用友云
yonyou cloud

迈向云端 创变未来

2017中国企业互联网大会

用友云
yonyou cloud

助力数据分析 加速企业创新

李宏 英特尔软件与服务事业部 资深经理

标题1：28号字体，左上角

正文： 18号字体，1.5倍行距，段落首行缩进2个字符。 18号字体，1.5倍行距，段落首行缩进2个字符。 18号字体，1.5倍行距，段落首行缩进2个字符。 18号字体，1.5倍行距，段落首行缩进2个字符。 18号字体，1.5倍行距，段落首行缩进2个字符。

18号字体，1.5倍行距，段落首行缩进2个字符。 18号字体，1.5倍行距，段落首行缩进2个字符。 18号字体，1.5倍行距，段落首行缩进2个字符。

The Next Big Wave

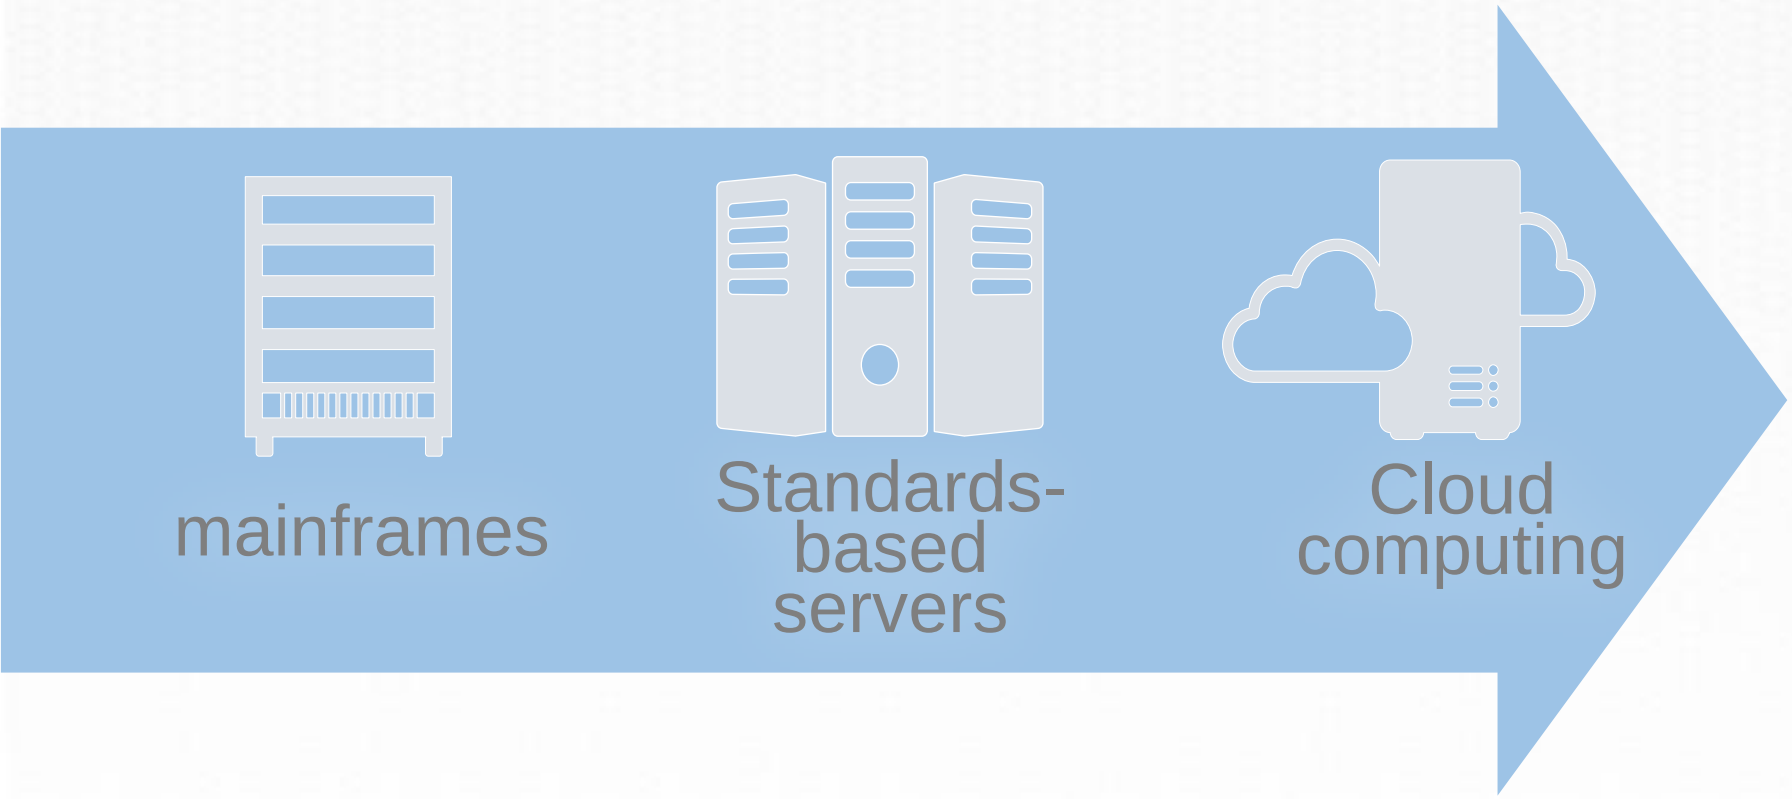

- Data deluge
- COMPUTE breakthrough
- Innovation surge

Artificial intelligence

AI Compute Cycles will grow 12X by 2020

Source: Intel forecast

Data Analytics is Evolving…

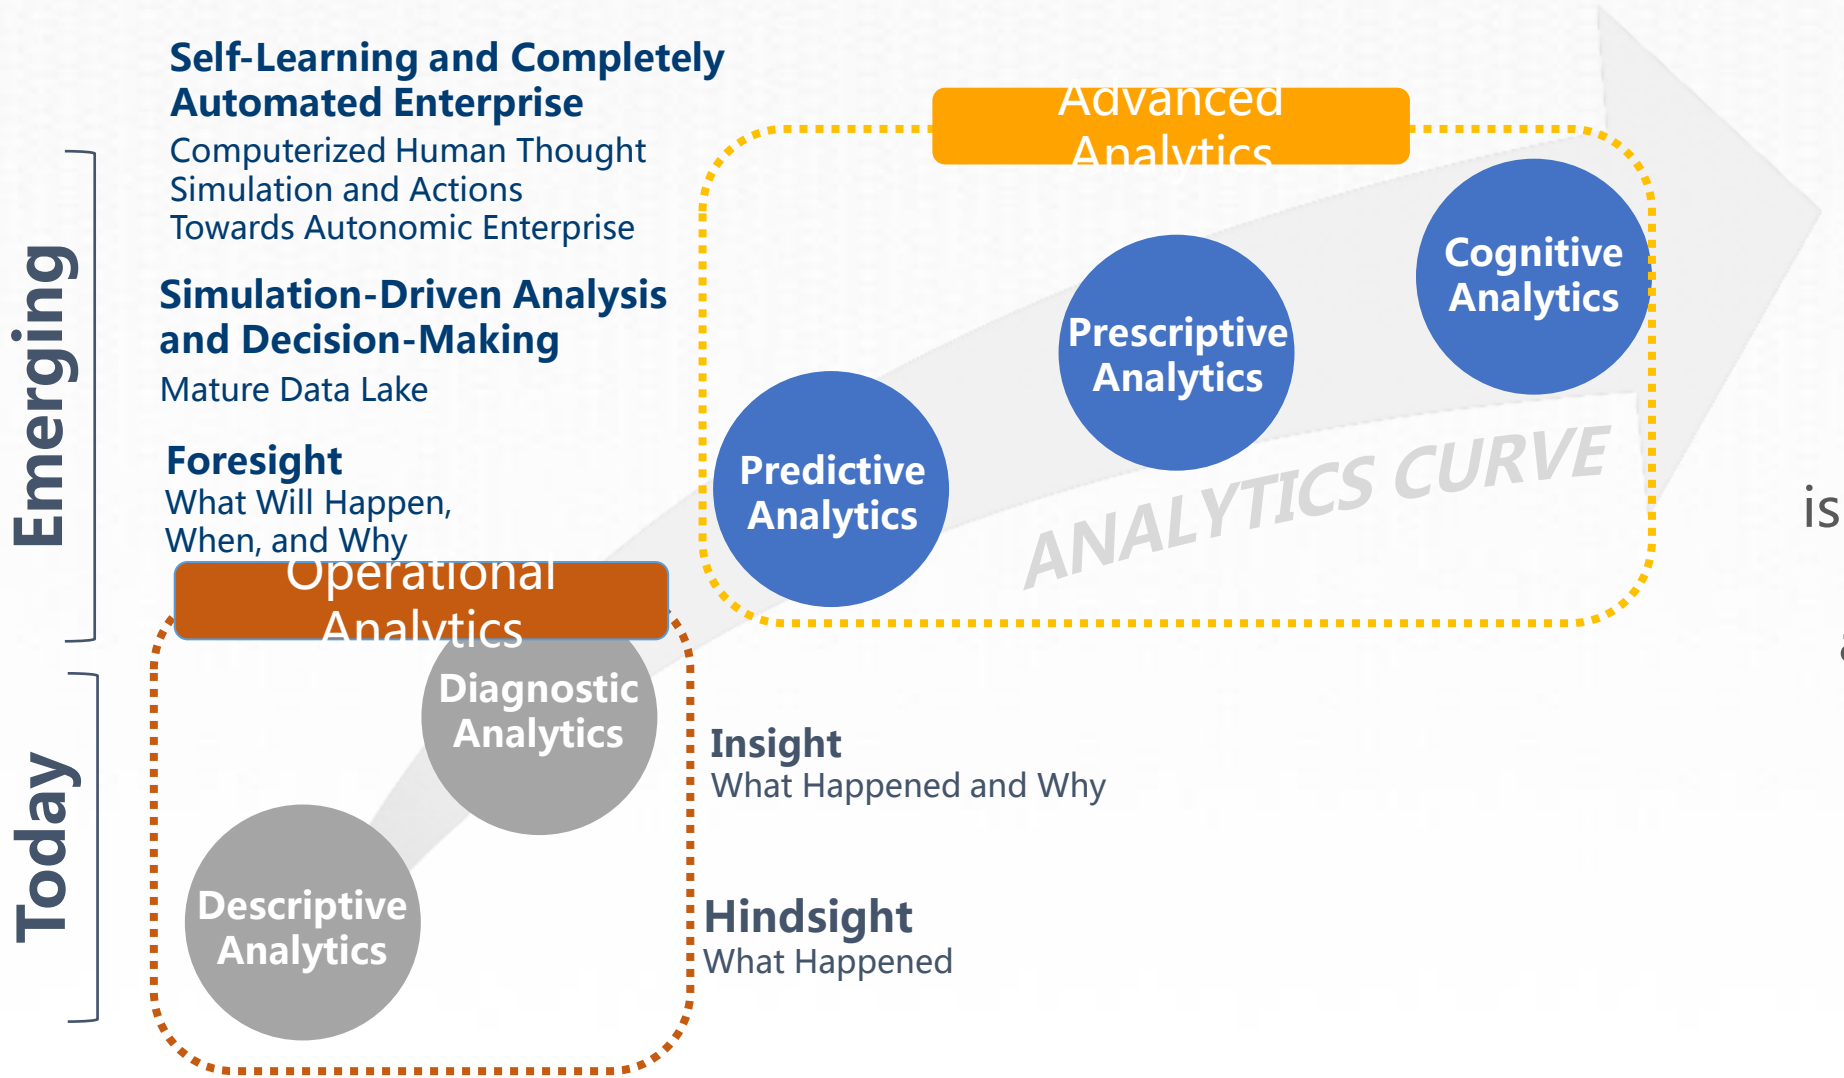

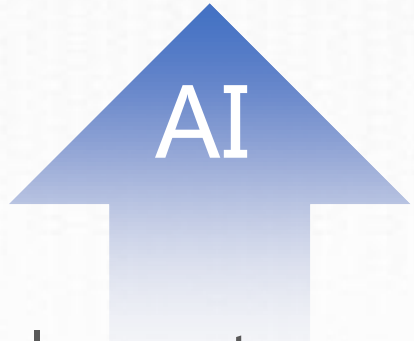

is a large category all on its own, and a vital tool for reaching higher maturity & scale data analytics

Data Analytics is Transformative

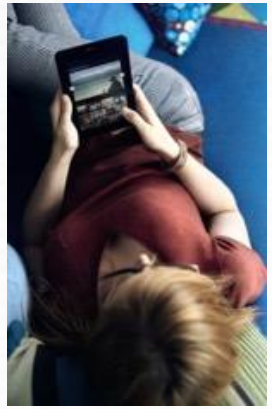

Consumer

- Smart Assistants
- Chatbots
- Search
- Personalization
- Augmented Reality
- Robots

Health

- Enhanced Diagnostics
- Drug Discovery
- Patient Care
- Research
- Sensory Aids

Finance

- Algorithmic Trading
- Fraud Detection
- Research
- Personal Finance
- Risk Mitigation

Retail

- Support
- Experience
- Marketing
- Merchandising
- Loyalty
- Supply Chain
- Security

Government

- Defense
- Data Insights
- Safety & Security
- Resident Engagement
- Smarter Cities

Energy

- Oil & Gas Exploration
- Smart Grid
- Operational Improvement
- Conservation

Transport

- Autonomous Cars
- Automated Trucking
- Aerospace
- Shipping
- Search & Rescue

Industrial

- Factory Automation
- Predictive Maintenance
- Precision Agriculture
- Field Automation

Other

- Advertising
- Education
- Gaming
- Professional & IT Services
- Telco/Media
- Sports

Source: Intel forecast

Data Analytics in Practice

Healthcare: Tumor detection

Normal

Tumor

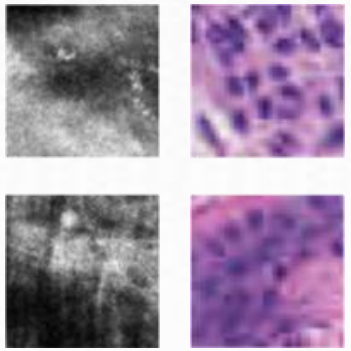

Industry: Agricultural Robotics

Plant

Weed

Energy: Oil & Gas

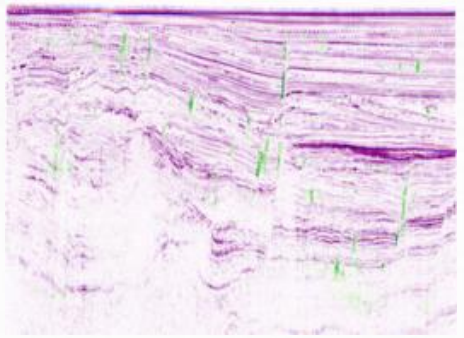

Automotive: Speech interfaces

Finance: Document Classification

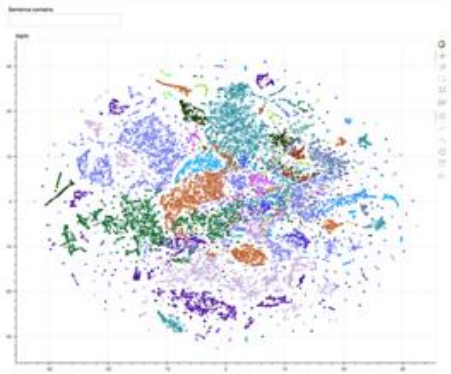

Genomics: Sequence analysis

End-to-End Data Analytics

Datacenter

Many-to-many hyperscale for stream and massive batch data processing

Ethernet & Wireless

Gateway

1-to-many with majority streaming data from devices

Wireless and non-IP wired protocols

- ✓ *Secure*
- ✓ *High throughput*
- ✓ *Real-time*

Edge

1-to-1 devices with lower power and often UX requirements

- Intel® Xeon® Processors
- Intel® Xeon Phi™ Processors*
- Intel® Core™ & Atom™ Processors
- CPU+ Intel® Processor Graphics
- Intel® FPGA
- Crest Family (Nervana ASIC)*
- Movidius VPU
- Intel® GNA (IP)*

*Future

Datacenter Product-line for Data Analytics

All purpose

Intel® Xeon® Processor Family

most agile AI Platform

Scalable performance for widest variety of AI & other datacenter workloads – including deep learning training & inference

Highly-parallel

Intel® Xeon Phi™ Processor (Knights Mill†)

Faster DL Training

Scalable performance optimized for even faster deep learning training and select highly-parallel datacenter workloads*

Flexible acceleration

Intel® FPGA

Enhanced DL Inference

Scalable acceleration for deep learning inference in real-time with higher efficiency, and wide range of workloads & configurations

Deep Learning

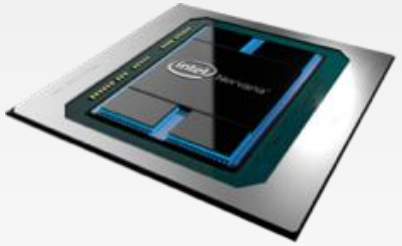

Crest Family†

Deep learning by design

Scalable acceleration with best performance for intensive deep learning training & inference

IA Optimized Frameworks

Intel®'s reference deep learning framework committed to best performance on all hardware
intelnervana.com/neon

✓Optimized for Intel Architecture

See **Roadmap** for availability

Intel® Data Analytics Acceleration Library

High Performance ML and Data Analytics library

Building blocks for all data analytics stages, including data preparation, data mining & machine learning

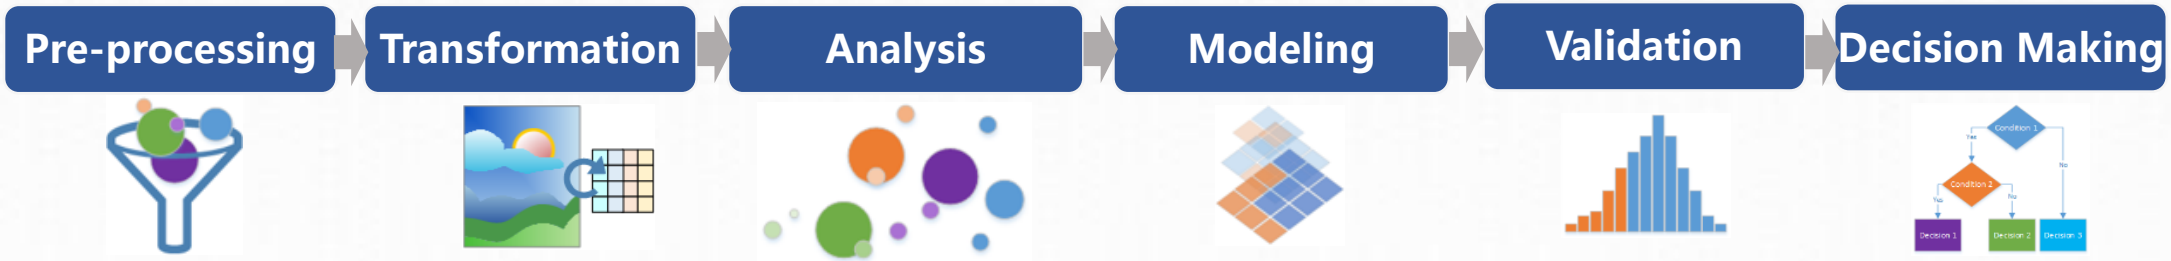

Open Source • Apache 2.0 License

Common Python, Java and C++ APIs across all Intel hardware

Optimized for large data sets including streaming and distributed processing

Flexible interfaces to leading big data platforms including Spark and range of data formats (CSV, SQL, etc.)

PCA Optimization by Intel DAAL

DAAL - IA optimized analytics building blocks

- Key ingredient for proprietary and open source data analytics platforms and applications
- Fundamental building blocks for all data analysis stages
- Delivers forward-scaling performance and parallelism on IA
- Built upon Intel® Math Kernel Library and Intel® Integrated Performance Primitives

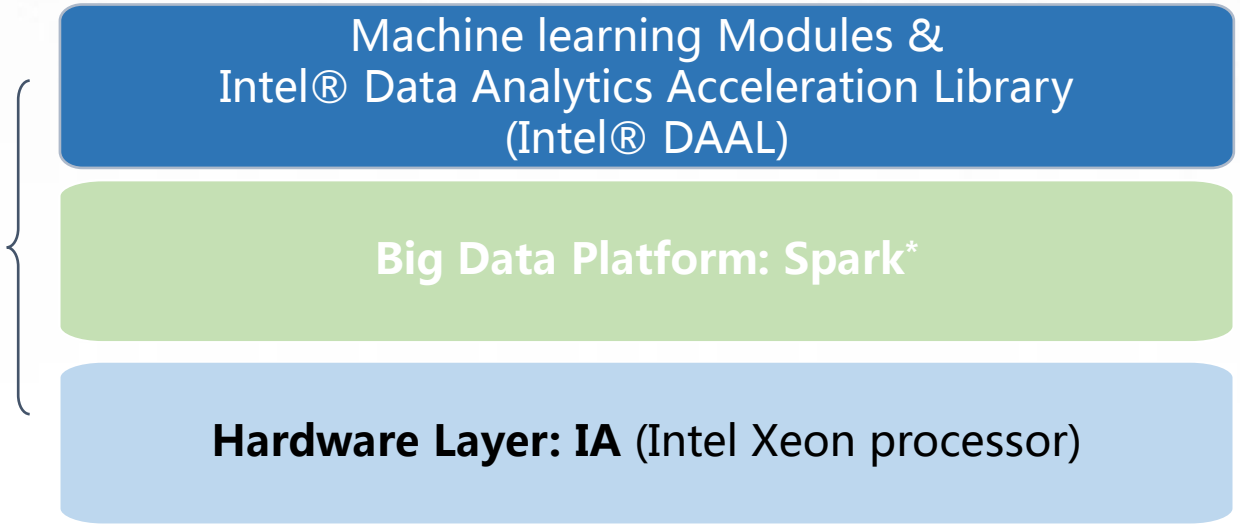

Intel DAAL dramatically speeds up PCA computing speed to reduce data dimension from 4096 to 128 on big data platform

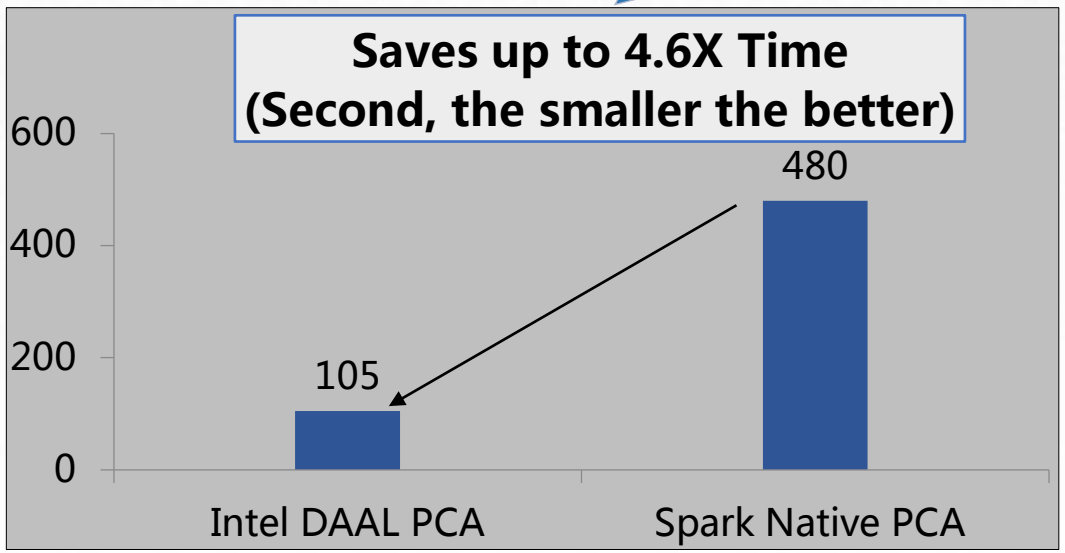

The test data is based on Intel® Xeon® E5620 @ 2.4GHz, 2 sockets, 64 GB RAM, 2 node

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance .

Intel® MKL-DNN

Math Kernel Library for Deep Neural Networks

For developers of deep learning frameworks featuring optimized performance on Intel hardware

Distribution Details

- Open Source
- Apache 2.0 License
- Common DNN APIs across all Intel hardware.
- Rapid release cycles, iterated with the DL community, to best support industry framework integration.
- Highly vectorized & threaded for maximal performance, based on the popular Intel® MKL library.

github.com/01org/mkl-dnn

Examples:

Illegal Video Detection Optimization

- Illegal video detection becomes a common requirement for 3rd party video cloud customers
- Originally, open source BVLC Caffe plus OpenBlas were used as CNN framework, but the performance was poor
- Using Caffe* + Intel MKL achieved up to 30X performance improvement

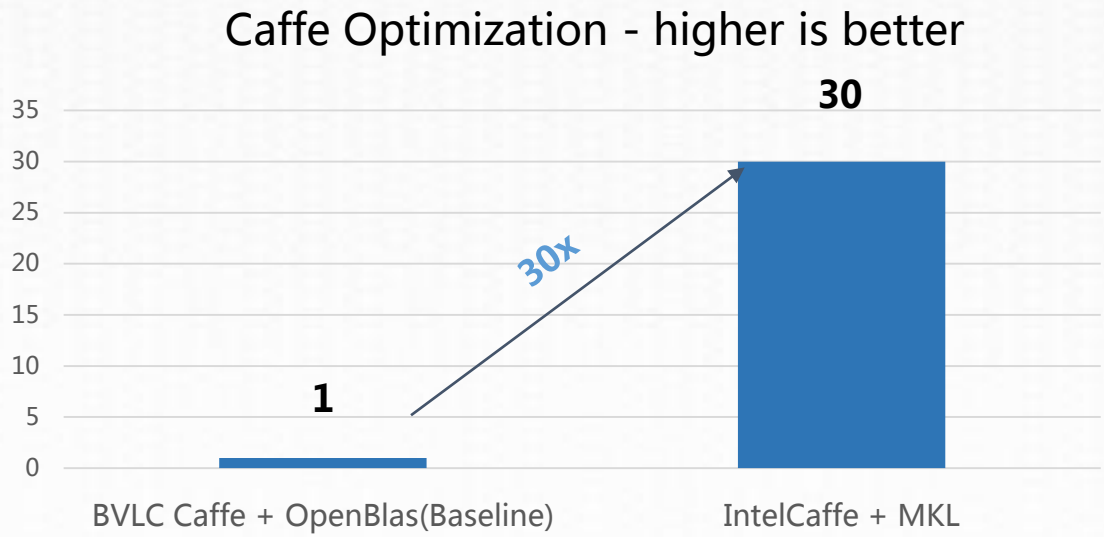

* The test data is based on Intel Xeon E5 2680 V3 processor

Yonyou BQ 性能优化

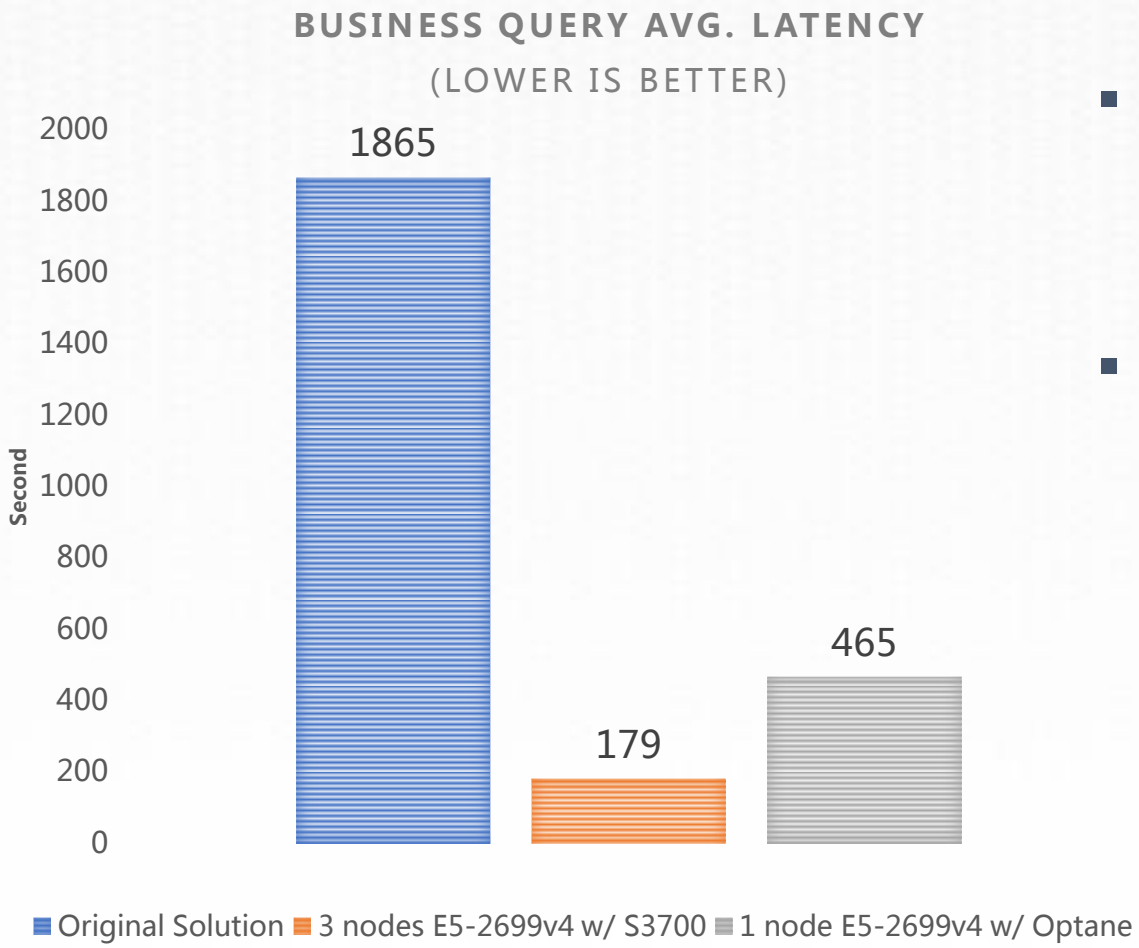

使用2个月真实数据 123GB

- **原方案**
  - 复杂业务查询、统计的平均时间超过30分钟
- **新方案**
  - 测试了3节点集群和单节点使用Optane的情况：
  - 3节点集群使用Spark内存计算，原始数据存放在HDD中，Intel SSD S3700 作为Shuffle存放中间结果，复杂查询时间缩短到179秒，速度提高了10倍
  - 单节点使用Optane存放数据和Shuffle存放中间结果，采用Spark内存计算，复杂查询时间缩短到465秒，速度提高了4倍

用友云 yonyou cloud

企业服务都在这