

Cluster GPU Instances Extend the Options for Running HPC Workloads in the AWS Cloud

SEATTLE -- (BUSINESS WIRE)-- Amazon Web Services LLC, an Amazon.com company, today announced Amazon Cluster GPU Instances, a new instance type designed to deliver the power of GPU processing in the cloud. GPUs are increasingly being used to accelerate the performance of many general purpose computing problems. However, for many organizations, GPU processing has been out of reach due to the unique infrastructural challenges and high cost of the technology.

Amazon Cluster GPU Instances remove this barrier by providing developers and businesses immediate access to the highly tuned compute performance of GPUs with no upfront investment or long-term commitment. To get started using Amazon EC2 GPU Instances, [visit](#): [1]

Amazon Cluster GPU Instances provide 22 GB of memory, 33.5 EC2 Compute Units, and utilize the Amazon EC2 Cluster network, which provides high throughput and low latency for HPC and data intensive applications. Each GPU instance features two NVIDIA Tesla® M2050 GPUs, delivering peak performance of more than one trillion double-precision FLOPS. Many workloads can be greatly accelerated by taking advantage of the parallel processing power of hundreds of cores in the new GPU instances.

Numerous industries – including oil and gas exploration, graphics rendering and engineering design – are using GPU processors to improve the performance of their critical applications.

Amazon Cluster GPU Instances extend the options for running HPC workloads in the AWS cloud. Cluster Compute Instances, launched earlier this year, provide the ability to create clusters of instances connected by a low latency, high throughput network.

Cluster GPU Instances give customers with HPC workloads an additional option to further customize their high performance clusters in the cloud. For those customers who have applications that can benefit from the parallel computing power of GPUs, Amazon Cluster GPU Instances can often lead to even further efficiency gains over what can be achieved with traditional processors. By leveraging both instance types, HPC customers can tailor their compute cluster to best meet the performance needs of their workloads.

For more information on HPC capabilities provided by Amazon EC2, [visit](#): [2]

Cluster GPU Instances Extend the Options for Running HPC Workloads in the Cloud

Published on Wireless Design & Development (<http://www.wirelessdesignmag.com>)

Source URL (retrieved on 12/27/2014 - 3:10am):

<http://www.wirelessdesignmag.com/news/2010/11/cluster-gpu-instances-extend-options-running-hpc-workloads-aws-cloud>

Links:

[1] <http://cts.businesswire.com/ct/CT?id=smartlink&url=http%3A%2F%2Faws.amazon.com%2Fec2%2Fhpc-applications%2F&esheet=6511154&lan=en-US&anchor=http%3A%2F%2Faws.amazon.com%2Fec2%2Fhpc-applications%2F&index=1&md5=f756a06b33c10c8d45350a8841b16c18>

[2] <http://cts.businesswire.com/ct/CT?id=smartlink&url=http%3A%2F%2Faws.amazon.com%2Fec2%2Fhpc-applications&esheet=6511154&lan=en-US&anchor=http%3A%2F%2Faws.amazon.com%2Fec2%2Fhpc-applications&index=2&md5=58e92843166f95ee7f1f9a2ce1997652>