

# Baidu PaddlePaddle on 3rd Gen Intel® Xeon® Scalable processors



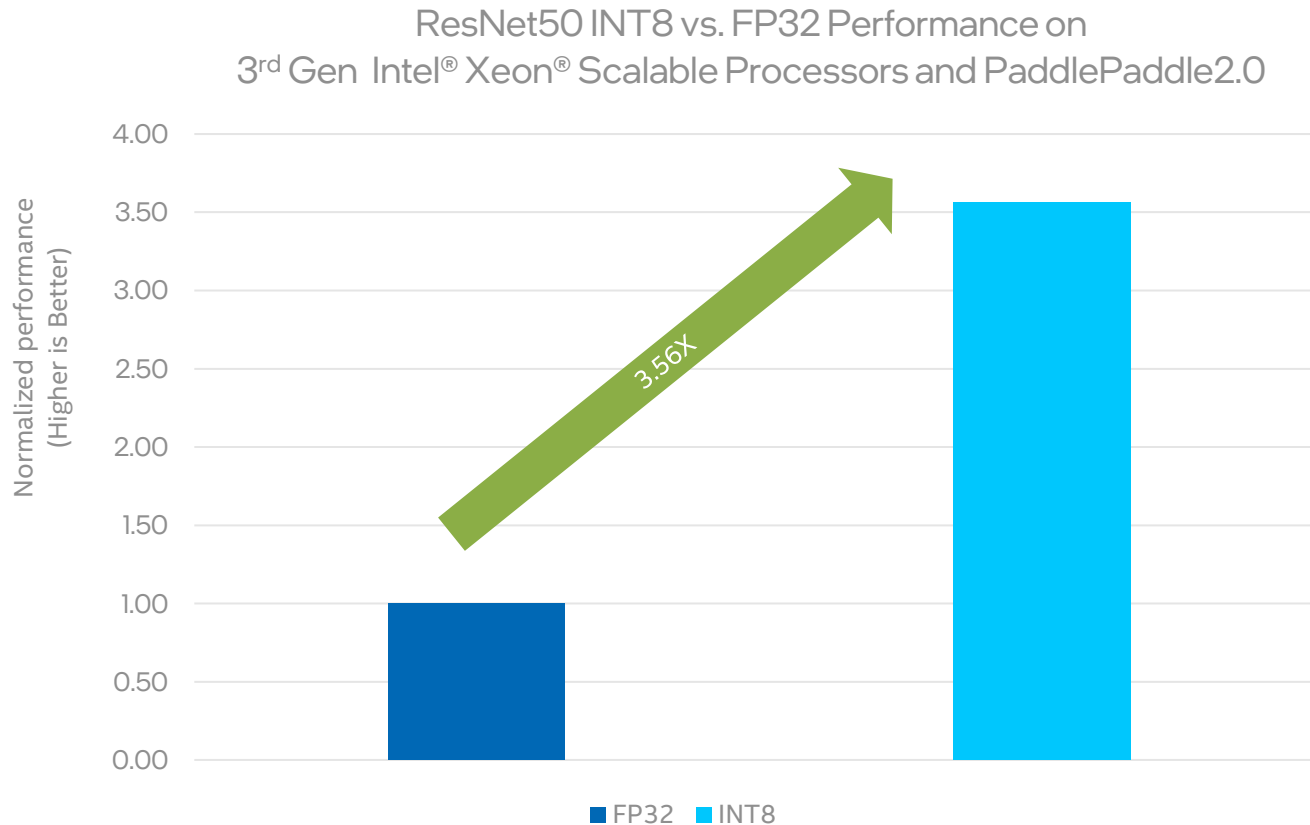
# Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at [www.Intel.com/PerformanceIndex](http://www.Intel.com/PerformanceIndex)

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

# 3<sup>rd</sup> Gen Intel® Xeon® Scalable Processors Baidu PaddlePaddle ResNet50



## Application

- ResNet50 is widely used for performance benchmarks in Baidu
- Many of Baidu optical character recognition (OCR) models in business line are using ResNet50 as the backbone. Its performance gain improves many related models.

## Customer Impact

- Unleashes the latest Intel® Xeon® Scalable processor features for Baidu models.
- Reduces inference latency on CPU and provides better experiences to Baidu's customers.<sup>1</sup>

## Performance Drivers

- **Up to 3.56X inference performance speedup** utilizing Baidu PaddleSlim tool with Intel® Deep Learning Boost technology on the latest Intel Xeon Scalable processors compared to a FP32 solution.<sup>1</sup>

1 – Performance results are based on testing done by Intel March 19, 2021. For complete testing configuration details, see Configurations section.

# Configurations

## Resnet50 Model on PaddlePaddle 2.0 Throughput Performance on 3rd Generation Intel® Xeon® Processor Scalable Family

Test configurations: Test by Intel as of 03/19/2021. 2-node, 2x Intel(R) Xeon(R) Platinum 8350C CPU @ 2.60GHz, 32 cores, HT On, Turbo Off. Total Memory 256 GB (4 slots/ 64GB/ 2933 MHz), BIOS: WLYDCRB1.SYS.0020.P92.2103170501 (ucode: 0x8d000270), Ubuntu 18.04.3 LTS, 4.15.0-55-generic, gcc 9.3.0 compiler, ResNet50 Model, Deep Learning Framework: PaddlePaddle2.0, download at [https://paddle-inference-lib.bj.bcebos.com/2.0.0-cpu-avx-mkl/paddle\\_inference.tgz](https://paddle-inference-lib.bj.bcebos.com/2.0.0-cpu-avx-mkl/paddle_inference.tgz) BS=1, ImageNet Val, 1 instances, Datatype: FP32/INT8

For more information visit:

<https://intel.com/content/www/us/en/artificial-intelligence/posts/deep-learning-with-baidu-paddlepaddle.html>

The Intel logo is centered on a solid blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter 'i'. To the right of the word "intel" is a registered trademark symbol (®).

intel®