

## ZenDNN 5.2.1 Support Matrix

	zentr	zentrorch
<b>Release Version</b>	5.2.1	5.2.1
<b>Release Date</b>	Apr, 2026	Apr, 2026
<b>manylinux Support (Based on Alma Linux GCC version 13.3)</b>	manylinux_2_28	manylinux_2_28
<b>glibc Version</b>	2.28	2.28
<b>Python Version</b>	<ul style="list-style-type: none"> <li>• 3.10</li> <li>• 3.11</li> <li>• 3.12</li> <li>• 3.13</li> </ul>	<ul style="list-style-type: none"> <li>• 3.10</li> <li>• 3.11</li> <li>• 3.12</li> <li>• 3.13</li> </ul>
<b>Supported Native Framework versions</b>	<ul style="list-style-type: none"> <li>• TensorFlow 2.21.0</li> <li>• TensorFlow 2.20.0*</li> </ul>	<ul style="list-style-type: none"> <li>• PyTorch 2.11.0, 2.10.0</li> <li>• vLLM 0.18.0 - 0.12.0</li> </ul>
<b>Compiler</b>	GCC v13.3	GCC v13.3
<b>Wheel File Release</b>	Yes	Yes
<b>Wheel File built with ZenDNN 5.2.1</b>	Yes	Yes
<b>Wheel File Built with FBGEMM 1.2.0</b>	Yes	Yes
<b>C++ Interface Release</b>	Yes	No
<b>Java Interface Support</b>	Yes, v1.2.0-SNAPSHOT*	No
<b>Compatible with OS</b>	<ul style="list-style-type: none"> <li>• Ubuntu 22.04, 24.04</li> <li>• RHEL 9.2, 9.5</li> </ul>	<ul style="list-style-type: none"> <li>• Ubuntu 22.04, 24.04</li> <li>• RHEL 9.2, 9.5**</li> </ul>
<b>Not Compatible with OS</b>	N/A	N/A
<b>Compatible H/W</b>	3 <sup>rd</sup> , 4 <sup>th</sup> , and 5 <sup>th</sup> Gen AMD EPYC™ processors	3 <sup>rd</sup> , 4 <sup>th</sup> , and 5 <sup>th</sup> Gen AMD EPYC™ processors
<b>Data Types</b>	<ul style="list-style-type: none"> <li>• FP32</li> <li>• BF16</li> </ul>	<ul style="list-style-type: none"> <li>• FP32, A8W8-W4, BF16, A16W4, DA8W8</li> </ul>
<p>** Use libstdc++.so.6 library which supports GLIBCXX_3.4.30</p> <p>* TF Java will be supported with TF 2.20.0 CPP package</p> <p>A16W4 - A16 refers to activation as BF16 dtype; we support both asymmetric and symmetric Quantization for weights</p> <p>DA8W8 - DA8 is dynamic INT8 quantization of activation; we support symmetric quantization for activation and weights.</p> <p>A8W8-W4 - Activation asymmetric and weights are symmetrically quantized for linear ops; Embedding Bag weights are asymmetrically quantized</p>		
<p><b>System's GCC version should be greater than or equal to 9.4 for torch.compiler flow</b></p>		