

CPUs are Key to Enterprise AI

Before choosing a hardware option, it's crucial to think about the business problem you're trying to solve. Deploying AI requires your best software development process, right from the start.

Enterprise respondents want to avoid porting AI applications from the training environment to a different chip architecture for inference; using a uniform fleet of machines allows a uniform software stack. Models losing performance in the transfer into production is very common indeed; it's therefore important to think about hardware early on in the process, and to see the choice of hardware as part of your software and model development process.



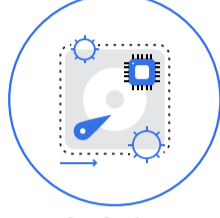
23%

On identical stacks



56%

On the same hardware architecture (e.g. GPU) but not the same chip (e.g. T4)



20%

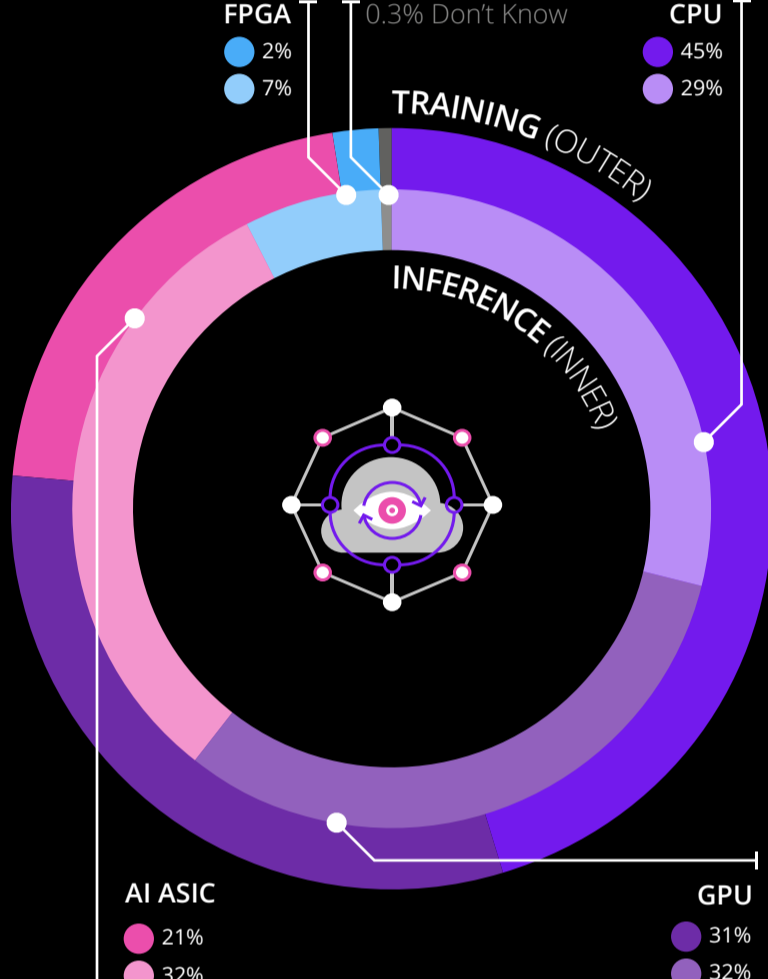
On different hardware architectures

1% Don't know

45%

of respondents said that CPUs were their primary AI model training option.

Primary hardware choices:

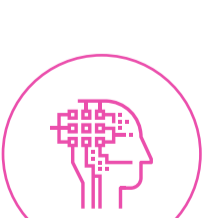
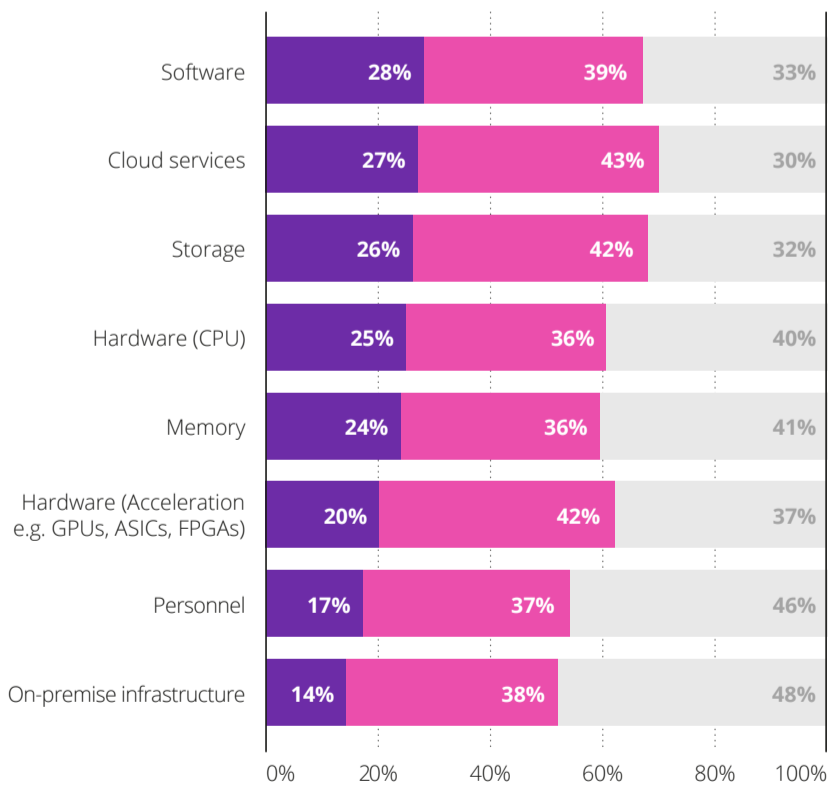


CPU also ranks as the favoured option for inference on-device, and almost alongside GPU for inference on the cloud.

It's therefore important to think about hardware early on in the process, and to see the choice of hardware as part of your software and model development process. Despite everything, our respondents are still more likely to increase their investment by over 20% in CPUs than they are in any other logic category.



How will spending on the following items change in the next two years?



Real enterprises deploying AI therefore need a wide selection of high performance CPUs to support their developers and data scientists' productivity.

Click here to view the full whitepaper

<https://www.intel.com/content/www/us/en/content-details/768032/content-details.html>

