AMD ROCm™ Star – Application Developer Certificate Practice Test

Sample Questions

1. **When porting CUDA code to HIP using the hipify tools, which of the following considerations is necessary for ensuring a successful conversion?**
    a. Rely on the hipify tools to automatically convert all aspects of the application, including inline PTX assembly and build scripts.
    b. Ensure that the warp size is hardcoded to 32 to maintain consistency between NVIDIA and AMD GPUs.
    c. Manually adjust any hardcoded constructs, such as warp size, by using WarpSize, #define WARPSIZE size, or props.warpSize to get the correct value from the runtime.
    d. Trust the hipify tools to convert the application's CUDA intrinsics and unsupported functions without any manual intervention.
2. **When setting up an environment for building and running ROCm™ applications on a platform with a supported Linux-based operating system and an AMD Instinct™ GPU, what are the required steps?**
    a. Install the linux-headers and linux-modules-extra, install the GPU drivers provided by the Linux kernel, install ROCm, and add users to the video and render groups.
    b. Install the linux-headers and linux-modules-extra, install the proprietary Nvidia GPU and CUDA, install ROCm, and add users to the video and render groups.
    c. Install the linux-headers and linux-modules-extra, install the AMDGPU drivers, install ROCm, and add users to the video and render groups.
    d. Install the linux-headers and linux-modules-extra, install the AMDGPU drivers, install ROCm, and add users to the GPU and render groups.
3. **Which of the following statements is the most accurate for HIP and ROCm™ software?**
    a. HIP is a hardware platform, while ROCm is a software API for writing GPU-accelerated code.
    b. HIP is a portable C++ runtime API and kernel language, while ROCm is an open-source software platform for running on AMD GPUs
    c. ROCm is a framework for developing machine learning models, while HIP is designed specifically for low-level GPU programming.
    d. HIP is a proprietary GPU language, and ROCm is the toolchain used exclusively for AMD graphics cards.
4. **An AMD Instinct™ GPU has 104 compute units and 64 stream processors. Each compute unit may have up to 40 active wavefronts, where each wavefront is comprised of 64 lanes. The GPU is used to compute a math operation on an image with one kernel instance per pixel. Assuming kernel resource requests are minimal, at what image size does the GPU attain full occupancy?**
    a. 64x104
    b. 64x2080
    c. 64x4160
    d. 64x4096
5. **Which component of ROCm can be used for the inference of pretrained models?**
    a. TensorRT
    b. DeepSpeed
    c. MIGraphX
    d. TensorFlow

AMD ROCm™ Star – Application Developer Certificate Practice Test

6. **A tech company is scaling its cutting-edge deep learning application to run on multiple AMD GPUs spread across several nodes. To efficiently perform scatter-gather operations with the fastest communication pathways, which of the following libraries would be the ideal choice?**
    a. RCCL
    b. GPU-aware MPI
    c. NCCL
    d. AMD Infinity Fabric

7. **Which programming language is used to program GPU kernels on AMD Instinct Accelerators directly?**
    a. CUDA
    b. HIP
    c. OpenMP
    d. HLS

8. **Which of the following is the dedicated deep learning primitive library for AMD Instinct Accelerators?**
    a. cuDNN
    b. MIOpen
    c. Vitis AI
    d. oneDNN

9. **Which of the following is the correct syntax of moving tensor X to a GPU device in PyTorch?**
    a. X.to('cuda:0')
    b. X.to('rocm:0')
    c. X.to('gpu:0')
    d. X.to('dpu:0')

10. **When engaging with customers about AMD GPU and APU products for computing systems, which of the following approaches is most effective in addressing their needs and ensuring they choose the right solution?**
    a. Assess the customer's specific workload requirements, including compute, memory, and power consumption needs, and explain how AMD's GPUs (with ROCm support) and APUs can optimize their performance and efficiency.
    b. Focus solely on explaining the theoretical compute power of AMD GPUs and APUs, as this is the most important factor for customers.
    c. Recommend the highest-end AMD GPU or APU for every customer, regardless of their workload size or budget, to ensure they have the best possible performance.
    d. Avoid technical discussions and only provide marketing material, as customers typically do not need in-depth information about how AMD GPUs and APUs work.

Answer Key:

1. **C**
2. **C**
3. **B**
4. **C**
5. **C**
6. **A**
7. **B**
8. **B**
9. **A**
10. **A**