



Modernize your Azure® Spark machine learning instances to save OpEx and increase asset utilization


We measured Apache Spark™ performance of two 64-vCPU Microsoft Azure cloud VMs:

- Standard_HB120-64rs_v3 VM, based on AMD EPYC™ 7V13 processors
- Standard_E64ds_v4 VM, based on 2nd Gen Intel® Xeon® Platinum 8272CL processors

We set up Hadoop® clusters with five VMs of each type—four worker VMs and one management VM. We then ran two HiBench workloads, Logistic Regression (LR) and Latent Dirichlet Allocation (LDA). The cluster of Standard_HB120-64rs_v3 VMs based on the AMD EPYC 7V13 processor finished both workloads in less time than the cluster of Standard_E64ds_v4 VMs based on the 2nd Gen Intel Xeon Platinum 8272CL processor. Plus, monthly estimated Pay-as-You-Go pricing was lower for the Azure Standard_HB120-64rs_v3 VMs based on the AMD EPYC 7V13 processor than for the VMs based on the Intel Xeon Platinum 8272CL processor.



HiBench Latent Dirichlet Allocation (LDA) Time in minutes to complete workload

Lower is better



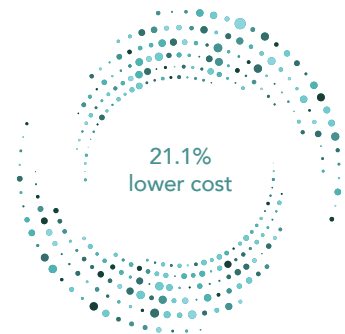
HiBench Logistic Regression (LR) Time in minutes to complete workload

Lower is better



Estimated monthly cost for HiBench test VMs in South Central Azure region (USD)*

Lower is better



*Five Azure Standard_HB120-64rs_v3 VM, based on AMD EPYC 7V13 processor, vs. five Azure Standard_E64ds_v4 VM, based on 2nd Gen Intel Xeon Platinum 8272CL processor, running 730 hours (24 hours a day for one month) in South Central US region. Source: Azure VM pricing calculator, accessed October 19, 2021, <https://azure.microsoft.com/pricing/calculator/>.

Learn more at <https://facts.pt/RzLQcyo>