

### Overview

### Models

NVIDIA Tesla K20c Compute Processor

C2J97AA

---

### Introduction

The NVIDIA Tesla™ K20 represents the next generation from NVIDIA of supercomputing level processing capability available in a wide range of HP Z-workstations. Based on the massively parallel CUDA architecture, it features many "must have" HPC capabilities including ECC memory for uncompromised accuracy and reliability, and an L1/L2 cache design for improved bandwidth and data sharing. Based on NVIDIA's new Kepler architecture, the K20 accelerators redefine the standard for energy-efficient computing and feature innovative technologies like SMX, Hyper-Q, and Dynamic Parallelism to boost application performance by up to 10x.

---

### Features and benefits

#### 2496 CUDA Cores

Delivers up to 1.17 Tflops of double-precision peak performance in each GPU, enabling a single workstation to deliver over 2 Tflops with two K20's. Single precision peak performance is 3.52 Tflops per GPU.

#### SMX (Streaming multiprocessor)

Delivers up to 3x more performance per watt compared to the previous generation Fermi based Tesla cards.

#### Dynamic Parallelism

Enables GPU threads to automatically spawn new threads. This allows elimination of unnecessary program control transfers between CPU and GPU and enables GPU acceleration of a broader set of algorithms.

#### Hyper-Q

This feature enables multiple CPU cores to simultaneously utilize the CUDA cores on a single Kepler GPU. The result is a dramatically increased level of average GPU utilization.

#### ECC Memory\*

Meets a critical requirement for computing accuracy and reliability for workstations. Offers protection of data in memory to enhance data integrity and reliability for applications. Register files, L1/L2 caches, shared memory, and DRAM all are ECC protected.

#### 5GB of GDDR5 Memory

Maximizes performance and reduces data transfers by keeping larger data sets in local memory that is attached directly to the GPU.

#### API's

Use OpenACC or CUDA toolkits for C, C++, or Fortran to express application parallelism and take advantage of the Kepler GPU's innovative architecture.

\*Enabling ECC will cause some of the memory to be used for the ECC bits so the user available memory will decrease by ~10%.

---

### Overview

#### Compatibility

The NVIDIA Tesla K20 Computing Processor is supported on the following HP Personal Workstation:

- Z820 with 1125W power supply, Z620, Z420

The NVIDIA Tesla K20 is supported with 64-bit operating systems only.

Subject to configuration restrictions.

---

#### Service and Support

The NVIDIA Tesla K20 has a one-year limited warranty or the remainder of the warranty of the HP product in which it is installed. Technical support is available seven days a week, 24 hours a day by phone, as well as online support forums. Parts and labor are available on-site within the next business day. Telephone support is available for parts diagnosis and installation. Certain restrictions and exclusions apply.

### Technical Specifications

<b>Form Factor</b>	4.376 inches by 10.5 inches Dual Slot
<b>System Interface</b>	PCI Express Gen2 ×16
<b>Video Outputs</b>	None
<b>Memory</b>	5 GB GDDR5, 320-bit memory path
<b>Peak Memory Bandwidth</b>	208 GB/s (with ECC off)
<b>Floating Point Formats</b>	IEEE 754 single & double
<b>Supported APIs</b>	CUDA and OpenACC API support includes: CUDA C, CUDA C++, Java, Python, and Fortran
<b>Supported Operating Systems</b>	Windows 8.1 (64-bit) Genuine Windows 7 Professional (64-bit) Red Hat Enterprise Linux (RHEL) 5, 6 Desktop/Workstation (64-bit) SUSE Linux Enterprise Desktop 11 (64-bit)
	HP qualified drivers may be preloaded or available from the HP support Web site: <a href="http://welcome.hp.com/country/us/en/support.html">http://welcome.hp.com/country/us/en/support.html</a>
	Novell SUSE Linux Enterprise drivers may also be obtained from: <a href="ftp://download.nvidia.com/novell">ftp://download.nvidia.com/novell</a> or <a href="http://www.nvidia.com">http://www.nvidia.com</a>
<b>Processor Cores</b>	GK110 GPU, 706 MHz clock 2496 CUDA cores
<b>Power Consumption</b>	~225 Watts

**Note 1:** A 1125W PSU is required for any K20 configuration on the Z820

---

© Copyright 2013 Hewlett-Packard Development Company, L.P.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.