TheFuturum Group

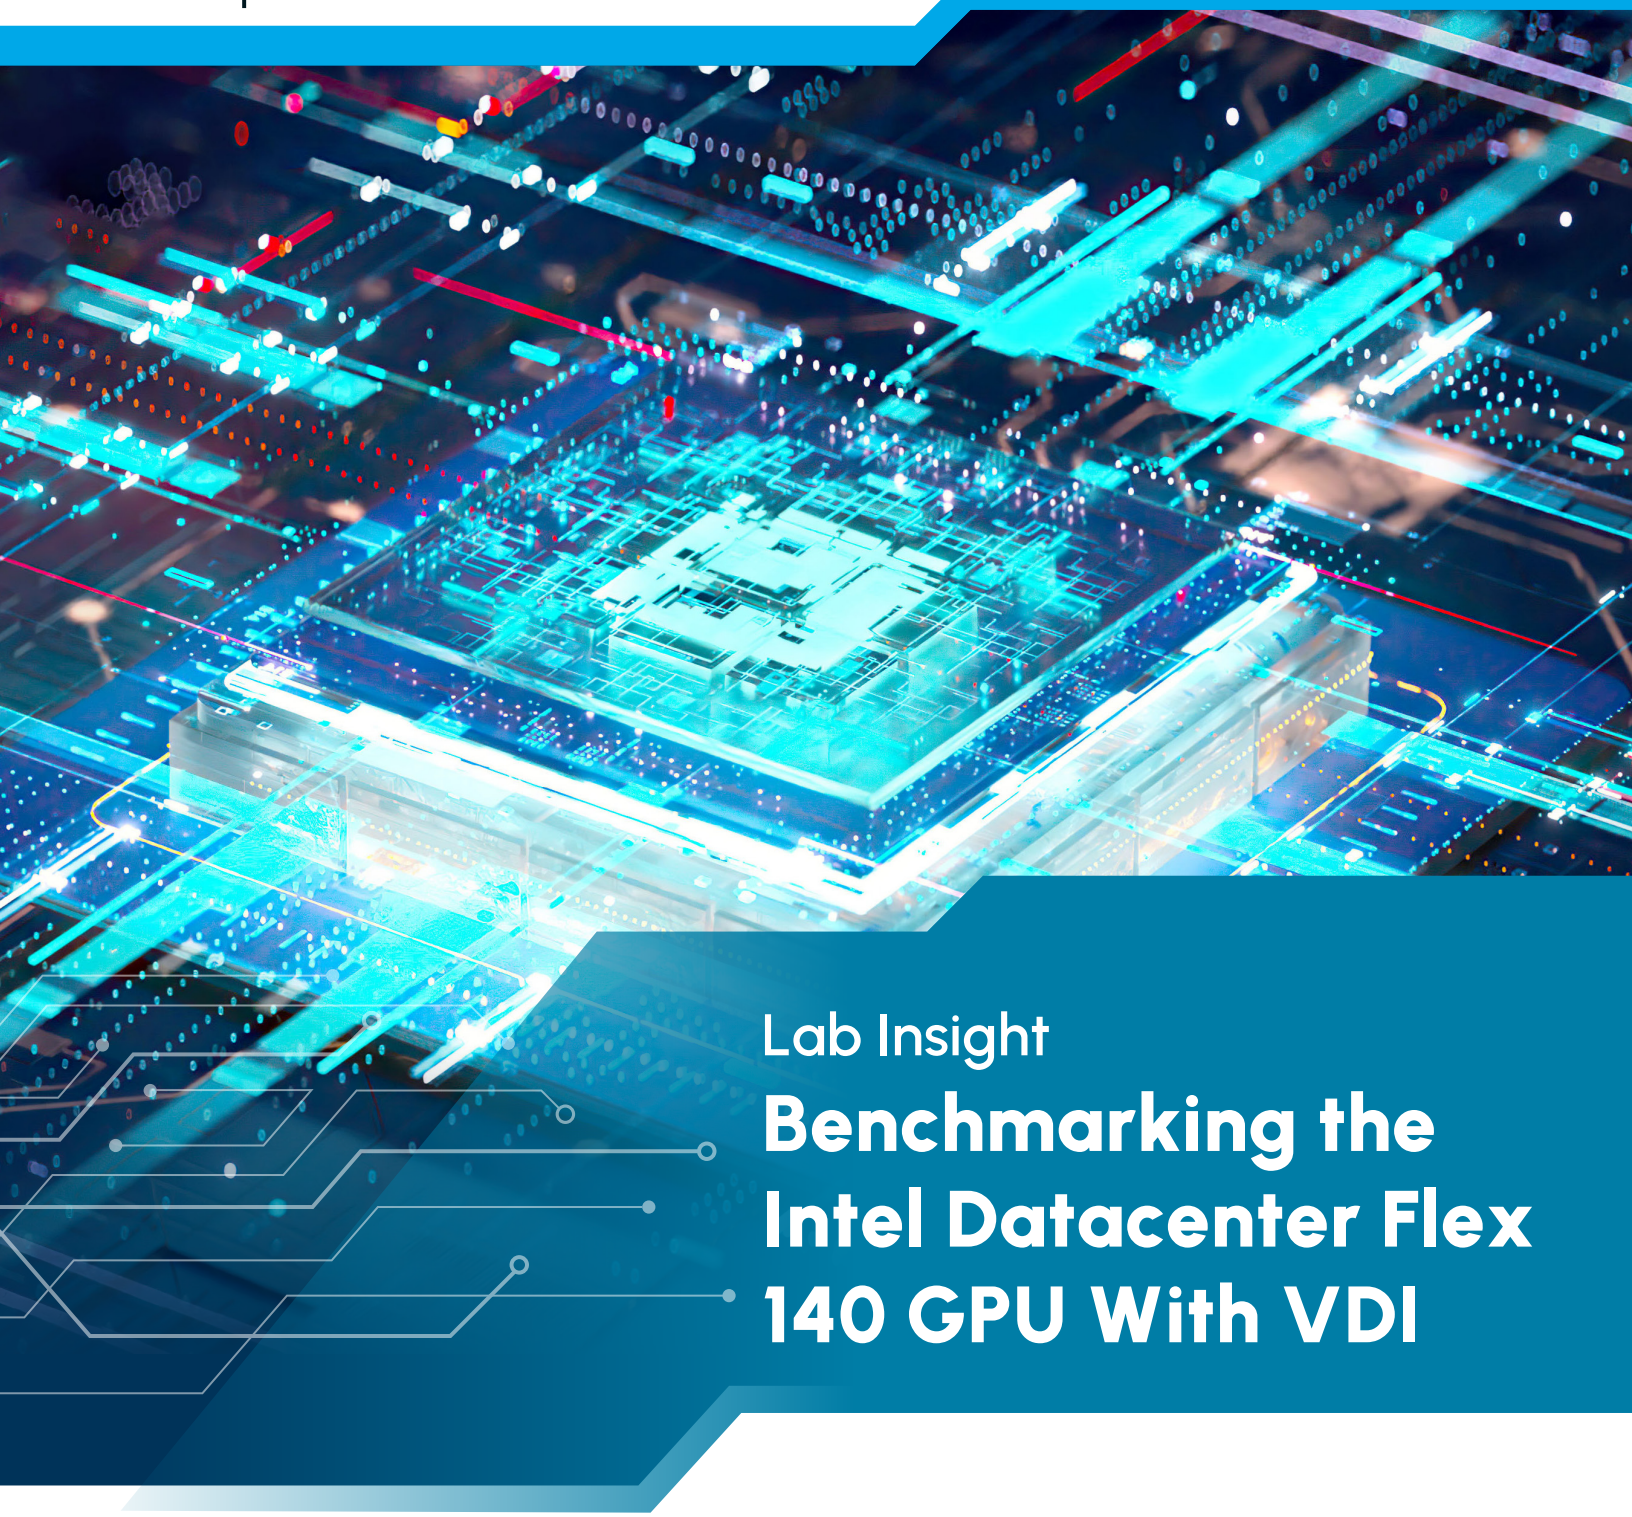

Lab Insight

Benchmarking the Intel Datacenter Flex 140 GPU With VDI

AUTHOR **Jonathan Fellows**
Performance Validation Engineer | The Futurum Group

APRIL 2024

This study was sponsored by Intel

Overview

Hyperconverged Infrastructure (HCI) continues to be a popular way for companies to utilize the compute, memory and storage resources needed for hosting large numbers of virtual machines for their employees. Due to its relative ease of use and management, whether on-premises or in the cloud, HCI is an excellent choice for companies looking to manage virtual infrastructure. Utilizing HCI as a method to deploy virtual desktops enables companies to create differentiated pools of VDI instances that match users' needs, including creating configurations that utilize hardware GPU cards in the infrastructure to enable higher graphics performance when desired.

Intel contracted the Futurum Group Labs to evaluate the performance, and price-performance of multiple HCI configurations supporting VMware Horizon VDI deployments running clusters of servers using Intel CPUs and GPUs, utilizing VMware vSphere together with vSAN storage. Adding GPUs to configurations is an effective way to enable higher graphics performance for VDI users that need it, while offloading CPU resources for other VDI users.

The application workloads were created with Login Enterprise from **Login VSI** using a standard workload profile known as the "Knowledge Worker" (KW) profile, along with a **GPU Reference Workload**. Login Enterprise was used to measure the performance of each configuration, including configurations using Intel Data Center GPU Flex Series 140.

Our testing found that installing Flex 140 GPUs into the servers resulted in a large increase in the performance, measured by FPS (Frames Per Second), for individual users in graphics intensive applications. The increase in FPS was complemented by a large reduction in CPU utilization. Important takeaways from our testing include:

- With the Microsoft 3D Viewer app FPS increased from ~15 to 60, a nearly 4x gain, while using ~35% less CPU resources when using Flex 140 GPUs for individual users
- Similarly, Flex 140 GPUs increased frame rates from ~15 to ~35 and reduced CPU utilization by ~48% for individual users during the Media Player application (CPU dependent)
- For the cost of about $14 more per user (CPU dependent) graphics performance is significantly improved for users with GPU access
- Across all users running the GPU workload, a 35-40% reduction in CPU utilization (CPU dependent) was seen with GPUs present

HCI for Virtual Desktops

Hyperconverged Infrastructure's ability to rapidly scale infrastructure significantly reduces overhead for IT departments compared to traditional infrastructure options. Nodes can be quickly added providing additional compute, memory and storage utilizing common building blocks. This enables rapid scaling of VDI instances to meet requirements for growing organizations. In the cloud, this may be as simple as requesting additional infrastructure through the cloud provider. On-premises, this can be done by physically adding another node to a cluster and provisioning more VDI instances.

With VMware vSAN, storage can be easily scaled along with CPU and memory resources. This allows for the management of all aspects of the HCI deployment through VMware vCenter, further reducing overhead and costs for IT departments. Adding additional vSAN nodes also increases resiliency due to more failure tolerations, and performance gains due to more resources being available. Another benefit of HCI is that clusters with different hardware can be added in the same environment, for example to add on the ability to use GPUs, for companies to distribute higher graphics performance for specific VDI users that need these capabilities.

Why HCI works well with VDI:

- HCI can be deployed with as few as two nodes, enabling lower costs for even a small VDI environment
- The scale-out design of HCI lines up with the scale out abilities of VDI deployments
- HCI can rapidly scale compute and storage capacity as much as is required
- Operational costs and time can be saved due to reductions in administrative resources
- GPUs can be added to servers for higher graphics performance needs

Testing Set Up

Intel asked the Futurum Group Labs to evaluate the performance and price / performance of several different CPU/GPU configurations using VMware Horizon VDI together with Login VSI's Login Enterprise product. With **Login Enterprise**, a customizable workload generation tool, workloads can be created to match what a company might see in terms of typical usage generated on their systems by users. The Login Enterprise profiles used were the standard Knowledge Worker (KW) profile, with the addition of the GPU Reference Workload. The KW profile consists of Microsoft Edge, and Microsoft Word / Excel / PowerPoint / Outlook applications. The GPU Reference Workload consists of a 1080p video playback using Windows Media Player, the 3D Viewer application running 3D animations, and an optional live Zoom meeting. This GPU profile was combined with the KW profile (KW+GPU) for 96 users (24 per host / 12 per GPU card).

Physical hardware tested included a 4-Node cluster with vSAN for storage. A second 5-node cluster was used to support the required Login VSI, Horizon, and other necessary VMs to create and manage the infrastructure. This second cluster ensures there are no unnecessary performance impacts on the tests. The specific CPUs tested were the Intel® Xeon Platinum 8562Y+

and 8462Y+ CPUs. The KW+GPU profile was tested with the 8562Y+ and 8462Y+ CPUs, both with and without the Flex 140 GPU enabled. Each node was equipped with 2 TB of DRAM, two 1.6TB P5800X drives along with six 3.84TB P5510 drives for storage. Nodes were equipped with two Intel® Ethernet Controller E810-C for QSFP 100Gb, one for vSAN traffic and one for all other traffic.

Desktop pools were created using VMware Horizon, one for the KW only users and another for the 96 KW+GPU users. For VMs that used GPUs, a Dynamic PCIe device assigned to the GPU was given to the VM via vCenter. This was either enabled or disabled depending on the configuration being tested.

For additional details on the environment, please refer to the Appendix.

Testing Background

A key measurement of the Login Enterprise tests is an EUX score, which is calculated by Login Enterprise. This score is meant to represent the "User Experience", or what an actual user would experience while using a VM in terms of response time of the logon process and applications, among other metrics. Login VSI considers a score of 7.5 or greater to represent "Very Good" performance*. Some of the tested systems could have handled more VDI users, but the EUX drops rapidly the more users that are added to the workload when the system is close to capacity**. Our testing was designed to achieve the highest number of users that would still result in an EUX score of 7.5 or higher. Some of the tested configurations ran out of memory resources which capped the number of VDI users at 956, otherwise more users may have been achieved.

* **Login Enterprise Feature Spotlight – EUX Score**

** **Login Enterprise EUX Score**

To measure VDI efficiency, we calculated the cost per unit of work, which for VDI workloads equates to cost per VDI user. This unit of measure is more important than raw performance, due to the scale-out nature of the VDI running on HCI. If a deployment requires more VDI users, simply adding additional HCI nodes will scale up the number of desktop instances linearly.

VDI Test Results

There are a variety of reasons why companies would have users that need higher graphics performance. These users could be in marketing, advertising, CAD users, video editing, etc. For companies with a large number of users and only a subset of those needing higher graphics performance, Intel's Datacenter Flex Series 140 GPUs offer an affordable choice to greatly improve the performance for those users. Important results from the GPU testing include:

- Use of Flex 140 GPUs can offload as much as 50% of CPU resources for individual users running graphics intensive applications
- Individual users can see up to a 3.75x increase in FPS during graphics intensive applications
- CPU Utilization was lowered up to 35% (8462Y+) and 40% (8562Y+) for all 96 KW+GPU users averaged together during the graphics intensive portion of the test
- Intel Flex 140 GPUs do not require an additional software license, unlike a leading competitor

Futurum Group Comments: *The addition of Flex 140 GPUs improved the performance of the graphics intensive applications significantly, increasing the frame rates up to nearly 4X, with an insignificant increase in cost per user, under 2%, going from $1048 per user without a GPU to a cost per user of $1060 with a GPU. Additionally with the cost per user of the 8562Y+ and 8462Y+ being the same with GPUs, the better choice is the 8562Y+ with its increased CPU efficiency.*

Test Details

Shown below in Table 1 are results for 4 test scenarios, with and without Flex GPUs for two different CPUs. We show the current fifth generation Xeon processors (8562Y+ CPUs), compared to prior fourth generation Xeon systems (8462Y+ CPUs). Two EUX scores are shown, one for the Knowledge Worker (KW) workload, and the other for the Knowledge Worker workload plus GPU workload (KW+GPU).

These results were collected when the systems had the maximum number of users supported for that configuration running. There are two EUX scores shown for each configuration because there were two Login Enterprise appliances needed to run two separate workloads, one for the KW only users and the other for the 96 KW+GPU users. All tests had 96 KW+GPU users, regardless of a GPU device being used or not. (As of this writing it is not possible to do two separate load tests from one Login Enterprise appliance.)

System Configuration Per Node (4 Node Cluster)	VM Configuration	Total Number of VDI Users	KW / KW+GPU EUX Score	3 Year Cost Per VDI User
2X Intel 8462Y+ CPU 2 TB DRAM	2 vCPU X 8GB RAM	936	7.5 / 7.5	$1,048.86
2X Intel 8462Y+ CPU 2 TB DRAM 2X Flex 140 GPUs	2 vCPU X 8GB RAM	956	7.6 / 7.5	$1,060.48
2X Intel 8562Y+ CPU 2 TB DRAM	2 vCPU X 8GB RAM	956	7.5 / 7.5	$1,045.42
2X Intel 8562Y+ CPU 2 TB DRAM 2X Flex 140 GPUs	2 vCPU X 8GB RAM	956	7.7 / 7.6	$1,060.48

Table 1: Intel CPU/GPU Comparison: Horizon VDI Performance

For an additional $12 to $15 more per user (CPU dependent), approximately 1% of the cost per user, a large improvement in graphics performance can be achieved. While the 8462Y+ and 8562Y+ had the same number of users with Flex 140 GPUs (956) due to memory resource limitations, the 8562Y+ did have a higher EUX score.

The following Table 2 summarizes GPU related metrics for a single user among 96 total KW+GPU users. A random user near the middle of the workload was chosen for this data to give a more accurate representation of what a typical user will see in terms of performance. Once again, these metrics were collected while the systems were under load with the maximum number of supported users running for each configuration respectively.

System Configuration Per Node (4 Node Cluster)	Max FPS During 3D Viewer Application	Max FPS During Media Player Application	Max CPU Utilization During 3D Viewer Application	Max CPU Utilization During Media Player Application
2X Intel 8462Y+ CPU 2 TB DRAM	17	15	100%	100%
2X Intel 8462Y+ CPU 2 TB DRAM 2X Flex 140 GPUs	60	39	67%	47%
2X Intel 8562Y+ CPU 2 TB DRAM	16	17	100%	99%
2X Intel 8562Y+ CPU 2 TB DRAM 2X Flex 140 GPUs	60	35	66%	50%

Table 2: Intel CPU/GPU Single User Performance Comparison

The following charts illustrate the benefits gained by using Flex 140 GPUs. Figures 1-4 are comparing results for all 96 KW+GPU users averaged together for 8462Y+ and 8562Y+ CPUs, both with and without GPUs.

When the GPU applications start up, the Media Player video playback runs first, followed by the 3D Viewer application, and then Zoom. The KW applications (not graphics intensive) start up again after the Zoom application. It is also important to note that Horizon caps FPS at a maximum of 60, and the FPS of video meeting applications such as Zoom, Skype, and Teams to 25 FPS.

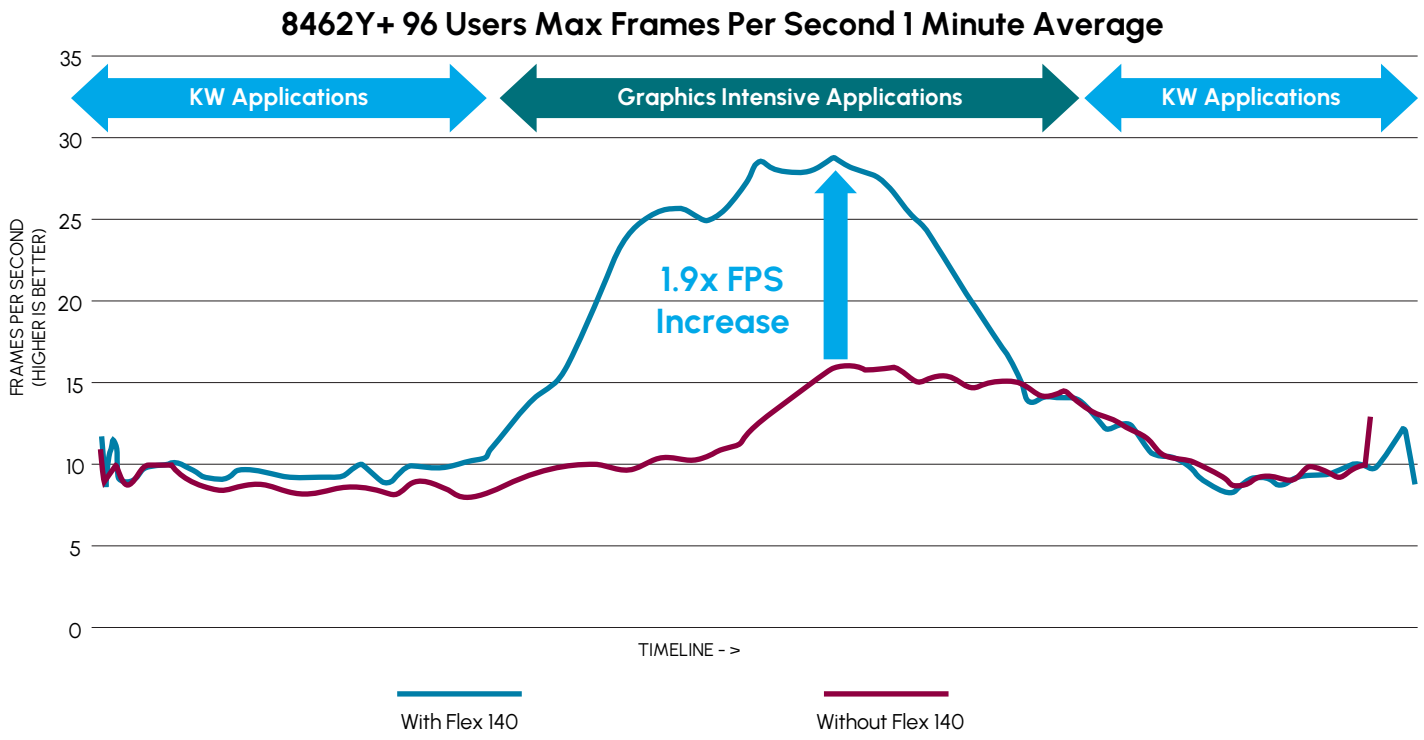

Figure 1: 8462Y+ Max FPS With and Without Flex 140 GPUs - 96 User Average

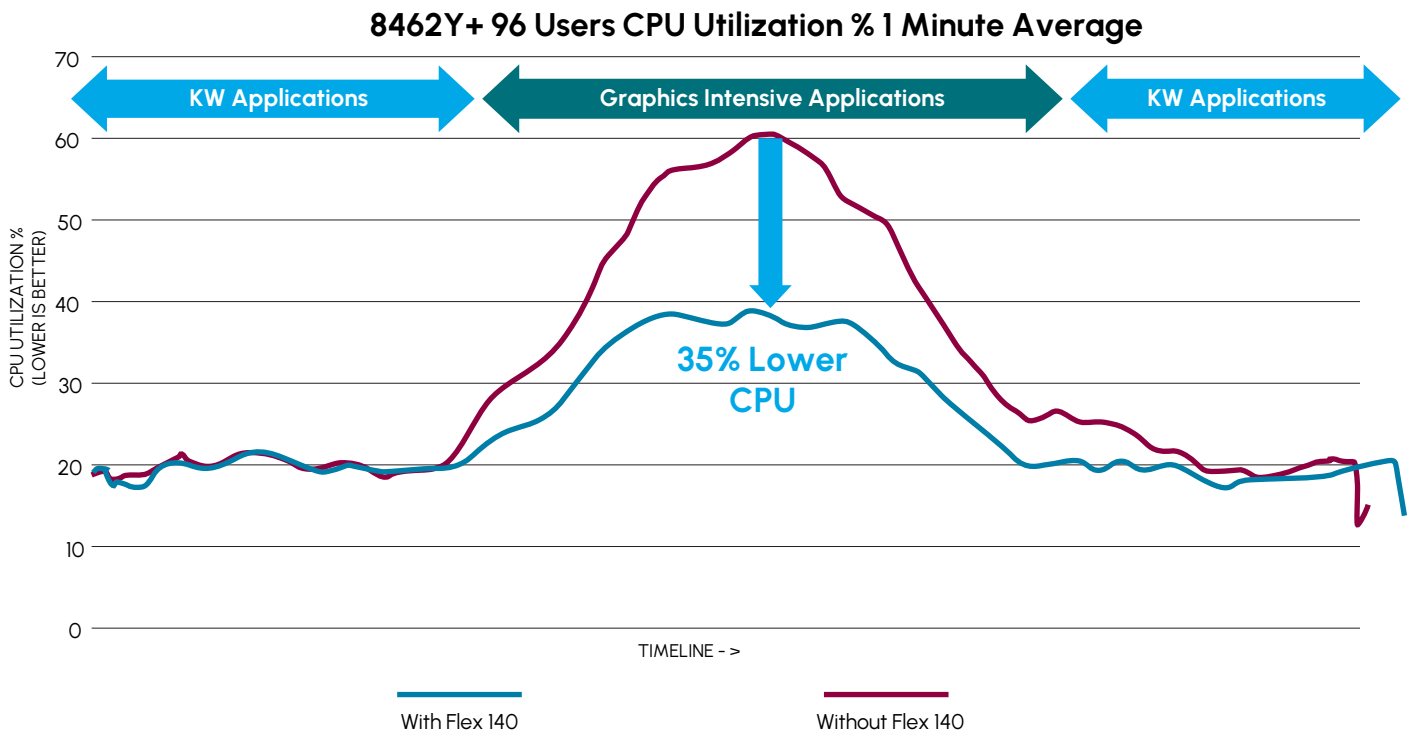

Figure 2: 8462Y+ CPU Utilization % With and Without Flex 140 GPUs - 96 User Average

Using Flex 140 GPUs results in nearly double the frame rates while using up to 35% less CPU across all 96 KW+GPU users with 8462Y+ CPUs.

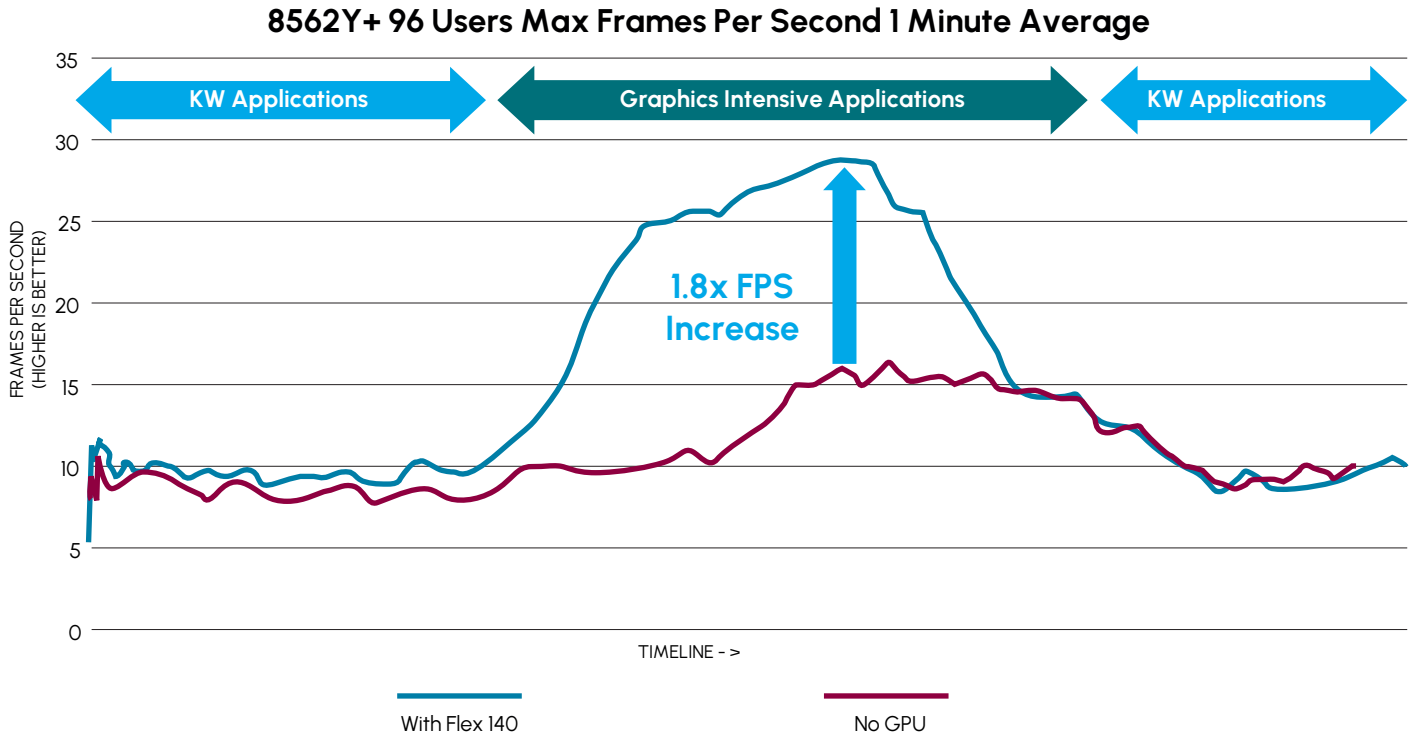

Figure 3: 8562Y+ Max FPS With and Without Flex 140 GPUs - 96 User Average

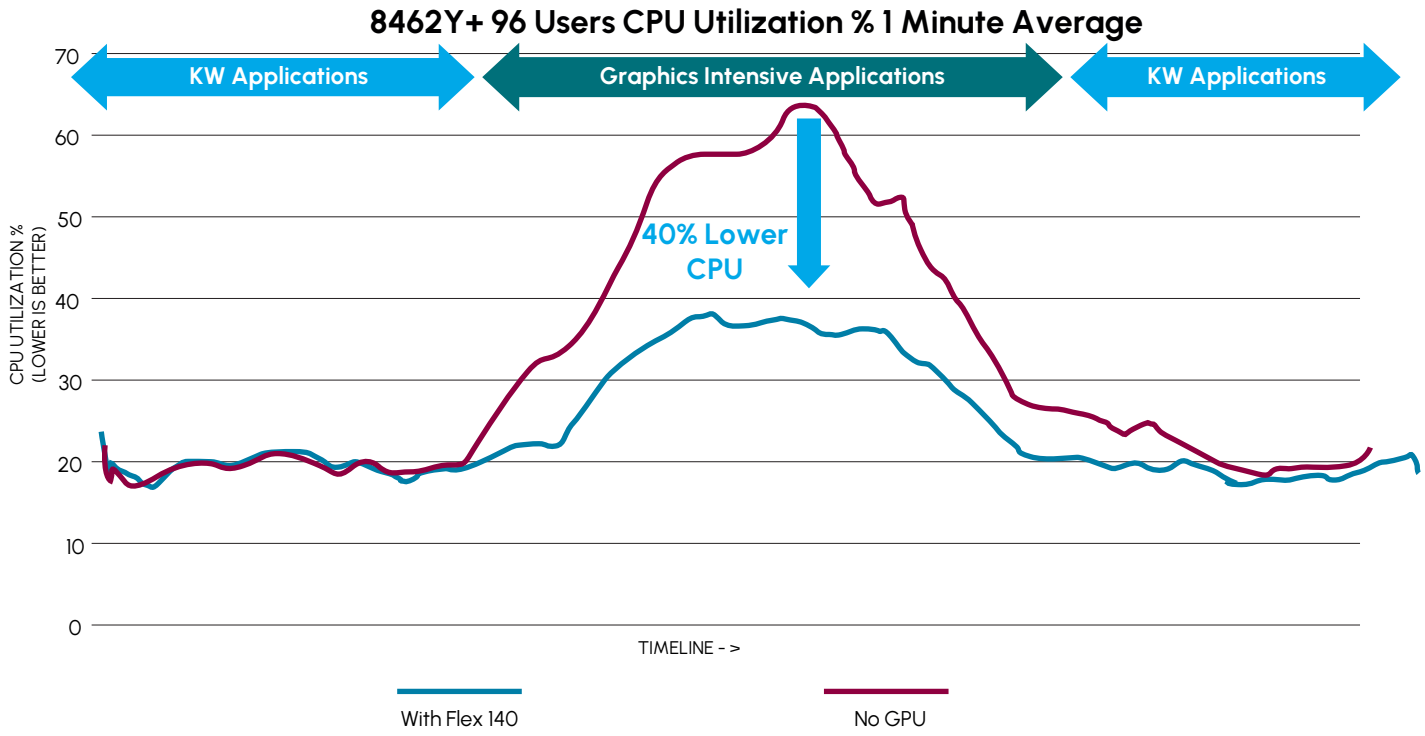

Figure 4: 8562Y+ CPU Utilization % With and Without Flex 140 GPUs - 96 User Average

Similar to the 8462Y+, Flex 140 GPUs with the 8562Y+ results in nearly double the frame rates and up to 40% less CPU utilization across all 96 KW+GPU users.

Figures 5-8 are comparing results *for a single user among the 96 KW+GPU users* for 8462Y+ and 8562Y+ CPUs, both with and without GPUs. A random user near the middle of the workload was chosen for these metrics. *Again, these results are collected with the systems under load, meaning the maximum number of users supported for that configuration are running.*

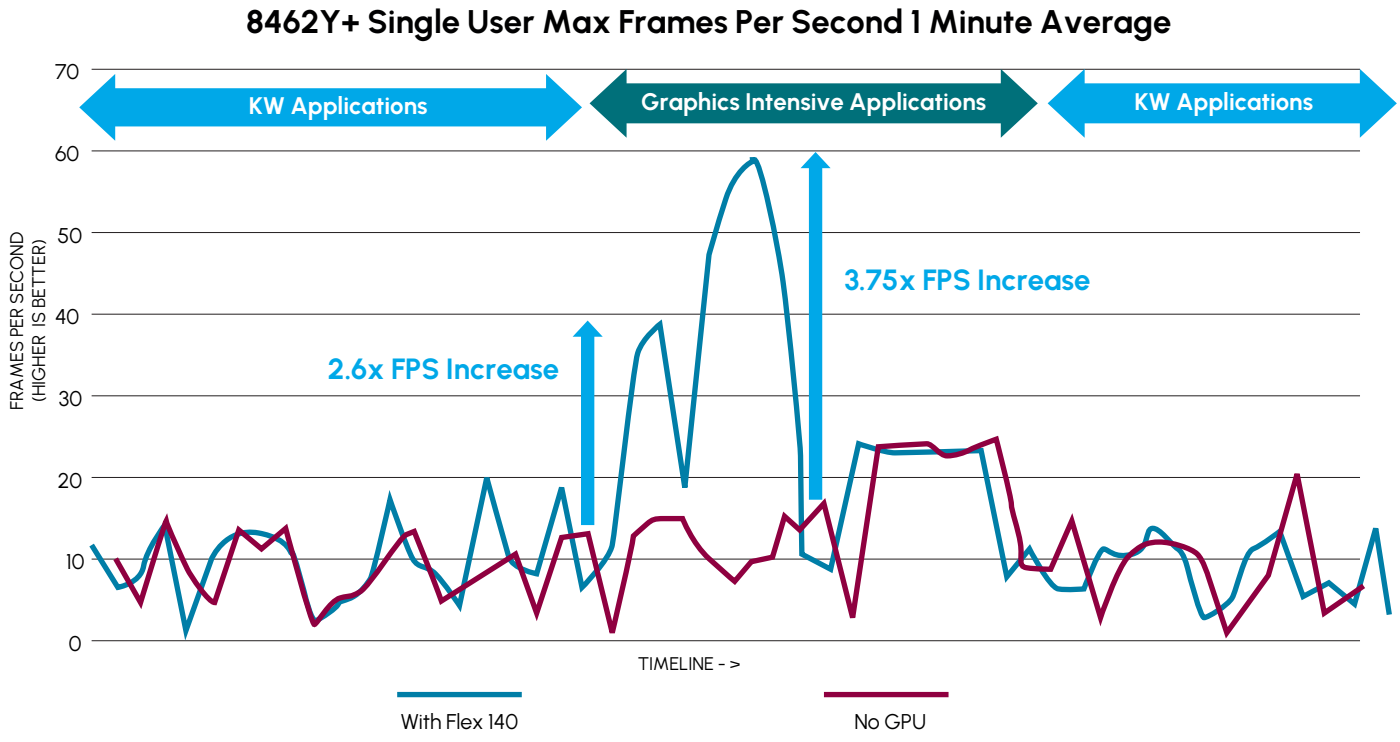

Figure 5: 8462Y+ Max FPS Comparison With and Without Flex 140 GPUs - Single User

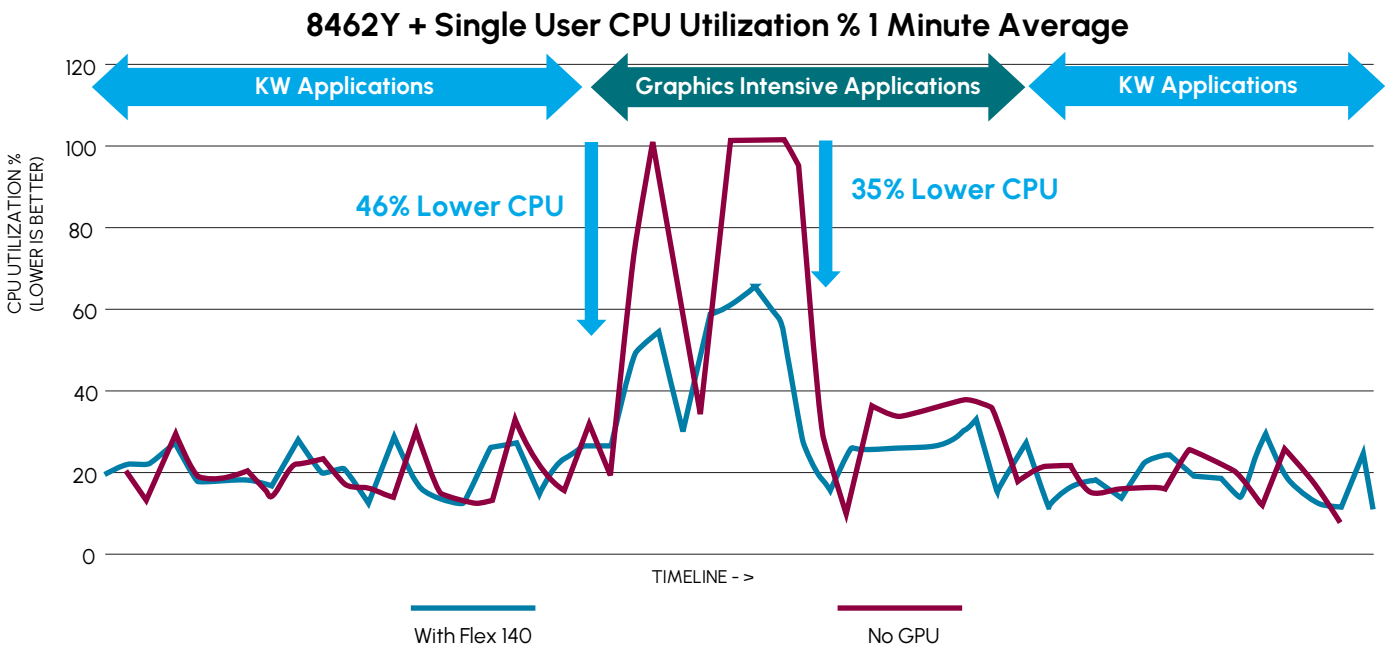

Figure 6: 8462Y+ CPU Utilization Comparison With and Without Flex 140 GPUs - Single User

Flex 140 GPUs with 8462Y+ CPUs provide a large increase in FPS for users during GPU intensive applications while using significantly less CPU resources.

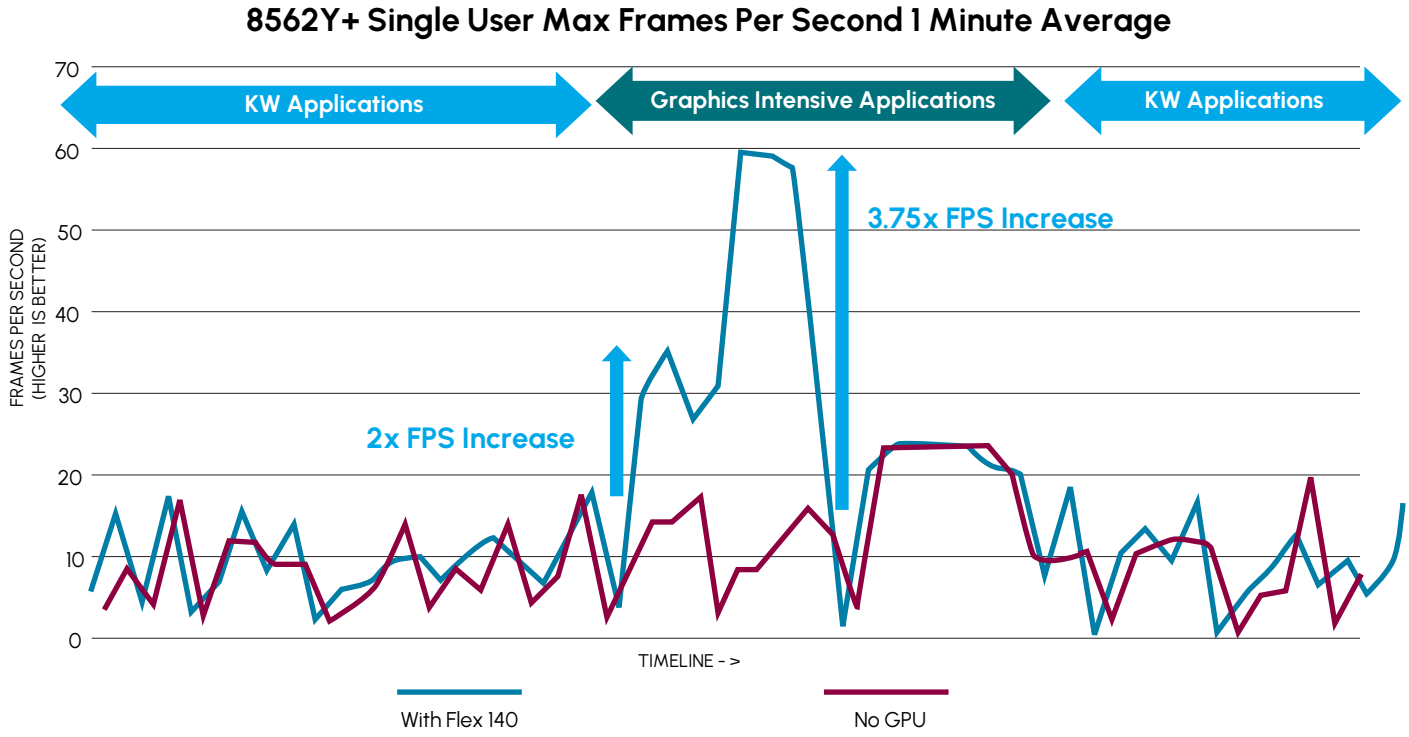

Figure 7 (8562Y+ Max FPS Comparison With and Without Flex 140 GPUs - Single User)

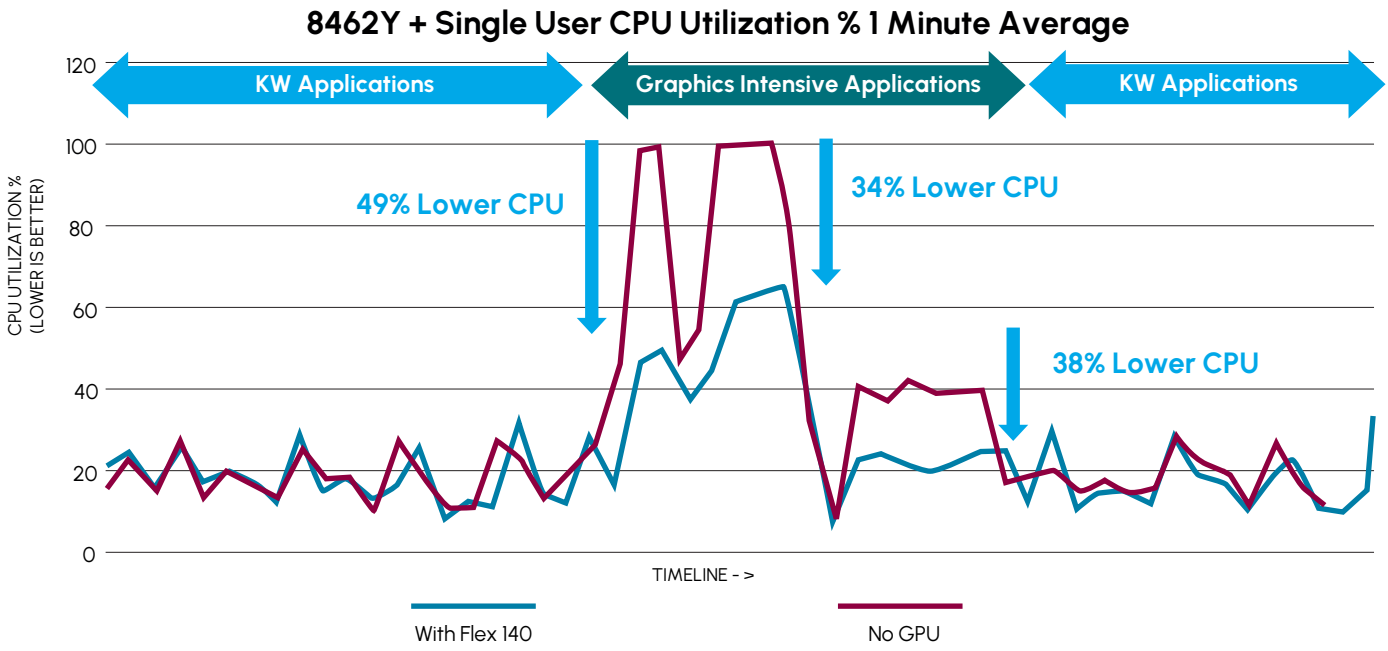

Figure 8 (8562Y+ CPU Utilization Comparison With and Without Flex 140 GPUs - Single User)

Flex 140 GPUs again show much improved frame rates and lower CPU usage with the 8562Y+.

Results Analysis

Configuring VDI instances with virtual GPUs, backed by physical Intel Flex 140 GPU cards in the test configurations provided a dramatically better graphic experience, increasing the frame rate up to 60 frames per second (FPS) during the graphics intensive 3D Viewer application compared to ~10 FPS for users without GPUs (Figures 5 and 7). This was the case for both the 8462Y+ and 8562Y+ CPUs.

Additionally, utilizing a GPU lowered the CPU utilization for individual users by 35% and 34% when using the 8562Y+ and 8462Y+ respectively during the 3D Viewer application. Similarly, CPU utilization was lowered by 49% and 46% during the Media Player HD Video playback when using GPUs (Figures 6 and 8). This allows for significant offloading of CPU resources to be available for other users. It is important to note that it is possible users could have seen above 60 FPS, however Horizon limits FPS to a maximum of 60.

Looking more closely at Figures 5 and 7 reveals details about how Flex 140 GPUs benefitted the users. The Media Player video playback (first large spike) shows a nearly 2x increase in frame rates for the duration of the app when using a GPU. During the 3D Viewer application (second large spike near the middle), while the max FPS of the non-GPU VMs was about 15 for the 3D Viewer app, the average frame rates during that time were closer to 10. Meanwhile the frame rates were close to 60 with Flex 140s for most of the duration of the 3D Viewer app. This indicates a much more stable performance when using GPUs. The FPS cap for Zoom (large plateau after the 3D spike) can clearly be seen with both GPU and non-GPU VMs coming in at 25 FPS. However, less CPU is used during this application with GPUs enabled.

Figures 6 and 8 show similar details in that CPU usage was flatlined at 100% for most of the 3D Viewer application when no GPU was present in the VMs. The Media Player application also saw a peak of 100% CPU usage. When an application uses 100% of the CPU capacity for long periods, it can lead to slow response times for users, among a variety of other issues. When GPUs are added, CPU utilization drops significantly making for a much smoother user experience.

Across all 96 users running the KW+GPU workload, the average CPU utilization was lowered up to 35% and 40% for the 8462Y+ and 8562Y+ respectively during the graphics intensive period of the test (Figures 2 and 4). This again highlights how using GPUs can significantly reduce CPU overhead for the cluster.

Futurum Group Comments: *Overall, we found that the benefits of adding a GPU were similar for both the 8462Y+ and 8562Y+ CPUs. Using Flex 140 GPUs allows for more users and better performance, due to offloading the CPU utilization. The use of Flex GPUs allows companies to add more users while increasing their performance, or provide a better graphics performance while maintaining the number of VDI users.*

Summary

With the continued popularity of VDI being used in HCI environments, companies are expected to invest in this technology for the foreseeable future. Intel's newest generation of the Platinum series CPU, 8562Y+ provides the better overall performance and cost per user than the 8462Y+. The newest 5th generation Xeon Platinum processors provide more CPU headroom to enable more users running more CPU intensive applications or increase the performance of the existing infrastructure.

Additionally, adding Intel Flex 140 GPU cards can enhance users experience, particularly for graphic intensive workloads while offloading the CPU resources for other tasks. Most importantly, these benefits can be obtained for an insignificant increase of approximately 1% of the total configuration cost.

Specifically, the use of Intel Datacenter Flex 140 GPUs can significantly improve the graphics performance for users while lowering CPU usage, freeing up resources for other users. Individual users seeing a nearly 4x increase in FPS in graphics intensive applications while using ~35% less CPU is no trivial data point. With CPU usage dropping by ~40% across all GPU users during graphics intensive applications, it is obvious why the use of GPUs is important for companies with users that need higher graphics performance.

These important takeaways lead to the conclusion that using Xeon Platinum CPUs along with Flex 140 GPUs can provide high performance for both CPU and GPU needs. For companies looking to increase their VDI performance capabilities, increase their VDI capacity, or add higher graphics performance for users that need it, the newer generation of Intel Xeon CPUs along with Flex 140 GPUs offer a clear method of meeting these needs cost effectively.

Appendix

The workload profiles were run using a 45-minute "Load Test". The Load tests runs each application in the list then loops for as long as is required. The Load Test included a 15-minute logon period, the time it takes for all users to be logged on. In total 60-minutes elapsed for the tests. The Load Test window was made long enough so the last VM to log was able to finish all applications in the list before being logged off at the end of the test period. The results collected were the median of 3 runs for each configuration.

The Media Player video provided by Login VSI was set on a loop to run for the allotted time. The video was encoded at 30 FPS.

For the live Zoom meeting, a Zoom account was created with an extended license for large virtual meeting rooms. The meeting was started when the load tests began, and a live web camera of an office setting with real people was used to ensure a realistic meeting scenario. Users would come in and out of the meeting as they reached and finished the Zoom application. There were approximately a max of 40 users in the meeting at any one time. The length users spent in the meeting was 5 minutes.

The tested environment was configured with a 4-node HCI deployment, along with vSAN and ESXi managed through vCenter. Each node had 2 CPU sockets with 2TB of DRAM.

For the tests using the KW+GPU profile, two Intel Datacenter Flex Series 140 GPU cards were installed in each node. The GPUs were configured with the V6 profile and SR-IOV passthrough was enabled for each graphics device in the ESXi settings. The V6 profile supports 12 users per card at 1080p definition with 1GB of buffer memory. Note that two separate Login Enterprise Appliance VMs were required for the KW+GPU tests, one to run the KW users and the other to run the 96 KW+GPU users. As of this writing, it is not possible to run two separate load tests at once from a Login Enterprise Appliance.

For all the KW+GPU tests, VMs were configured with 2 vCPUs and 8GB of RAM. 96 VMs were given Dynamic PCIe graphics devices through the VM settings in vCenter when testing GPUs. When not testing GPUs, the device was removed from the VMs. This was to show the benefits of using GPU devices for users running a graphics intensive workload. For the GPU tests (both with and without the device present in the VMs), 96 users ran the KW+GPU profile, while the rest of the users ran the original KW profile.

Registry Keys were set in the Horizon config to allow for up to 60 FPS.

VMs had Windows 11 21H2 as the OS, including Microsoft Office products and the Edge internet browser. Windows Media Player, 3D Viewer and Zoom applications were installed for the KW+GPU users.

Test Environment Details

Testing was performed from August 2023 through February 2024, based upon equipment availability. The test environment utilized the following hardware, software and applications outlined below.

Hardware and Infrastructure

- A Test cluster of 4-nodes, with a 6-node cluster for supporting automation, Figure 4 below
- VMware vSphere 8.0 U2, with vSAN Advanced (ESXi + vSAN on hosts, with vCenter server)
- ESXi 8.0 U2 Build 22380479; vCenter 8.0 U2 Build 22385739
- Multiple VMs configured as required to support Horizon Virtual Desktop as required

A high-level overview of the test configuration is shown in Figure 9 below.

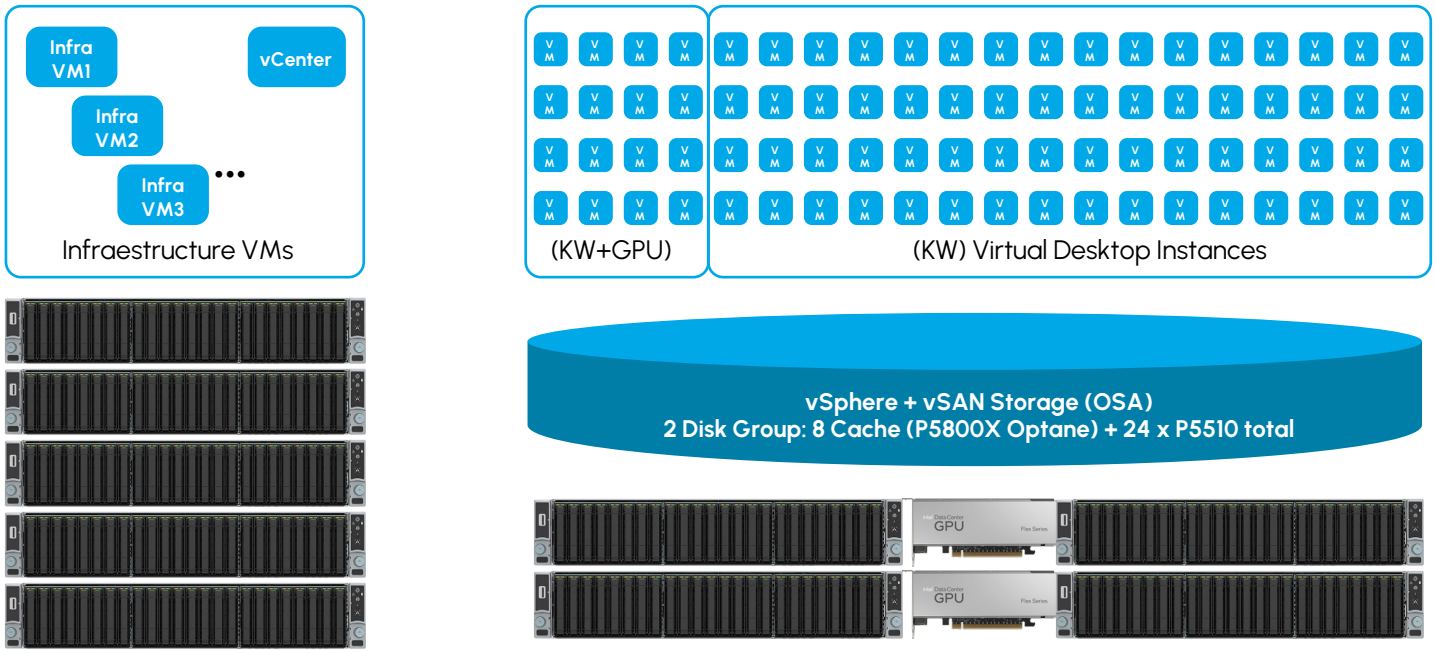

Figure 9: Test Setup for Horizon VDI on Intel

Windows, VMware and LoginVSI Version

- VDI Client OS: Windows 11 Pro for Workstations 21H2
- Server OS for VM's: Windows Server 2022 Enterprise
- VMware Horizon 8 Desktop, Version 2303
- VMware 8.0U2; ESXi build 22380479; vCenter build 22385739
- Login Enterprise, version 5.7.2
- Microsoft Office Version – 2304 Build 16327.20214 Click-to-Run

SUT Cluster Configuration

- 1 cluster, 4-node HCI cluster for VDI apps with vSAN storage
- Intel Server configuration:
  - 2 socket Intel Xeon 8462Y+, 2 Socket Intel Xeon 8562Y+Support for 32 DIMM slots
  - Support for 32 DIMM slots
  - 2 Intel E810-CQDA2 100G NICs per host
  - 2 TB DRAM
- Intel vSAN configuration: 2 disk groups per node, each with 1 cache + 3 capacity SSDs
  - 2 x 1.6TB P5800X Optane NVMe SSDs
  - 6 x 3.84TB D7-P5510 NVMe SSDs

Client Cluster Configuration – Infrastructure Applications

- One, 5-node HCI cluster of Intel systems with vSAN was used for supporting infrastructure
  - System hardware and VMware licenses were not included in pricing
  - All required supporting infrastructure ran on Client Cluster
  - Client hosts also used the same ESXi versions as SUT hosts
- Infrastructure VM's included following:
  - Horizon Connection Server VM to support managing virtual desktops
  - Horizon version 2309 - Build 22629722
  - Windows VM with PDC role, along with AD, DHCP and DNS services
  - Other required VM's, including vCenter and 2 Login Enterprise Appliances
  - 96 Login VSI "Launcher" VMs to initiate VDI workloads

Application Workload

The application workload tool, Login Enterprise was used to generate an application workload

- Login VSI Enterprise version 5.7.2
- Knowledge Worker (KW) profile application workload applications:
  - Edge Browser
  - MS Excel
  - MS Outlook
  - MS PowerPoint
  - MS Word
- Knowledge Worker + GPU (KW+GPU) profile application workload applications:
  - Edge Browser
  - MS Excel
  - MS Outlook
  - MS PowerPoint
  - MS Word
  - Windows Media Player 1080p Video Playback
  - 3D Viewer Application with bee.glb animations
  - Zoom client with a live meeting running, including live desktop video from host (VDI users did not have cameras or audio)
- VM Hardware configuration settings
  - 2 vCPU and 8 GB vRAM, full 8 GB memory reservation
  - Dynamic PCIe graphics device added for VMs running KW+GPU workload
  - Intel Datacenter Flex 140 GPU device driver installed on VMs running KW+GPU workload
- Windows 11 OS with VMware tools installed (Windows 11 Pro for Workstations 21H2)
- VMware Horizon agent installed (agent for Horizon 8 2309)
- Multiple VMs configured as required to support Horizon Virtual Desktop as required
- Horizon version 2309 - Build 22629722

Test Setup Overview

- Setup HCI infrastructure per HCI hardware and VMware recommendations
  - vSAN storage policy utilized was 'vSAN Default' which uses RAID-10, without deduplication, compression or encryption enabled.
- Setup and deploy Microsoft Domain and Active Directory, along with DHCP and DNS servers
  - A single primary domain controller (PDC) was configured with AD, DHCP and DNS
- Setup and deploy required VMware Horizon infrastructure applications on "Client Cluster"
  - Horizon 8 controller
- Setup initial virtual desktop VM
  - Create a "golden image" Windows 11 VM
  - Optimize for Horizon VDI per recommendations
  - Install required applications (listed previously)
  - Add group policy object and other customization required for Login Enterprise
- Create Login Enterprise "Launcher VMs" on "Client Cluster" as required
- Utilize Horizon manager to create clones of "golden image" VM on "SUT Cluster" as needed
- Use "Login Enterprise" tool to generate application workload on "SUT Cluster" per test

TCO Pricing Data

To provide comparable estimated pricing using widely available pricing data, we utilized ThinkMate.com1, configuring servers consistent with our testing. We varied processors, system memory, and in some cases storage components. When tested components were not available on ThinkMate, we used pricing data from other sources (i.e. NewEgg, Amazon, etc.). When using previous generation of a current product which was no longer available, we used pricing for the current generation product, which was very similar to previous gen pricing. This methodology was applied consistently across the configurations tested.

VMware Horizon pricing was calculated using the following sku(s) and pricing available from CDW: (HAH-CRCUB-36PT0-C1S & HAH-ADCUA-36PT0-C1S, per the VMware Horizon Packaging and Licensing white paper2.

All pricing data was obtained during Sep-Nov 2023 & Feb 2024 and was accurate at the time of publication. All Costs are calculated for a 3-year period of ownership or TCO.

Pricing does not include Windows 11, Zoom, or Office Professional Plus client licensing.

System Configuration Per Node (4 Node Cluster)	Total Number of VDI Users	HW Costs incl 3 yr Support	SW Costs incl 3 yr Support	Total cost - 3 years	$ / Horizon User
2X Intel 8462Y+ CPU 2 TB DRAM	936	$154,023	$827,709	$981,731	$1,048.86
2X Intel 8462Y+ CPU 2 TB DRAM 2X Flex 140 GPUs	956	$168,423	$845,395	$1,013,817	$1,060.48
2X Intel 8562Y+ CPU 2 TB DRAM	956	$154,023	$845,395	$999,417	$1,045.42
2X Intel 8562Y+ CPU 2 TB DRAM 2X Flex 140 GPUs	956	$168,423	$845,395	$1,013,817	$1,060.48

1 - https://www.thinkmate.com/systems/storage

2 - https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/docs/vmw-horizon-pricing-and-packaging-whitepaper.pdf

Important Information About this Report

CONTRIBUTORS

Jonathan Fellows
Performance Validation Engineer | The Futurum Group

PUBLISHER

Daniel Newman
CEO | The Futurum Group

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

CONTACT INFORMATION

The Futurum Group LLC | futurumgroup.com | (833) 722-5337