

## Broad range of OCI HPC Compute Shapes and Applicable Use Cases

Instance	Specifications	Application	Billing	Learn more
BM. Optimized3.36	<ul style="list-style-type: none"> <li>• Intel Ice Lake processors</li> <li>• 100 Gbps RDMA networking</li> <li>• 3.2 TB NVMe storage</li> </ul>	Geared toward applications that are optimized for CPUs only and can be run in parallel across a large number of cores.	Per-second billing with a one-hour minimum. No charge for RDMA or NVMe SSD local storage.	<a href="#"><u>Announcing Compute instances with 3rd Gen Intel Xeon Ice Lake processors</u></a>
VM. Optimized3.Flex	<ul style="list-style-type: none"> <li>• Intel Ice Lake processors</li> <li>• Up to 50 Gbps networking</li> <li>• Up to 1 PB of block storage</li> </ul>	Geared toward applications that are optimized for CPUs only and can be run in parallel across a large number of cores.	Per-second billing with the flexibility of choosing OCPU and memory.	<a href="#"><u>Announcing Compute instances with 3rd Gen Intel Xeon Ice Lake processors</u></a>
BM. HPC2.36	<ul style="list-style-type: none"> <li>• Intel Skylake processors</li> <li>• 100 Gbps RDMA networking</li> <li>• 384 GB memory</li> <li>• 6.4 TB NVMe SSD storage</li> </ul>	Geared toward applications that are optimized for CPUs only and can be run in parallel across a large number of cores.	Per-second billing with a one-hour minimum. No charge for RDMA or NVMe SSD local storage.	<a href="#"><u>Virtual Humans on Oracle Cloud Infrastructure HPC</u></a>

Instance	Specifications	Application	Billing	Learn more
BM. GPU4.8	<ul style="list-style-type: none"> <li>8 NVIDIA A100</li> <li>AMD EPYC 7542 processors</li> <li>8X200 Gbps RDMA</li> <li>27.2 TB NVMe SSD storage</li> </ul>	Geared toward applications designed to run on GPUs. RDMA cluster networking is an added bonus for applications that can run in parallel.	Per-second billing with a one-hour minimum. No charge for RDMA or NVMe SSD local storage.	<a href="#">NVIDIA A100 Bare Metal Performance in Oracle Cloud Infrastructure</a>
BM. GPU3.8	<ul style="list-style-type: none"> <li>8 NVIDIA V100 GPUs</li> <li>Intel processors</li> <li>2X25 Gbps networking</li> <li>Block storage only</li> </ul>	Geared toward applications designed to run on GPUs.	Per-second billing with a one-hour minimum. Block storage attachment is defined by the user and billed separately.	<a href="#">GPU Benchmarking for Drug Cardiotoxicity Prediction on Oracle Cloud</a>
BM. StandardE4.128	<ul style="list-style-type: none"> <li>AMD Milan processors</li> <li>2X50 Gbps networking</li> <li>2048 GB memory</li> <li>Block storage only</li> </ul>	Geared toward high-memory workloads that do not need to scale beyond a couple of instances.	Per-second billing with the flexibility of choosing OCPU and memory.	<a href="#">Announcing Oracle Cloud Compute E4 platform on third gen AMD EPYC processors</a>

Instance	Specifications	Application	Billing	Learn more
VM. Standard.E4. Flex	<ul style="list-style-type: none"> <li>• AMD Milan processors</li> <li>• Up to 64 OCPU</li> <li>• Up to 1024 GB memory</li> <li>• Up to 40 Gbps</li> </ul>	Geared toward high-memory workloads that do not need to scale beyond a couple of instances.	Per-second billing with a one-hour minimum. Block storage attachment is defined by the user and billed separately.	<a href="#"><u>Announcing Oracle Cloud Compute E4 platform on third gen AMD EPYC processors</u></a>
BM. Standard.E3.128	<ul style="list-style-type: none"> <li>• AMD Rome processors</li> <li>• 2X50 Gbps networking</li> <li>• 2048 GB memory</li> <li>• Block storage only</li> </ul>	Geared toward high-memory workloads that do not need to scale beyond a couple of instances.	Per-second billing with a one-hour minimum. Block storage attachment is defined by the user and billed separately.	<a href="#"><u>Oracle Cloud Infrastructure Compute E3 Platform on 2nd Gen AMD EPYC Processors</u></a>
VM. Standard.E3. Flex	<ul style="list-style-type: none"> <li>• AMD Rome processors</li> <li>• Up to 64 OCPU</li> <li>• Up to 1024 GB memory</li> <li>• Up to 40 Gbps</li> </ul>	Geared toward high-memory workloads that do not need to scale beyond a couple of instances.	Per-second billing with the flexibility of choosing OCPU and memory.	<a href="#"><u>Cloud Infrastructure Compute E3 Platform on 2nd Gen AMD EPYC Processors</u></a>