

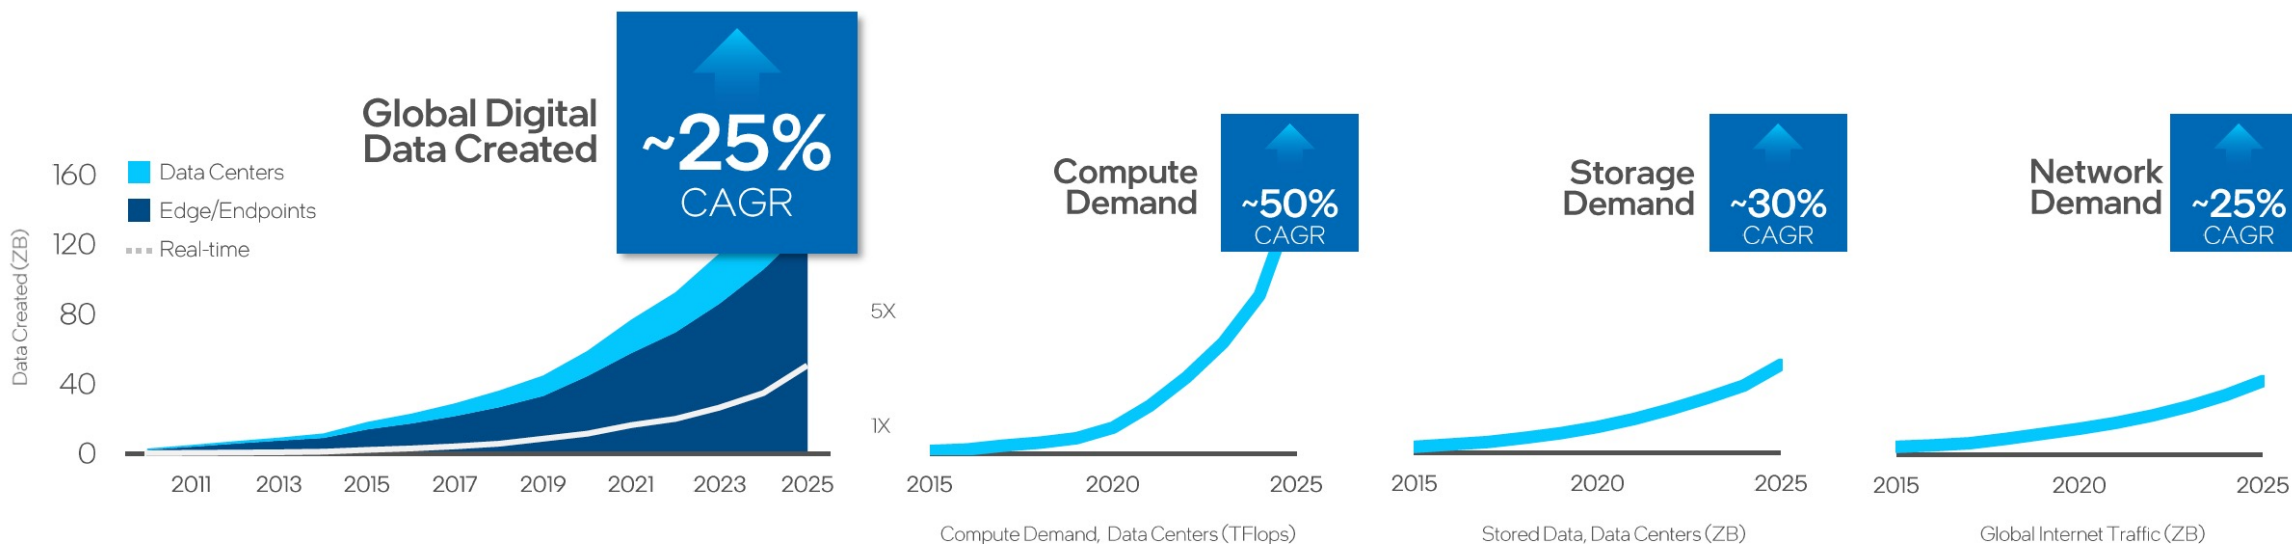
클라우드를 위한 인텔 제온 서버

인텔코리아 이인구 전무






데이터 증가로 인해 컴퓨팅, 스토리지, 네트워크 수요 증가



클라우드로의 전환과, 4차 산업혁명.

Sources: Global Internet Traffic, Cisco VNI Forecasts (2016, 2017, 2018); Stored Data, IDC Data Age 2025 (2018); Compute demand, Intel analysis

INTEL TECHNOLOGIES FOR CLOUD COMPUTING



Move **Faster**

Store **More**

Process **Everything**

Intel® Ethernet

Intel® Silicon Photonics

Intel® Tofino

intel.

OPTANE™

PERSISTENT
MEMORY

intel.

OPTANE™

SSD

intel.

XEON™

PLATINUM

intel.

ATOM

intel.

AGILEX™


intel.

eASIC™



Software & System Level **Optimized**

INTEL TECHNOLOGIES FOR CLOUD COMPUTING



Move **Faster**

Intel® Ethernet
Intel® Silicon Photonics
Intel® Tofino

Store **More**

intel.
OPTANE™
PERSISTENT MEMORY

intel.
OPTANE™
SSD

Process **Everything**

intel.
XEON™
PLATINUM

intel.
ATOM

intel.
AGILEX™

intel.
eASIC™

intel.
MOVIDIUS™



X^e

Software & System Level **Optimized**

INTEL CLOUD ADVANTAGE

Performance Gains with Acceleration

Relative performance of 2-socket 2nd Gen Intel Xeon Scalable (24C, 2.9GHz) processor with feature on/off. (Higher is Better)



AI-ready: Intel® DL Boost, Intel® Arria FPGAs, optimized frameworks



In-Memory Computing: Intel® Optane™ DC Persistent Memory



HPC: Intel® AVX-512, fabric, software



Networking: Intel® Ethernet, Intel® QuickAssist, Intel® DDIO



Storage: Intel® QuickAssist, Intel® Optane™ Solid State Drives, Intel® VMD




Security: Intel® QuickAssist, Intel® SGX, Intel® TXT

Cloud Tools & Resource Optimization

CLOUD OPTIMIZATION

Intel Migration Advisor
by  Cloud General
Inventory & Business Case

Intel Cloud Optimizer
by  Densify
Performance, Cost, SLA

Intel Workload Optimizer
by  GRANULATE
Performance & Latency

WORKLOAD TUNING

Intel Cloud DevOps Toolkit
Characterize workloads to analyze and debug performance
bottlenecks

Intel Optimized Cloud Stacks
Optimized configurations delivered as Infrastructure as code or
container

Intel Optimized Containers
Pre-defined optimized containers for specific applications and/or use
cases

APPLICATION DEVELOPMENT

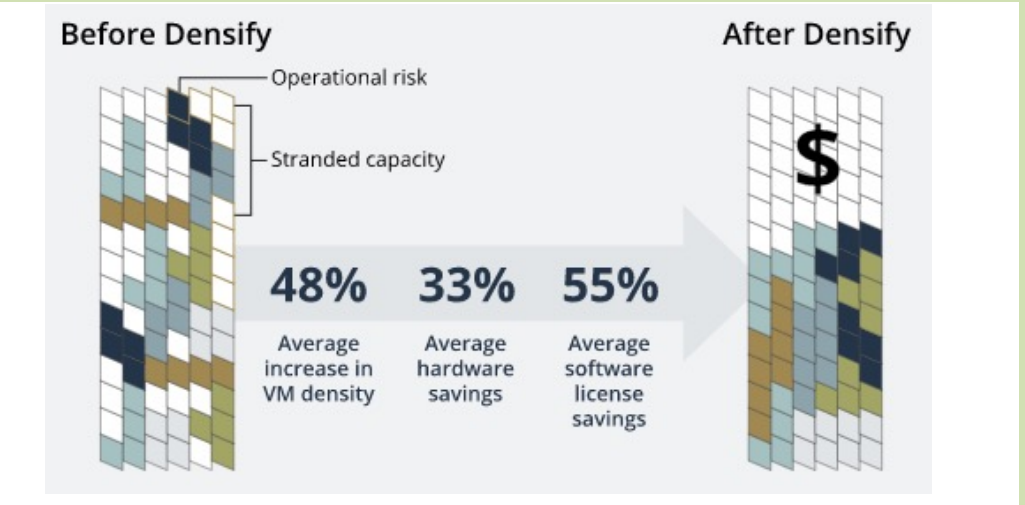
Intel oneAPI Toolkits
Develop performant data-centric applications across Intel® CPUs,
GPUs, and FPGAs with this foundational set

Intel Distribution of OpenVINO™ Toolkit
Develop high performance AI inference applications
from edge to cloud

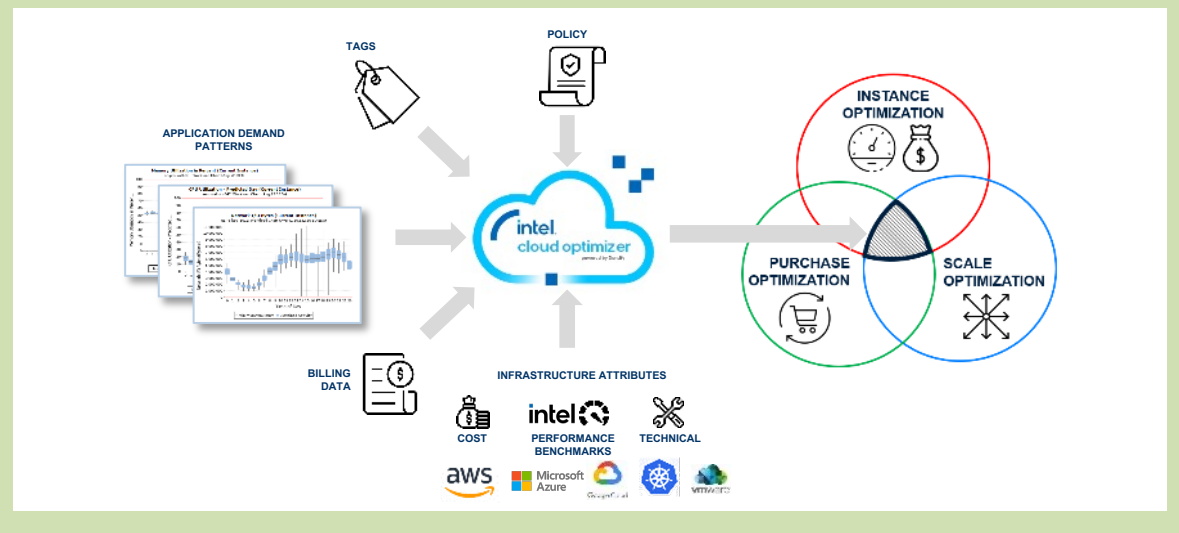
Intel Cloud Optimizer

by  Densify

1 Intel Cloud Optimizer by Densify Increases Cloud Utilization While Minimizing Performance Risk



2 Following The Activation, Machine Learning Analyzes Workload Patterns



Application-Driven

Optimization Directives on Steps to Reduce Cost while Preserving Cloud Performance

Optimization

Optimization of the Configuration of Auto Scaling Groups and Automatic Setting of Cloud Computing Instance Types


Visibility

Impact Analysis and Recommendation Reports on Optimization Recommendations, Predicted Utilization, Effort, and Cost


Intel Workload Optimizer


by GRANULATE

1 Intel Workload Optimizer by Granulate Learns and Analyses the Application Dataflow and Processing Stages Patterns



2 Following The Activation, The Agent Optimizes OS-Resource Scheduling and Prioritization



Application-Driven

Tailored Resource Management Based On The Specific Application Running

Dynamic

Real-Time Scheduling and Prioritization Decisions
Continuous Adapted To The Load Applied

Automatic

Automatic and Autonomous Decisions Without Configuration or Maintenance

Intel Cloud DevOps Toolkit

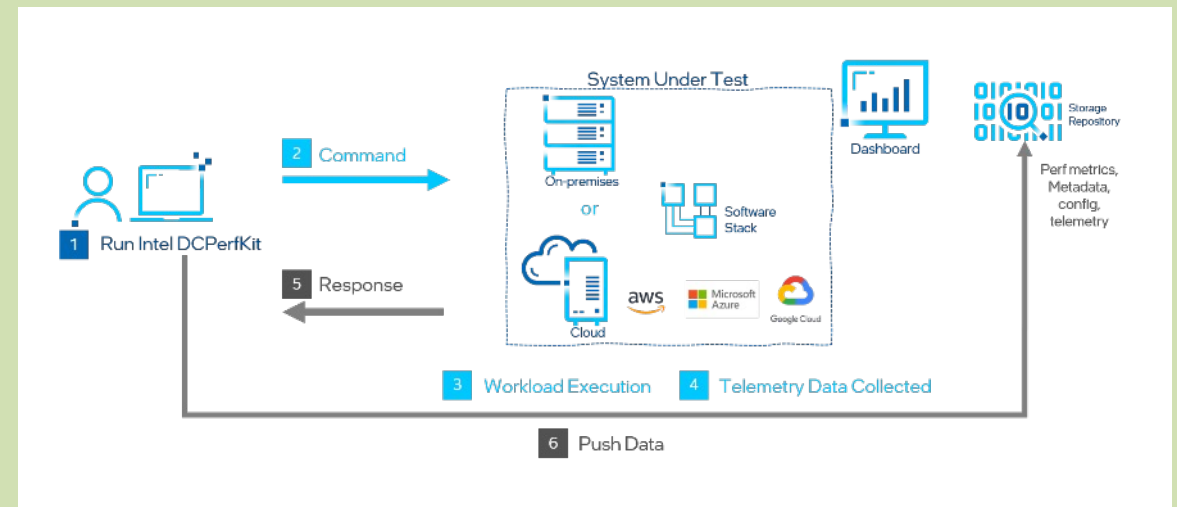
1

Suite of compilers, developer tools, code libraries, reference models, and guides to optimize performance, portability, and cost of cloud instances

- [Intel® VTune™ Platform Profiler](#)
- [Intel® Data Center Performance Kit \(Intel® DC Perf Kit\) – \(NDA\)](#)
- [Intel® Telemetry Collector - \(NDA\)](#)
- [Intel® Perfspect – \(NDA\)](#)
- [Intel® SKU Emulator – \(NDA\)](#)
- [Intel® System Health Inspector \(svr_info\) – \(NDA\)](#)
- [Intel® Workload Similarity Analyzer – \(NDA\)](#)
- [Intel® Memory Latency Checker \(MLC\) – \(NDA\)](#)
- [Process Counter Monitor \(PCM\) – \(NDA\)](#)
- [Intel® Workload Optimization Service – \(NDA\)](#)
- [cAdvisor with Intel® PMU extension](#)
- [Workload Co-location Agent](#)
- [3rd gen Intel® Xeon® Processor Performance Tuning Guides](#)
- [gProfiler](#)

2

Suite of compilers, developer tools, code libraries, reference models, and guides to optimize performance, portability, and cost of cloud instances



Flexible

Supports on-premises and public cloud performance benchmarks and system configuration recommendations

Proven

Decades of co-engineering across on-premises, cloud, and end-user communities with proven results

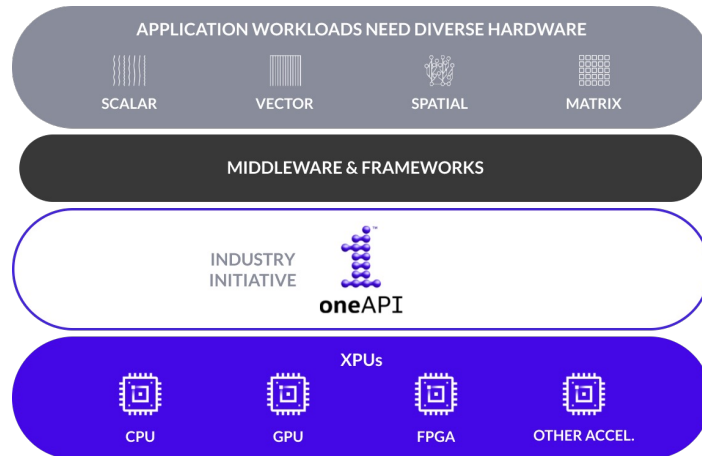
Optimized

Optimization of recommendation for system and application configuration settings. Automated benchmarks to verify changes

Intel oneAPI Toolkit(s)

1

A single, unified programming model that aims to simplify development across multiple architectures, such as CPUs, GPUs, FPGAs accelerators.



2

Extract the most application performance on multiple types of Intel® architecture by using advanced, cross-architecture software development tools from Intel.

- Intel® oneAPI Base Kit
- Intel® oneAPI AI Analytics Toolkit
- Intel® oneAPI HPC Toolkit
- Intel® oneAPI Rendering Toolkit
- Intel® oneAPI IOT Toolkit
- Intel® System Bring-up Toolkit
- Intel® Distribution of OpenVINO Toolkit (powered by oneAPI)

Open

Removes code barriers with one programming model for all architectures. Apply your skills to the next innovation, and not rewriting software.

Efficient

Achieve fast and efficient development with a complete set of cross-architecture libraries and tools that use familiar languages.

Optimized

Expose and exploit the latest hardware's cutting-edge features to unleash application performance.

Open Source Collaboration



#1 corporate contributor

Intel features frequently enabled in Linux kernel at platform time-to-market



Up to **15.21X AES** and **9.16X SHA** crypto/hash performance gain on 3rd Gen Intel Xeon Scalable processor



Intel optimized distribution.
Up to **3.7X** AI inference performance gain on 2nd Gen Intel Xeon Scalable processor¹



Intel optimized distribution
Up to **1.93X** AI training performance gains on 3rd Gen Intel Xeon Scalable processor¹



...and many more

Workload Benchmarking



LB/Webtier
TLS connections

Up to
4.2X
Higher


NGINX web server connections/s with Intel Crypto Acceleration vs prior gen



Web Services

Up to
1.48X
Higher

secure requests to content management systems (Wordpress with HTTPS) vs prior gen



Artificial Intelligence

Up to
11X
Higher

batch AI inference perf on ResNet50 using TensorFlow and Intel DL Boost vs prior gen (FP 32)



Analytics

Up to
1.65X
Higher

performance on CloudXPRT cloud data analytics usage vs prior gen



Search

Up to
5.1X
Higher

Splunk search performance scaling containers vs prior gen



Databases

Up to
1.64X
Higher

MySQL transactions per minute vs prior gen

3세대 인텔® 제온® 스케일러블 프로세서

Performance made flexible

AI 및 보안 솔루션을 내장한
유일한 x86 데이터 센터 프로세서



1S-2S 시스템용

첨단 보안 솔루션



인텔 Software
Guard Extensions



인텔 암호화
가속 기능




인텔 Total Memory
Encryption



인텔 Platform
Firmware
Resilience

확장성, 유연성, 고객맞춤



인텔 딥 러닝 부스트



인텔 스피드
셀렉트 기술



인텔
AVX-512



최적화된
소프트웨어

차세대 인텔® 제온® 스케일러블 플랫폼

최대 **6TB** 최대 **8CH** 최대 **2.6X** 최대 **64**

시스템 메모리
용량
(소켓 별)
DRAM + PMEM

DDR4-3200
2 DPC
(소켓 별)

2세대 제온
스케일러블 프로세서
대비
메모리 용량 증가

라인
PCI Express 4
(소켓 별)

획기적인 데이터 성능

OPTANE



인텔 옵테인
퍼시스턴트 메모리
200 시리즈

OPTANE



인텔 옵테인
SSD P5800X
시리즈



인텔 SSD
D 시리즈

더 빠르고, 유연하며, 데이터 확장 가능



인텔 이더넷
800 시리즈
네트워크 어댑터



인텔
앨리스
FPGA
솔루션



인텔
셀렉트
솔루션

3세대 인텔® 제온® 스케일러블 프로세서

Performance made flexible

AI 및 보안 솔루션을 내장한
유일한 x86 데이터 센터 프로세서



1S-2S 시스템용

첨단 보안 솔루션



인텔 Software Guard Extensions



인텔 암호화 가속 기능




인텔 Total Memory Encryption



인텔 Platform Firmware Resilience

확장성, 유연성, 고객맞춤



인텔 딥 러닝 부스트



인텔 스피드 선택 기술



인텔 AVX-512



최적화된 소프트웨어

차세대 인텔® 제온® 스케일러블 플랫폼

최대 6TB

시스템 메모리 용량 (소켓 별)
DRAM + PMEM

최대 8CH

DDR4-3200 2 DPC (소켓 별)

최대 2.6X

2세대 제온 스케일러블 프로세서 대비 메모리 용량 증가

최대 64

레인 PCI Express 4 (소켓 별)

획기적인 데이터 성능

OPTANE



인텔 옵테인 퍼시스턴트 메모리 200 시리즈

OPTANE



인텔 옵테인 SSD P5800X 시리즈



인텔 SSD D 시리즈

더 빠르고, 유연하며, 데이터 확장 가능



인텔 이더넷 800 시리즈 네트워크 어댑터



인텔 애질렉스 FPGA 솔루션



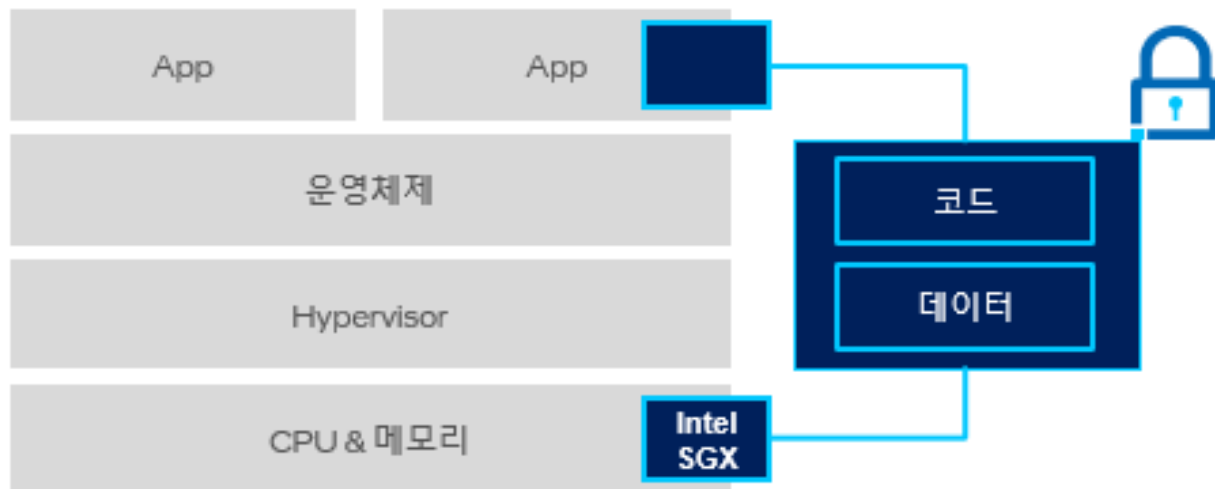
인텔 셀렉트 솔루션

인텔 SGX (Software Guard Extensions)



성능 저하 없이 민감한 데이터 영역의
개인 정보 보호 가능

메인스트림 워크로드를 위한 더 크고 보안이 강화된
엔클레이브 제공(최대 1TB 메모리 공간)



내장된 암호화 가속기를 통한 개선된 암호화 처리

암호화 알고리즘의 연산 비용을 줄이기 위해 관련 업계를 이끌고 있으며,
획기적인 성능으로 더 많은 데이터를 보다 안전하게 보호해 줍니다.



공개키 암호화

이전 세대 대비

5.6X

단일 스레드 공개키 암호화를 위한
OpenSSL RSA Sign 2048

대칭키 암호화

이전 세대 대비

3.3X

단일 스레드 대칭키 암호화를 위한
AES-GCM

워크로드 및 구성에 대한 정보는 웹사이트(<https://www.intel.com/3gen-xeon-config>)의 [70,71]을 참조합니다. 결과는 다를 수 있습니다.

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®