intel®

Intel® Select Solution for BigDL on Apache Spark*

Accelerate and simplify deep learning development and deployment on an optimized, verified infrastructure based on Apache Spark.

(intel®) select solution

In the past few years, organizations have seen a convergence of massive amounts of data with the compute power and large-capacity storage needed to process it all. The right infrastructure can provide modern businesses with new ways of harnessing data for innovative apps and services built on artificial intelligence (AI). The opportunities are nearly infinite and stretch across almost every field—from financial services to manufacturing to healthcare and beyond.

But organizations with on-premises infrastructures or using hybrid cloud models face several challenges on the road to AI. They need to research, select, deploy, and optimize infrastructure that can provide efficient resource utilization while scaling on demand to meet changing business requirements. Beyond scalability, organizations seek easier ways to implement AI initiatives. Many businesses lack sufficient in-house expertise and infrastructure to get started with AI, particularly for deep learning (DL). The road to deploying DL in production environments is time-intensive and complex. Managing the data for AI initiatives can also be a challenge: organizations struggle to extract value from their "data swamps," and it can be complex and resource-intensive to move data from on premises to the cloud for analytics.

The Intel® Select Solution for BigDL on Apache Spark* can help businesses overcome these key challenges to achieve their AI initiatives faster and more easily. The pre-tested and tuned solution eliminates the need for organizations to research and manually optimize infrastructure to efficiently pursue their AI initiatives. The solution reduces the need for specialized in-house expertise to deploy and manage AI infrastructure. And it can help IT organizations improve infrastructure utilization, while ensuring scalability to meet the growing needs of their companies.

BigDL

Apache Spark helps solve the IT challenges of DL, data, and specialized expertise by providing for standardized big-data storage and compute, with scalability, by enabling the addition of hundreds of nodes without degrading performance and without changing the fundamental architecture.

BigDL—a distributed DL library that augments the storage and compute capabilities of Apache Spark—provides efficient, scalable, and optimized DL development. BigDL enables the development of new DL models for training and serving on the same big data cluster. It also supports models from other frameworks, including TensorFlow*, Keras*, and others, so you can import other trained models into the BigDL framework or use BigDL trained models in other frameworks. BigDL is supported by Analytics Zoo, which provides a unified AI platform and pipeline with built-in reference use cases to further simplify your AI-solutions development.

BigDL is optimized for Intel®-based platforms with software libraries like Intel® Math Kernel Library (Intel® MKL) and Intel® Math Kernel Library for Deep Learning Networks (Intel® MKL-DNN) to increase computational performance. Other supporting software includes the Intel® Distribution for Python*, which accelerates popular machine learning libraries such as NumPy*, SciPy*, and scikit-learn* with integrated Intel® Performance Libraries such as Intel MKL and Intel® Data Analytics Acceleration Library (Intel® DAAL). On the hardware side, the Intel Select Solution for BigDL on Apache Spark uses Intel® Xeon® Scalable processors for high performance and Intel® Solid State Drives (SSDs) for better performance and improved reliability compared to traditional hard-disk drives (HDDs).

The Intel Select Solution for BigDL on Apache Spark

The Intel Select Solution for BigDL on Apache Spark helps optimize price/performance while significantly reducing infrastructure evaluation time. The Intel Select Solution for BigDL on Apache Spark combine Intel Xeon Scalable processors, Intel SSDs, and Intel® Ethernet Network Adapters to empower enterprises to quickly harness a reliable, comprehensive solution that delivers:

- **The ability to prepare** your machine learning (ML)/DL infrastructure investments for the future with scalable storage and compute

- **Excellent total cost of ownership (TCO)** with multi-purpose hardware that your IT organization is used to managing in a verified, tested solution that simplifies deployment

- **Accelerated time to market** with a turnkey solution that includes a rich development toolset and that is optimized for crucial software libraries

- **The ability to run analytics** on data where it is stored

---

Intel® Xeon® Scalable Processors

Intel Xeon Scalable processors:

- Offer high scalability for enterprise data centers

- Deliver performance gains for virtualized infrastructure compared to previous-generation processors

- Achieve exceptional resource utilization and agility

- Enable improved data and workload integrity and regulatory compliance for data center solutions

The Intel® Select Solution for BigDL on Apache Spark* features Intel Xeon Gold processors.

Improved Patient Care with AI

Built on BigDL, Apache Spark*, and Intel® Xeon® Scalable Processors

The Intel® Select Solution for BigDL on Apache Spark began as a **partnership with the University of California, San Francisco (UCSF)**. Intel and UCSF deployed BigDL on Apache Spark, running on a data-analytics cluster based on Intel Xeon Scalable processors. The solution was used to develop and train a DL model that could examine sophisticated 3D medical images to improve medical diagnoses for patients with osteoarthritis.

Intel expanded on the success of that project to create an optimized Intel Select Solution, based on next-generation technology.

Hardware Selections

The Intel Select Solution for BigDL on Apache Spark combines a platform based on the Intel Xeon Gold processor or a higher number Intel Xeon Scalable processor, Intel® 3D NAND SSDs, and the Intel® Ethernet 700 Series, so that your business can quickly deploy reliable, BigDL-based DL solutions on a performance-optimized infrastructure.

Intel® Xeon® Scalable Processors

Intel Xeon Gold processors provide Intel Select Solutions for BigDL on Apache Spark with an excellent price/performance ratio. Specifically, Intel selected the Intel Xeon Gold 6148 processor as the minimum processor required to power Intel Select Solutions for BigDL on Apache Spark in order to optimize cost and performance for representative DL workloads while taking into consideration the number of cores and the speed of the processor. Solutions incorporating the latest Intel Xeon Gold 6248 processors deliver the same performance or incremental performance gains as compared to similarly configured solutions based on previous-generation Intel Xeon Scalable processors. Intel Xeon Scalable processors can provide up to 1.4x faster AI training and 5.4x faster AI inferencing compared to previous generations of Intel Xeon processors.[1]

Intel® SSD Data Center Family

Storage latency can be a major bottleneck for DL performance. For this reason, the Intel Select Solution for BigDL on Apache Spark uses the Intel SSD DC S4500 and Intel SSD DC P4510. Based on Intel 3D NAND technology, these enterprise data center SSDs provide a 3.2x lower annualized failure rate (AFR) than HDDs.[2]

Intel® Ethernet Connections and Intel® Ethernet Adapters

The Intel Ethernet 700 Series accelerates the performance of Intel Select Solutions for BigDL on Apache Spark. These solutions feature the Intel Ethernet 700 Series with 10 gigabit Ethernet (GbE) for validated performance ready to meet high quality thresholds for data resiliency and service reliability for most media types and port speeds, and they're backed by extensive testing, validation, and worldwide product support.[3,4,5,6]

What Are Intel® Select Solutions?

Intel Select Solutions are verified hardware and software stacks that are optimized for specific software workloads across compute, storage, and network. The solutions are developed from deep Intel experience with industry solution providers, in addition to extensive collaboration with the world's leading data center and service providers.

To qualify as an Intel Select Solution, solution providers must:

1. Follow the software and hardware stack requirements outlined by Intel (see **Appendix A**)

2. Replicate or exceed Intel's reference benchmark-performance threshold

3. Publish a detailed implementation guide to facilitate customer deployment

Solution providers can develop their own optimizations to add further value to their solutions.

Verified Performance through Benchmark Testing

All Intel Select Solutions are verified through benchmark testing to meet a pre-specified minimum capability level of workload-optimized performance. Intel chose the GoogLeNet* topology running an ImageNet* dataset in order to meet stringent performance standards representative of DL workloads.

GoogLeNet (Inception v1*) won the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) in 2014. ImageNet is a large dataset of more than one million images belonging to multiple classes. This combination of DL topology and image set enables Intel to characterize system-resource utilization for long-running DL training workloads in order to find the optimal configuration for the Intel Select Solution for BigDL on Apache Spark.[7,8,9]

Base Configuration

The Intel Select Solution for BigDL on Apache Spark is available in the configuration shown in Appendix A. The "Base" configuration specifies the minimum required performance capability for the Intel Select Solution for BigDL on Apache Spark.

Technology Selections for the Intel Select Solution for BigDL on Apache Spark

In addition to the Intel hardware foundation used for the Intel Select Solution for BigDL on Apache Spark, Intel technologies integrated in Intel Xeon Scalable processors deliver further performance and reliability gains:

- **Intel® Volume Management Device (Intel® VMD):** Enables hot-swap replacement of NVM Express* (NVMe*) SSDs from the Peripheral Component Interconnect Express*

(PCIe*) bus without shutting down the system, while standardized LED management helps provide much faster identification of SSD status. This standardization brings enterprise reliability, availability, and serviceability (RAS) features to NVMe SSDs, enabling you to deploy next-generation storage with confidence. IT professionals can now service these drives online without an outage, which minimizes interruptions and improves uptime and serviceability. The unique value of Intel VMD is that Intel is sharing this technology across the ecosystem for broad enablement.

- **Intel® Virtual RAID on CPU (Intel® VROC) Technology:** Delivers excellent performance, low power/TCO; supports full-featured RAID levels 0, 1, 5, and 10; and is a host bus adapter (HBA)-less RAID solution.

- **Intel® QuickAssist Technology (Intel® QAT):** An offload engine to accelerate some critical workloads such as bulk cryptography, public key exchange, and data compression on Intel architecture–based platforms. Intel QAT on Intel Xeon Scalable processors offers outstanding capabilities: up to 100 gigabits per second (Gbps) encryption, 100 Gbps compression, and 100,000 decryption operations per second using a 2,048-bit RSA key.

- **Internet Wide Area Remote Protocol (iWARP) Remote Direct Memory Access (RDMA):** A host-offload, host-bypass technology that enables a secure direct memory-to-memory data communication between two applications across a network. iWARP RDMA can make use of current Ethernet infrastructure without lossless network support. It also provides flow control and congestion management, and it is highly scalable.

- **Intel® Platform Trust Technology (Intel® PTT):** Root of trust with full Trusted Platform Module (TPM) 1.2 functionality integrated into platform firmware. The new Intel PTT feature is available as an option versus a discrete chip to simplify integration and activation.

- **Intel® Boot Guard (Security):** Hardware-based boot integrity protection that prevents unauthorized software and malware takeover of boot blocks critical to a system's function, thus providing an added level of platform security based on hardware.

Deploy an Enterprise-Ready DL Solution with the Intel Select Solution for BigDL on Apache Spark

Intel Select Solutions are a fast path to data center transformation with workload-optimized configurations verified for Intel Xeon Scalable processors. When organizations choose Intel Select Solutions for BigDL on Apache Spark, they get pre-tuned and tested configurations that are workload-optimized and proven to scale with Intel Xeon Scalable processors so that they can deploy DL solutions quickly and efficiently with less tuning.

Visit intel.com/selectsolutions to learn more, and ask your infrastructure vendor for Intel Select Solutions.

Learn More

Intel Select Solutions: **intel.com/selectsolutions**

Intel Xeon Scalable processors: **intel.com/xeonscalable**

Intel SSD Data Center Family: **intel.com/content/www/us/en/products/ memory-storage/solid-state-drives/data-center-ssds.html**

Intel Ethernet 700 Series: **intel.com/ethernet**

Intel Select Solutions are supported by Intel® Builders: **http://builders.intel.com**. Follow us on Twitter: **#IntelBuilders**

BigDL: **https://software.intel.com/en-us/articles/bigdl-distributed-deep-learning-on-apache-spark**

Appendix A: Base Configuration for the Intel Select Solution for BigDL on Apache Spark

To refer to a solution as an Intel Select Solution, a server vendor or data center solution provider must meet or exceed the defined minimum configuration ingredients and reference minimum benchmark-performance thresholds listed below.

INGREDIENT	INTEL® SELECT SOLUTION FOR BIGDL ON APACHE SPARK* **BASE CONFIGURATION**
MINIMUM OF ONE MASTER NODE	
PROCESSOR	Intel® Xeon® Gold 6148 processor (2.40 GHz, 20 cores, 40 threads), Intel Xeon Gold 6248 processor (2.5 GHz, 20 cores), or a higher number Intel Xeon Scalable processor
MEMORY	192 GB or higher (12 x 16 GB DDR4-2666)
BOOT DRIVE**	1 x 240 GB or larger Intel® SSD DC S4500/S4510 (M.2 or 2.5-inch) or higher
STORAGE HBA CONTROLLER PER NODE**	Not applicable (N/A)
DATA TIER	4 x 960 GB or larger Intel SSD DC S4500/S4510 or higher
DATA NETWORK	10 Gb Intel® Ethernet Converged Network Adapter X722
MANAGEMENT NETWORK PER NODE	Integrated 1 GbE port 0/RMM port
MINIMUM OF FOUR WORKER NODES	
PROCESSOR	Intel Xeon Gold 6148 processor (2.40 GHz, 20 cores, 40 threads), Intel Xeon Gold 6248 processor (2.5 GHz, 20 cores), or a higher number Intel Xeon Scalable processor
MEMORY	384 GB or higher (12 x 32 GB DDR4-2666)
BOOT DRIVE**	1 x 240 GB or larger Intel SSD DC S4500/S4510 (M.2 or 2.5-inch) or higher
STORAGE HBA CONTROLLER PER NODE**	N/A
DATA TIER	4 x 3.8 TB or larger Intel SSD DC S4500/S4510 or higher
DATA NETWORK	10 Gb Intel Ethernet Converged Network Adapter X722
MANAGEMENT NETWORK PER NODE	Integrated 1 GbE port 0/RMM port
NETWORK SWITCHES	
TOP OF THE RACK (ToR) SWITCH	10 Gbps 48x port switch
MANAGEMENT SWITCH	1 Gbps 48x port switch

SOFTWARE	
LINUX* OS	CentOS* Linux release 7.5.1804 or Red Hat* Enterprise Linux (RHEL*) 7
APACHE SPARK	2.2.0
APACHE HADOOP*	2.7.3
JAVA DEVELOPMENT KIT* (JDK*)	Oracle* JDK 1.8.0 update 181 or later
BIGDL	0.6
ANALYTICS ZOO	0.2
INTEL® DISTRIBUTION FOR PYTHON*	2.7
INTEL® MATH KERNEL LIBRARY (INTEL® MKL)	Intel MKL version 2018 Update 3
APPLIES TO ALL NODES	
TRUSTED PLATFORM MODULE (TPM)	TPM 1.2 discrete or firmware TPM (Intel® Platform Trust Technology [Intel® PTT])
FIRMWARE AND SOFTWARE OPTIMIZATIONS	Intel® Volume Management Device (Intel® VMD) enabled**
	Intel® Boot Guard enabled**
	Intel® Hyper-Threading Technology (Intel® HT Technology) disabled
	Intel® Turbo Boost Technology enabled
	P-states enabled**
	C-states enabled**
	Power-management settings set to performance**
	Workload configuration set to balanced**
	Memory Latency Checker (MLC) streamer enabled**
	MLC spatial prefetch enabled**
	Data Cache Unit (DCU) data prefetch enabled**
	DCU instruction prefetch enabled**
	Last-level cache (LLC) prefetch disabled**
	Uncore frequency scaling enabled**

MINIMUM PERFORMANCE STANDARDS

Verified to meet or exceed the following minimum performance capabilities:

IMAGENET* TRAINING THROUGHPUT	375 images per second with Top-5 Accuracy of 85% or higher[10]

**Recommended, not required

[1] Intel® Optimization for Caffe ResNet-50* performance does not necessarily represent other framework performance. Based on Intel internal testing: 1x baseline (July 11, 2017), 1.4x training (August 2, 2018), and 5.4x inferencing (July 26, 2018) performance improvement based on Intel® Optimization for Caffe ResNet-50 inference throughput performance on Intel® Xeon® Scalable processors. Performance results are based on testing as of July 11, 2017 (1x), August 2, 2018 (1.4x), and July 26, 2018 (5.4x) and may not reflect all publicly available security updates.

July 11, 2017 details: (Previous-generation processor) 2S Intel Xeon processor E5-2699 v4 at 2.20 GHz (22 cores), Intel® Hyper-Threading Technology (Intel® HT Technology) enabled, Intel® Turbo Boost Technology disabled, scaling governor set to "performance" via acpi-cpufreq driver, 256 GB DDR4-2133 ECC RAM. CentOS* Linux* release 7.3.1611 (core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3500 (480 GB, 2.5-in Serial ATA [SATA], 6 gigabits per second [Gbps], 20 nanometer [nm], Memory Latency Checker [MLC]). (Current-generation processor) 2S Intel Xeon Platinum 8180 processor at 2.50 GHz (28 cores), Intel HT Technology disabled, Intel Turbo Boost Technology disabled, scaling governor set to "performance" via intel_pstate driver, 384 GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel SSD DC S3700 (800 GB, 2.5-in SATA, 6 Gbps, 25 nm, MLC). Performance measured with environment variables set to: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU frequency set with cpupowerfrequency-set -d 2.5G -u 3.8G -g performance. Caffe framework details: revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffetime --forward_only" command, training measured with "caffetime" command. For "ConvNet" topologies, dummy dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (GoogLeNet*, AlexNet*, and ResNet-50), https://github.com/intel/caffe/tree/master/models/default_vgg_19 (VGG-19), and https://github.com/soumith/convnet-benchmarks/tree/master/caffe/imagenet_winners (ConvNet benchmarks; files were updated to use newer Caffe prototxtformat).

August 2, 2018 details: 2S Intel Xeon Platinum 8180 processor at 2.50 GHz (28 cores), Intel HT Technology enabled, Intel Turbo Boost Technology enabled, 376 GB DDR4-2666 ECC RAM, resnet_50_v2 topology. Engine: Intel® Math Kernel Library for Deep Neural Networks (Intel® MKL-DNN) version 464c268e544bae26f9b85a2acb9122c766a4c396.

July 26, 2018 details: 2S Intel Xeon Platinum 8180 processor at 2.50GHz (28 cores), Intel HT Technology enabled, Intel Turbo Boost Technology enabled, 192 GB DDR4-2666 ECC RAM, best (resnet--50, gnet_v3 --224, ssd--224) topology. Caffe framework details: inference measured with "caffe time --forward_only -phase TEST" command, training measured with "caffe train" command. Environment variables: KMP_AFFINITY='granularity=fine, compact,1,0', OMP_NUM_THREADS=28, CPU frequency set with cpupowerfrequency-set -d 2.5G -u 3.8G -g performance. Caffe run with "numactl-l".

[2] Based on initial product AFR of 0.66 percent vs. industry AFR average (2.11%). Source: Backblaze. "Hard Drive Stats for Q1 2017." May 2017. backblaze.com/blog/hard-drive-failure-rates-q1-2017/.

[3] The Intel® Ethernet 700 Series includes extensively tested network adapters, accessories (optics and cables), hardware, and software, in addition to broad operating system support. A full list of the product portfolio's solutions is available at intel.com/ethernet. Hardware and software is thoroughly validated across Intel® Xeon® Scalable processors and the networking ecosystem. The products are optimized for Intel® architecture and a broad operating system ecosystem: Windows*, Linux* kernel, FreeBSD*, Red Hat* Enterprise Linux (RHEL*), SUSE*, Ubuntu*, Oracle Solaris*, and VMware ESXi*.

[4] The Intel® Ethernet 700 Series is backed with global support infrastructure for customers pre- and post-sales.

[5] Supported connections and media types for the Intel® Ethernet 700 Series are: direct-attach copper and fiber SR/LR (QSFP+, SFP+, SFP28, XLPPI/CR4, 25G-CA/25G-SR/25G-LR), twisted-pair copper (1000BASE-T/10GBASE-T), backplane (XLAUI/XAUI/SFI/KR/KR4/KX/SGMII). Note that Intel is the only vendor offering the QSFP+ media type.

[6] The Intel® Ethernet 700 Series supported speeds include 10 GbE, 25 GbE, 40 GbE, and 100 GbE.

[7] Olga Russakovsky**, Jia Deng**, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg and Li Fei-Fei. (** = equal contribution) ImageNet Large Scale Visual Recognition Challenge. IJCV, 2015. image-net.org/challenges/LSVRC/.

[8] Szegedy, Christian, Liu, Wei, Jia, Yangqing, et al. "Going deeper with convolutions." arXiv.org. September 2014. https://arxiv.org/pdf/1409.4842.pdf.

[9] ImageNet. About ImageNet page. http://image-net.org/about-overview.

[10] Performance results are based on testing as of October 10, 2018, and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. **Configuration:** *One master node:* 2 x Intel® Xeon® Gold 6148 processor, Intel® Server Board S2600WFT, total memory: 192 GB, 12 slots/16 GB/2,666 megatransfers per second (MT/s) DDR4 RDIMM, Intel® Hyper-Threading Technology (Intel® HT Technology) disabled, Intel® Turbo Boost Technology enabled; storage (boot): 1 x 240 GB Intel® SSD DC S4500; storage:4 x 960 GB Intel SSD DC S4500; network devices: 1 x 10 Gb Intel® Ethernet Network Adapter X722, network speed: 10 GbE, ucode: 0x043, OS/software: CentOS Linux release 7.5.1804*; *four worker nodes:* 2 x Intel Xeon Gold 6148 processor, Intel Server Board S2600WFT, total memory: 384 GB, 12 slots/32 GB/2,666 megatransfers per second (MT/s) DDR4 RDIMM, Intel HT Technology disabled, Intel Turbo Boost Technology enabled; storage (boot): 1 x 240 GB Intel SSD DC S4500; storage: 4 x 3.8 TB Intel SSD DC S4500; network devices: 1 x 10 Gb Intel Ethernet Network Adapter X722, network speed: 10 GbE, ucode: 0x043, OS/software: CentOS Linux release 7.5.1804.