White Paper

Data Center
Virtualization

Revolutionizing VDI Performance: The Impact of Intel® Data Center GPU Flex Series and SR-IOV Technology

Unlock the full potential of your VDI environment with Intel® Data Center GPU Flex Series and SR-IOV technology, delivering unparalleled performance, scalability, and user experience.

Authors:
Azhar Shaikh
Cloud Software Development Engineer
Anitha Suresh
Director of Engineering

Table of Contents

Introduction

In the dynamic landscape of Virtual Desktop Infrastructure (VDI), the demand for high-performance solutions has reached unprecedented levels. As organizations increasingly rely on VDI for flexible work environments, the imperative to deliver exceptional user experiences while optimizing resource utilization has never been more critical. To meet these evolving needs, the integration of dedicated hardware accelerators has become a cornerstone of modern VDI server solutions.

Dedicated hardware accelerators, specifically the Intel® Data Center GPU Flex Series, play a pivotal role in enhancing VDI performance. From delivering consistently high frame rates and uncompromised image quality to supporting modern remote display topologies, the Intel® Data Center GPU Flex Series sets a new standard for VDI performance.

In addition to showcasing the capabilities of the Intel® Data Center GPU Flex Series, this whitepaper also provides an overview of the SR-IOV driver within VMware's ESXi stack, GPU profiles and Intel® Device Manager for VMware* vCenter Server.

Furthermore, this paper presents insights from practical demonstrations using VDI environments. By utilizing various profiles tailored for specific use cases, we demonstrate the real-world impact of Intel® Data Center GPU Flex Series on VDI performance. These demonstrations serve as tangible examples of the enhanced user experiences and optimized resource utilization achievable through the integration of dedicated hardware accelerators.

Through comprehensive analysis and empirical evidence, this paper delves into various aspects of VDI performance, including end-user experience, GPU utilization, hardware encode performance, scalability using Single Root I/O Virtualization (SR-IOV), CPU-GPU offload, and Total Cost of Ownership (TCO). By examining these key metrics, organizations can gain valuable insights into optimizing their VDI deployments for maximum efficiency and productivity.

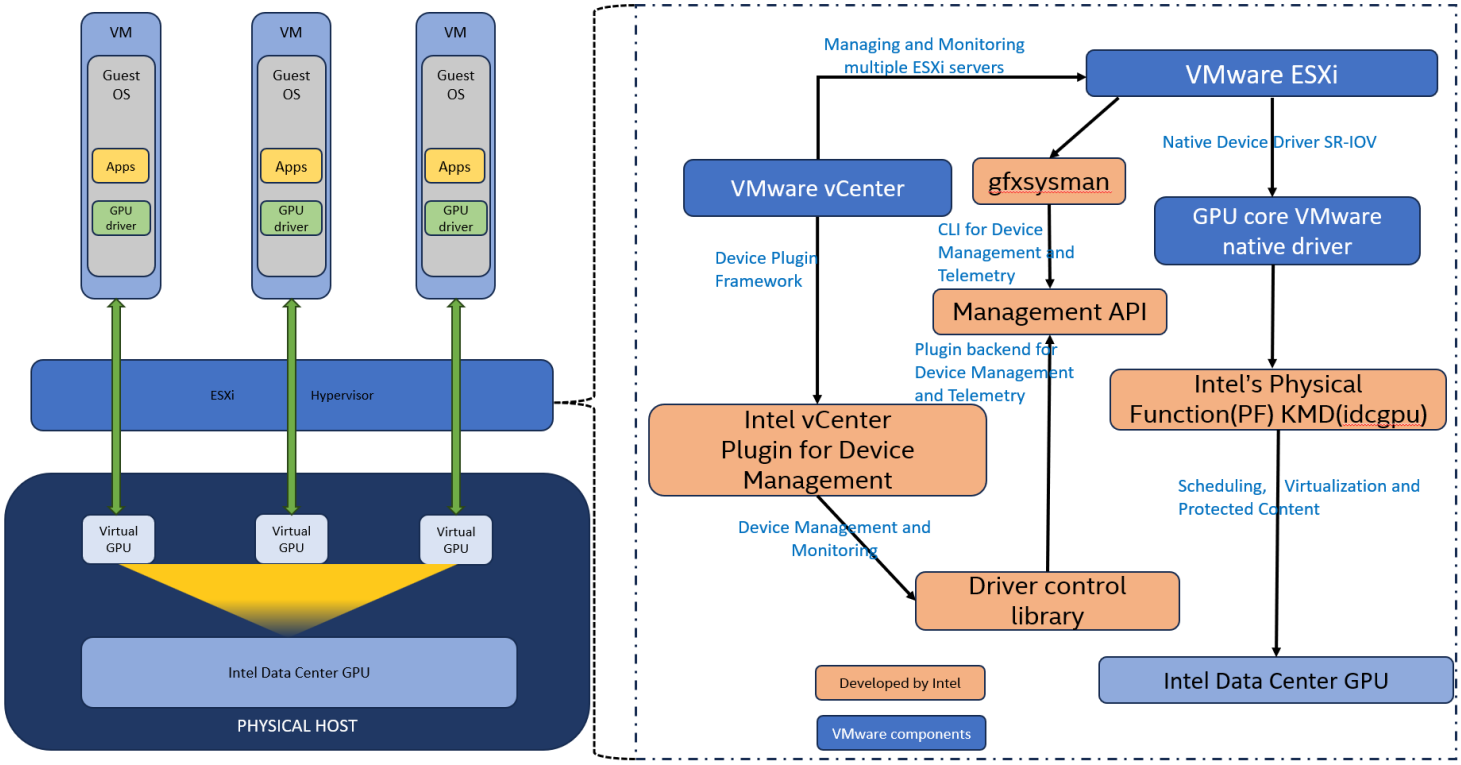

Figure 1. High Level Architecture of PF KMD and other components in VMware's ESXi

VDI with SR-IOV Solution Architecture

1. SR-IOV Kernel Mode Driver (KMD)

i. There are multiple VMs connected to the physical host using the VMware ESXi hypervisor.
ii. Each VM has its own guest OS running which loads the GPU driver and runs apps on it. Each GPU driver is connected to a virtual GPU created by SR-IOV.
iii. The PF KMD and all its associated components shown on the right-hand side of the figure above run within ESXi.
iv. The native driver is provided by VMware for managing GPU resources within the ESXi environment.
v. The PF KMD enables Intel® Data Center GPU on VMware's vSphere using SR-IOV technology along with PCI vSphere DirectPath I/O to share a single PF GPU device between multiple Virtual Functions (VFs), and it provide graphics acceleration to workloads in numerous Virtual Machines (VMs). With direct access to the hardware, this driver takes care of:
   - Scheduling
   - Distributing the GPU resources among the virtual machines
   - Playing protected content videos
   - Updating Integrated Firmware Image (IFWI)

   Using VMware's native driver API, this PF KMD driver is managed by ESXi.
vi. gfxsysman is a command line interface for managing devices and collecting telemetry data.
vii. VMware's vCenter is the centralized management platform for VMware's vSphere suite, used for managing multiple ESXi hosts and their virtual machines.
viii. Intel's vCenter plugin for Device Management uses the driver control library, which is the API for managing devices and collecting telemetry data.
   a. This plugin then runs on the vCenter instance allowing the administrator to configure the virtual machines and use the correct profiles as per the desired workload to be run on a particular virtual machine.

2. GPU Profiles

i. The Intel® Data Center GPU Flex Series software solution can control GPU resources and parameters knobs for VFs that belong to each PF. Since the list of parameters is quite large, Intel has organized them in a series of profiles. These profiles are statically loaded in the driver and can be specified for each Intel® Data Center GPU device in the system.
ii. GPU profiles can be tailored to different usage scenarios to optimize performance, memory consumption, and other factors.
iii. Based upon the type of worker, the administrator can create a profile for a Knowledge Worker or a Power User.
iv. The profiles are named as <PRODUCT_NAME>_<LMEM_SIZE>, for example, Flex140_6. LMEM_SIZE is the memory allocated to each VF in that profile.
v. Flex140_1, Flex140_2, Flex140_3, Flex140_6, Flex170_1, Flex170_2, Flex170_4, Flex140_8 and Flex170_16 are the supported profiles.

3. Intel® Device Manager for VMware vCenter Server (vCenter plugin)

i. To scale, VDI solution needs telemetry/management solutions like vSphere Client, vCenter Plugin (made by Intel®) and Horizon Server.

ii. Intel® Device Manager for VMware vCenter Server [i] is a user-friendly way to manage Intel devices on VMware's vCenter platform. It is deployed as a plug-in to vCenter. The plug-in lets you manage and configure devices right from the interface of VMware vCenter. You can:

   b. View information on supported devices
   c. Monitor device counters and sensors.
   d. Customize device configuration.
   e. Save configurations as recipes for future reuse and apply these recipes to all your devices with a single click.

iii. The Intel® Accelerator Management Daemon[1] part of the vCenter plugin needs to be installed on each host to discover and allow device management.

iv. Unigine Heaven benchmark, which is a 3D workload, is run on a virtual machine attached to the Intel® Data Center GPU Flex series using SR-IOV. The below graphs can be accessed from the vCenter. Please check the user guide referenced below, under section 'Viewing Device Information' to navigate, to see the below graphs. These graphs can be seen for each GPU device attached to the vCenter host. GDDR6 is the LMEM_SIZE of each VF.

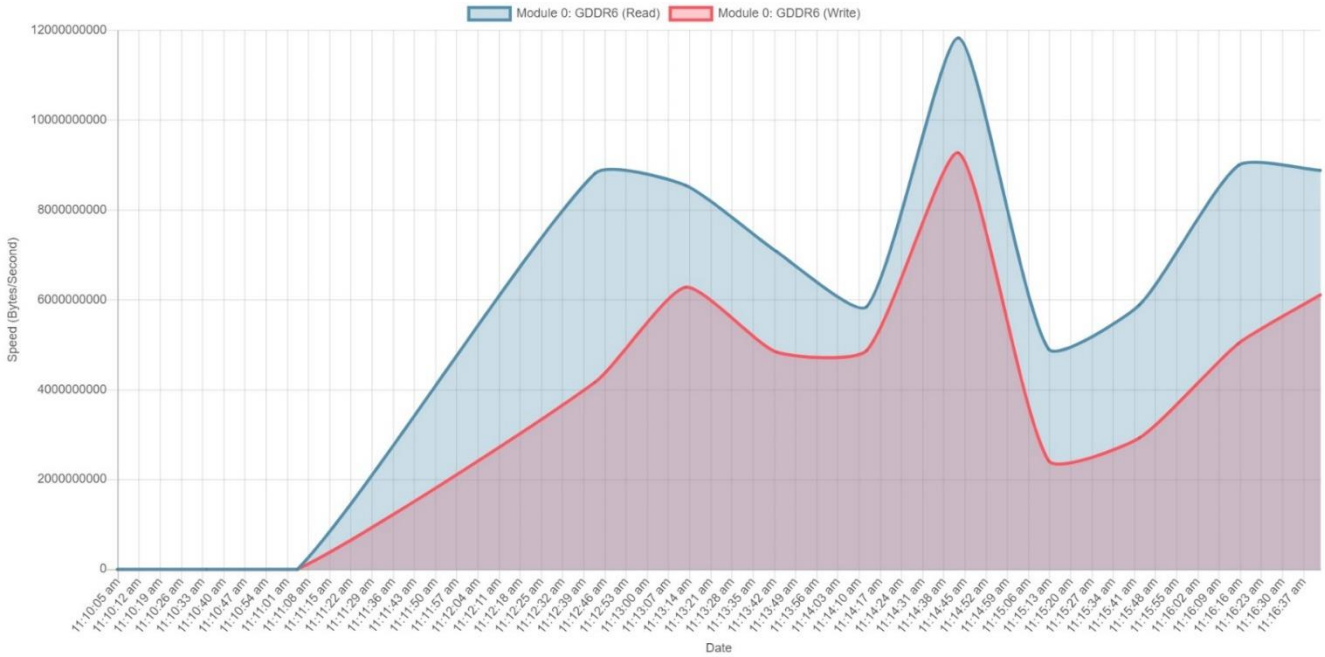

Figure 2. The GPU memory read/write speeds are shown in bytes/second.

[1] Intel® Device Manager for VMware* vCenter Server

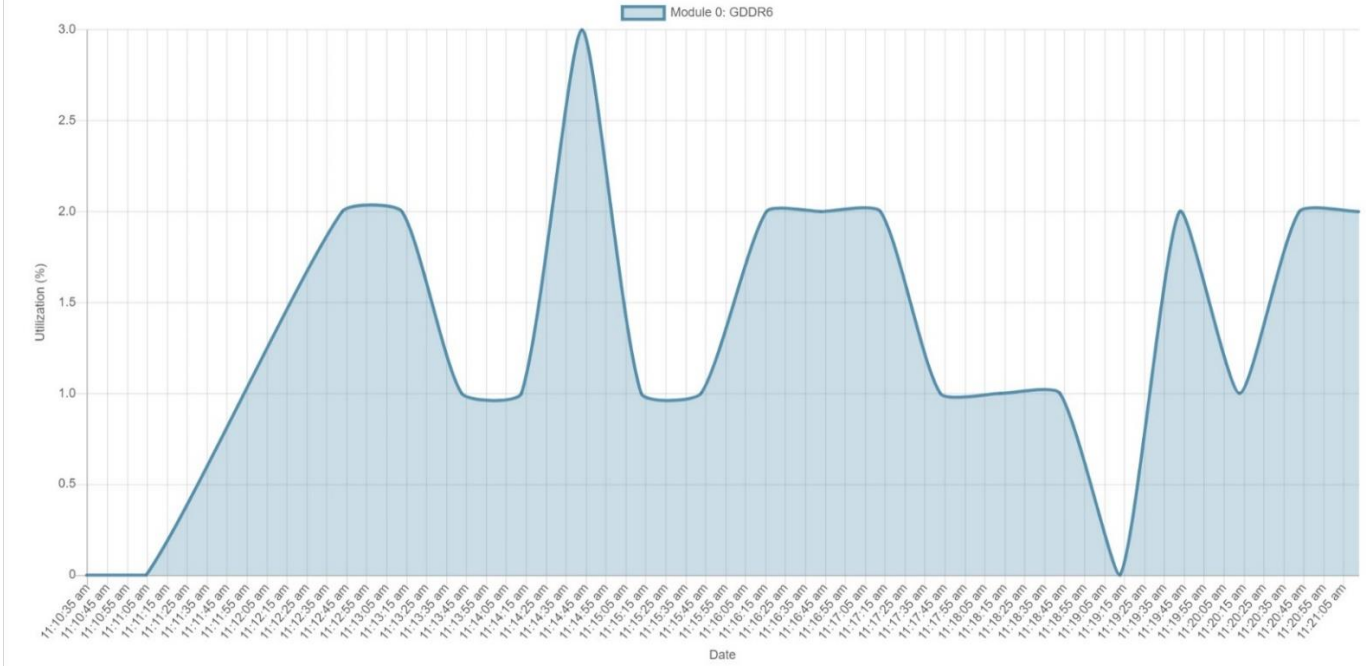

Figure 3. Percentage of GPU memory used.

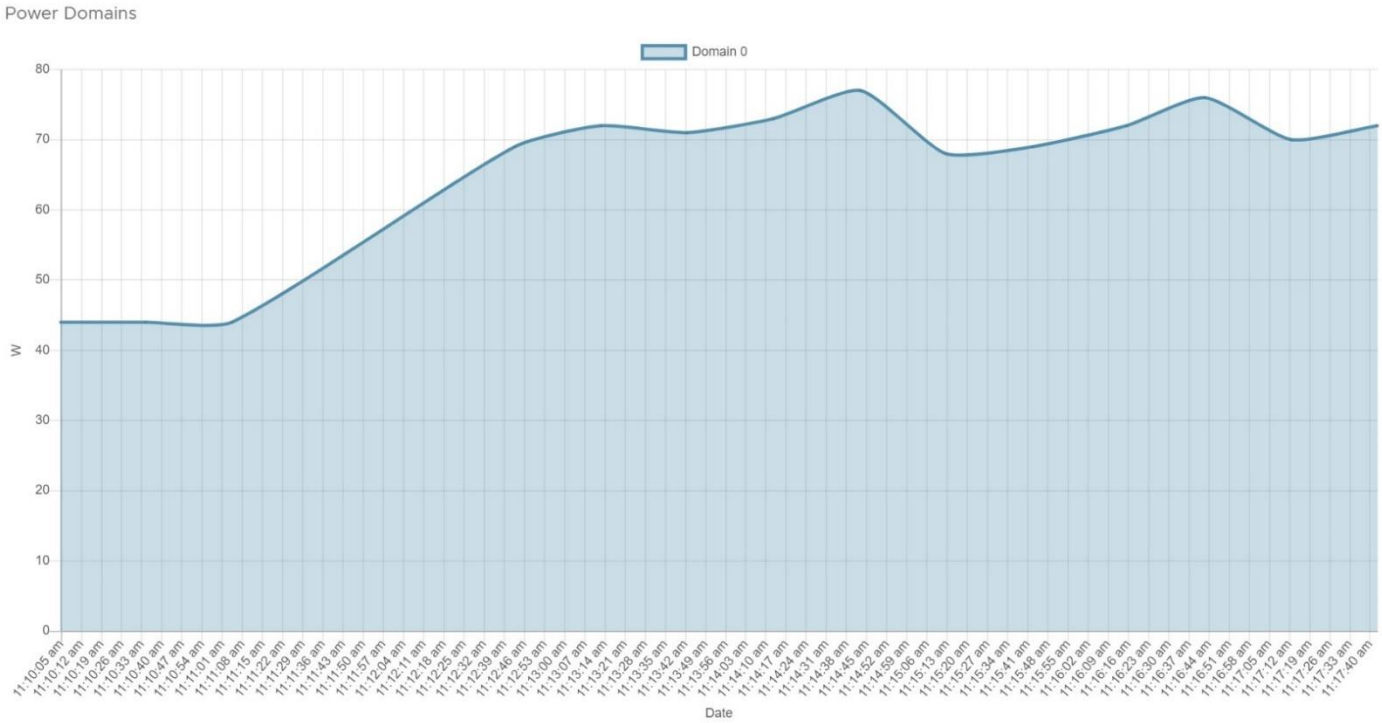

Figure 4. Power consumption of the GPU

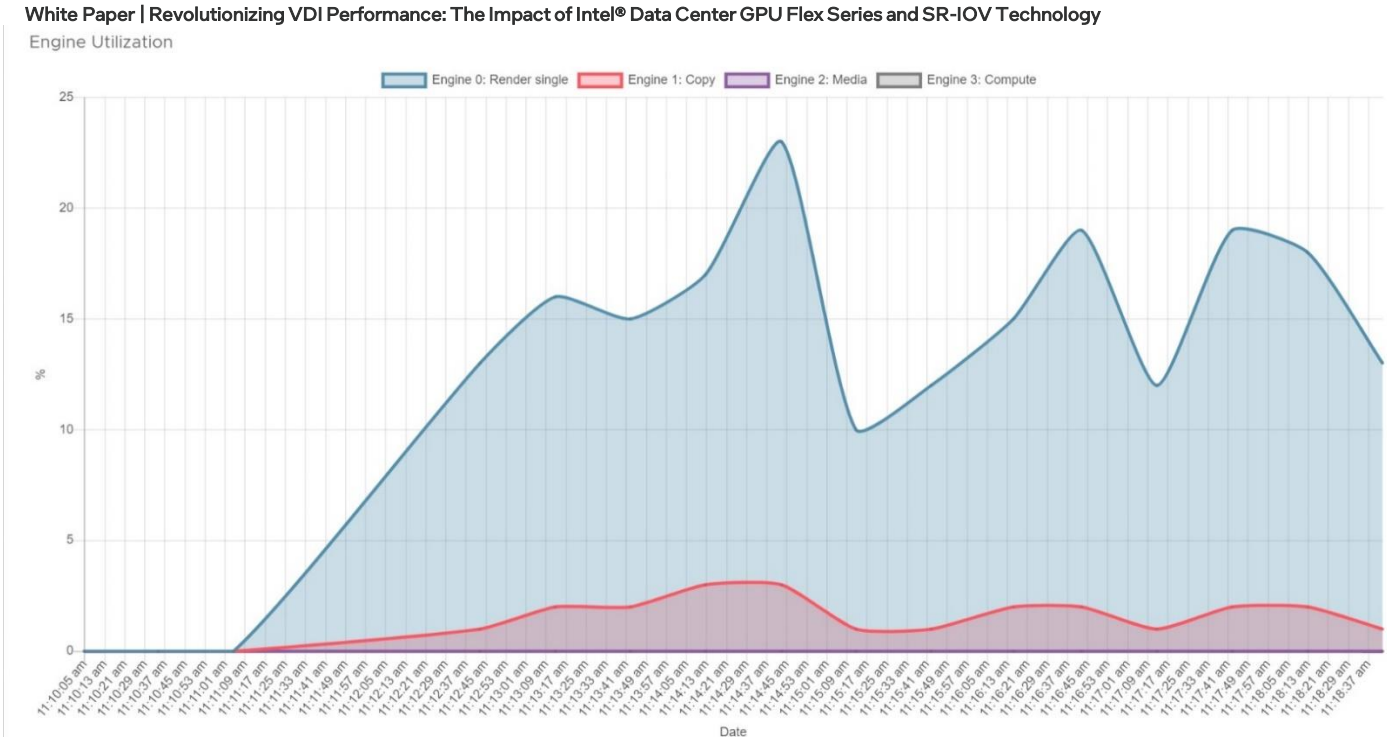

Figure 5. Engine utilization of all VMs on a GPU

As seen in Fig.5 the GPU has render, copy, media, and compute engines. The utilization of each engine can be tracked. In this case since it is a 3D workload, the render and copy engines are active.

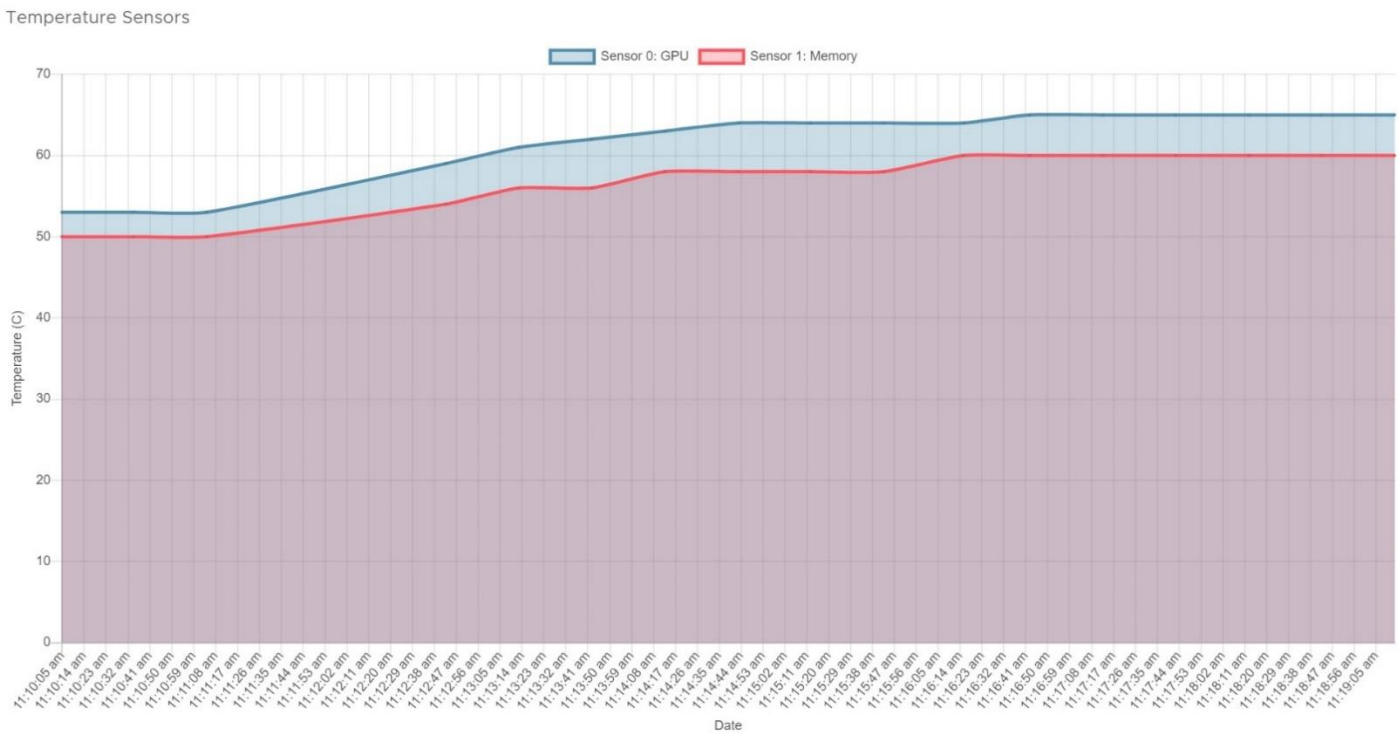

Figure 6. GPU's core and memory temperature

GPU and its local memory's core temperature is shown below. As seen below the temperature increases as the workload starts running.

A user can set a temperature threshold on a GPU and when the temperature crosses the threshold an alert is generated on vCenter. To set the threshold refer 4

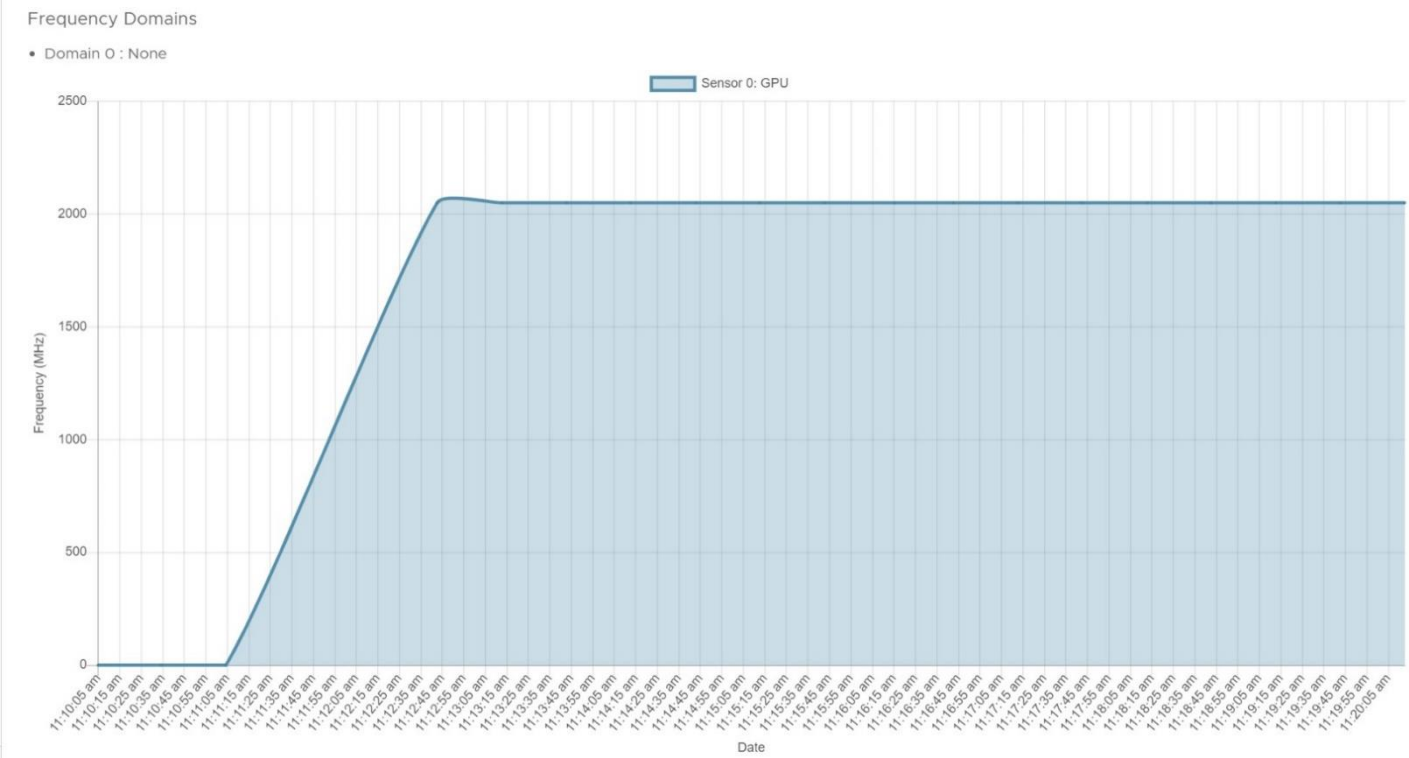

Figure 7. Operating frequency of the GPU.

As seen in Fig7. As the workload increases, the frequency increases and stays at maximum to get the best possible performance.

Benchmarks/Performance

This section goes through the test carried out to demonstrate improved user experience while maximizing users with Intel® Data Center GPU Flex Series in a VDI environment, as measured by higher frame rate, lower latency, and GPU offload. Login Enterprise was used in this testing. It is the de facto industry standard tool for testing VDI environments and server-based computing (RDSH - Remote Desktop Services Host).

1. Test Setup

i. 4 node vSAN Cluster running 32c (8462Y+ (Sapphire Rapids), 8562Y+ (Emerald Rapids)) with 2TB system memory per node and 2 Intel® Data Center GPU Flex 140 per node. Each user was allocated 2 vCPUs and 8 GB memory.
ii. Each Flex Series 140 GPU was configured to support 12 virtual GPUs per card for a total of 96 GPU-enabled VMs across the 4-node cluster. Each Flex 140 has 12GB of memory.
iii. Using the Flex140_1 profile, each virtual GPU will be allocated less than 1GB/user, as some memory is provisioned to the host driver (PF KMD). Virtual GPUs were configured to support an HD display at 1080p. Each user has a persistent session.
iv. Workload 1 is the combination of the standard Knowledge Worker profile and the GPU Reference Workload, which runs on the GPU-enabled VMs (have virtual GPU assigned).
v. Workload 2 is the standard Knowledge Worker profile, which runs on the CPU-only VMs (GPUs are disabled).

2. Results

The frames per second (FPS) for the graphics workload increased from **15FPS to 59FPS**, which is an almost **300%** improvement when using GPU enabled VMs compared to using only CPU enabled VMs

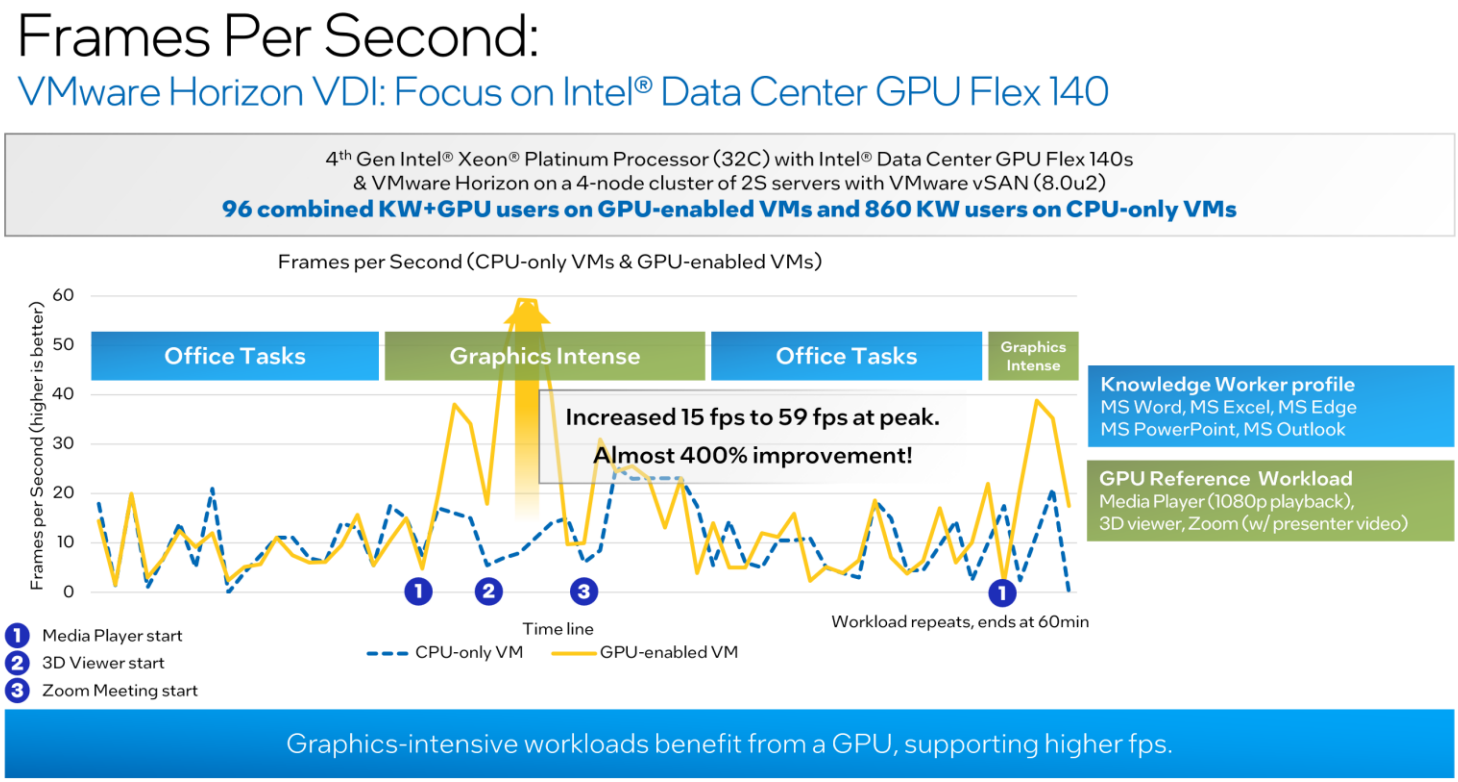

Figure 8. FPS variations captured while running the tests.

The CPU Utilization also decreased with the work offloaded to the GPUs. At peak, the **CPU utilization decreased by 40%.**

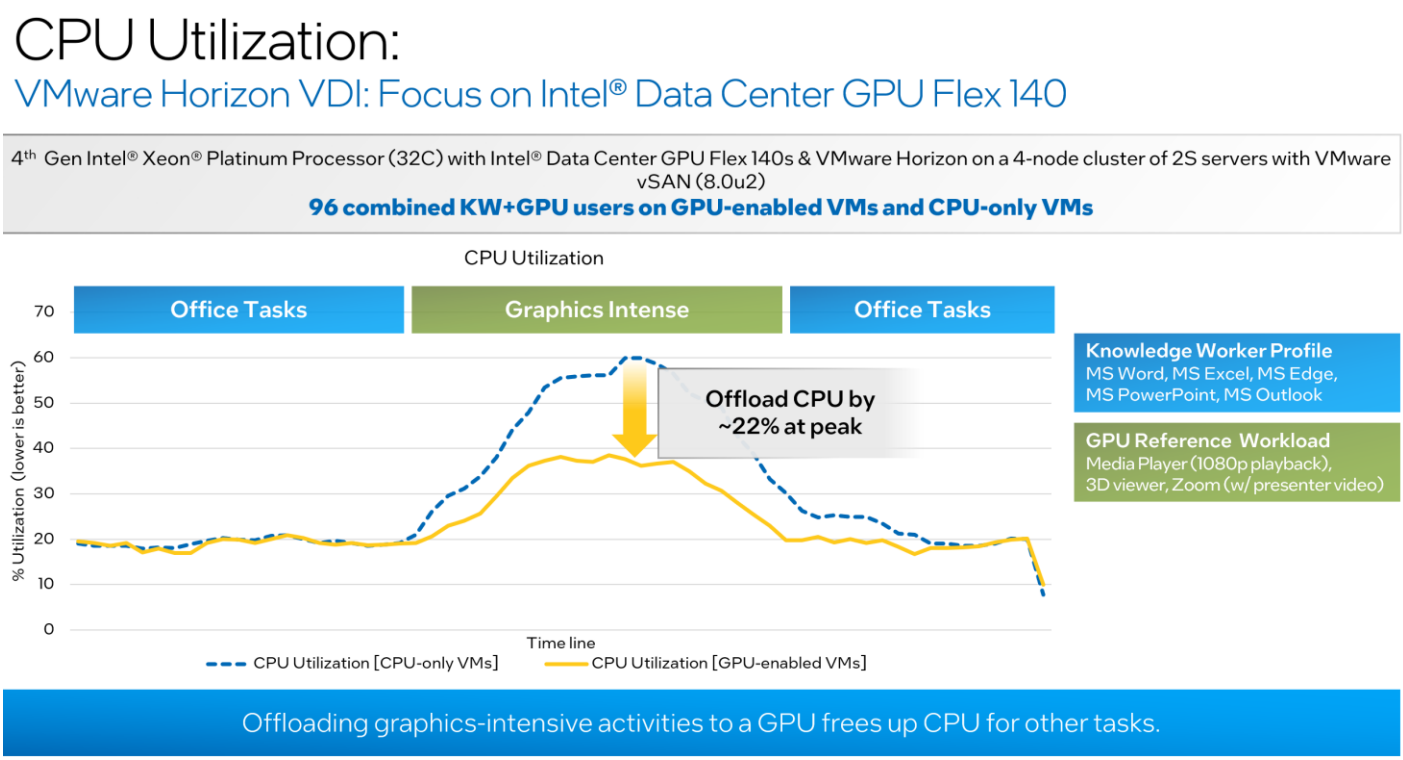

Figure 9: Offloading CPU

Conclusion

From the results above, the following observations are made.

1. Graphics intensive workloads will benefit from using a GPU as it supports higher frame rates and lower CPU utilization. Typical Microsoft Windows users will not see increased frame rates when using a GPU unless there is GPU-intensive workload running.
2. Running a graphics-intensive workloads on a GPU-enabled VM offloads the task to CPU to an extent that it can support additional users or workloads.
3. Intel® Data Center GPU Flex Series can scale across a customer's infrastructure as needed.
4. Intel® Data Center GPU Flex 140 is best for VDI environments, supporting traditional Knowledge Workers.
5. Intel® Data Center GPU Flex 170 is best for heavy graphic workloads like AI (inference), 3D CAD and rendering.

The integration of virtual GPUs in VDI sessions provides ample GPU resources, reducing CPU usage and enhancing VM density, scalability, and potential Total Cost of Ownership. SR-IOV based GPU virtualization enables flexible virtual GPU configurations without additional licensing costs, and the low pricing compared to competitors gives an added advantage to win customers. Intel's Flex series GPUs support dynamic virtual GPU scheduling [2], allowing administrators to rebalance performance on VMs. With strong encode performance per watt, Flex Series GPUs facilitate the delivery of high-performance, scalable, and future-proof VDI solutions with high user density.

[2] Intel® Data Center Graphics Driver for VMware ESXi* Section 7 of README.txt, which is part of driver package.

Resources

The VMware certified drivers can be downloaded from:

VMware Compatibility Guide - Shared Pass-Through Graphics

The setup guides can be downloaded from:

VMware Horizon* Virtual Desktop Infrastructure (VDI) Guide for Intel® Data Center GPU Flex Series - Setup Reference Guide

Citrix* VDI Guide for Intel® Data Center GPU Flex Series Setup Reference Guide

[i] Intel® Device Manager for VMware* vCenter Server User Guide