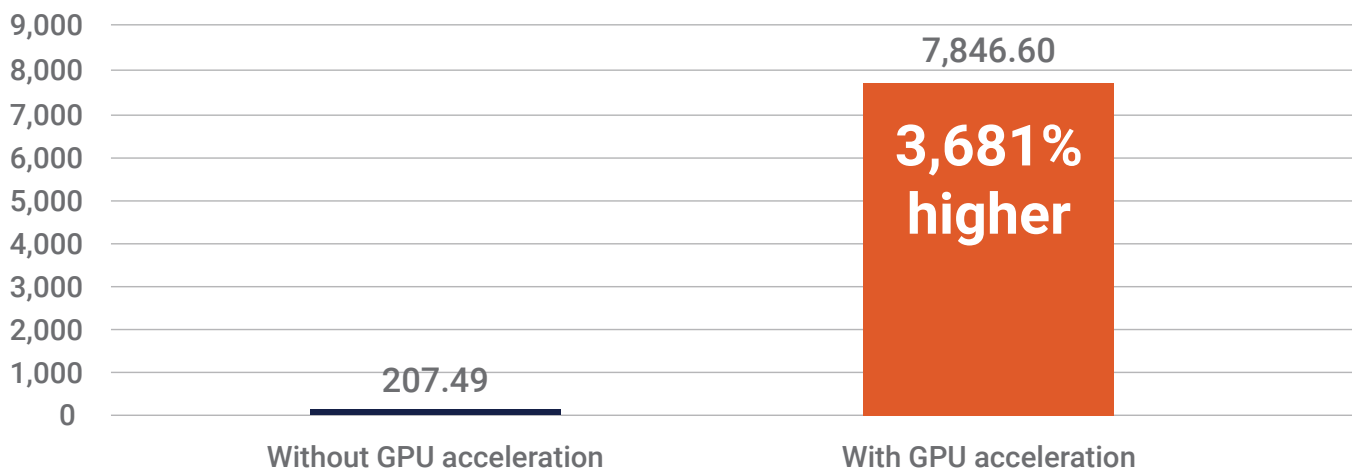


Figure 1. Performance test results for a non-accelerated versus an accelerated Blender® benchmark workload on a single VM

We also ran iPerf in conjunction with the Blender workload to measure file-transfer performance in gigabits per second (Gbps), both with and without DPU acceleration. As Figure 2 demonstrates, the DPU-accelerated configuration resulted in a 36% gain in file transfer speed.

Transfer Rate (Gbps) Between Two VMs with a Blender® Workload Running Concurrently (higher is better)

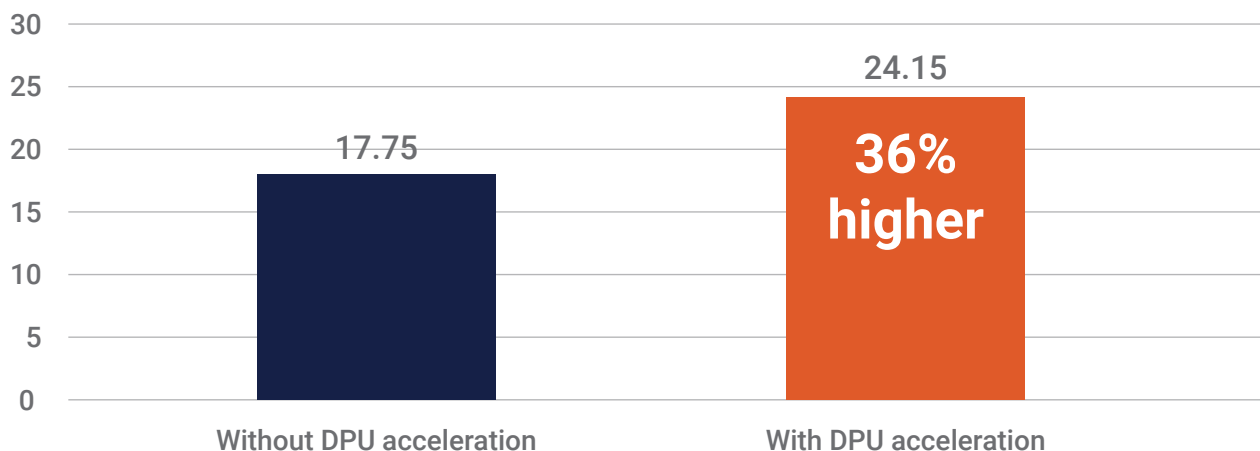


Figure 2. File-transfer rates recorded with iPerf®, showing performance without and with DPU acceleration, measured on a single VM in conjunction with a Blender® workload

In addition to significantly improving performance, the accelerated configuration resulted in much lower CPU utilization, as shown in Figure 3. This drop in CPU utilization demonstrates the ability of the DPU and GPU to free up the CPU to handle additional work for other workloads as needed.

Note that we recorded CPU utilization across all of the threads in all CPUs deployed in each configuration. Because there were 128 threads total, the resulting CPU utilization percentage can be greater than 100.

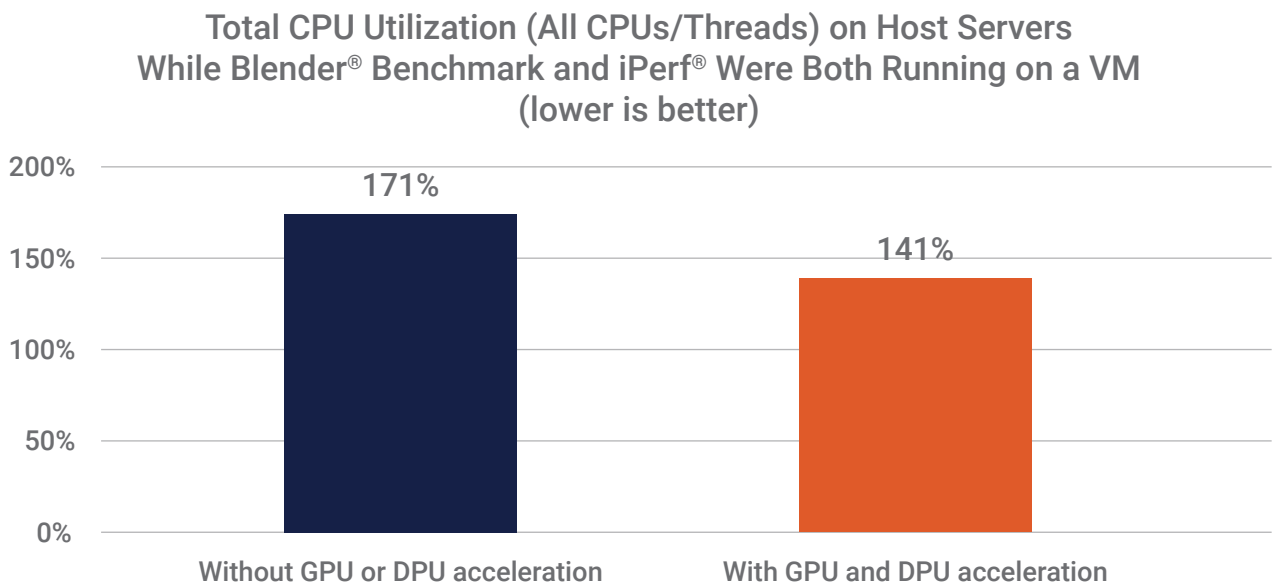


Figure 3. Host CPU utilization (combined CPUs/threads) without and with both GPU and DPU acceleration, in conjunction with Blender® and iPerf® workloads running on a single VM

Get the full story by reading the technical research report

[“Accelerate Demanding Graphics and AI Workloads with GPUs and DPUs, Even in Virtualized Environments”](#)

¹ 2024. Source: Prowess Consulting. “Accelerate Demanding Graphics and AI Workloads with GPUs and DPUs, Even in Virtualized Environments.” Commissioned by Dell Technologies. February 2024.



The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies.

Results have been simulated and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC.

Copyright © 2024 Prowess Consulting, LLC. All rights reserved.

Other trademarks are the property of their respective owners.