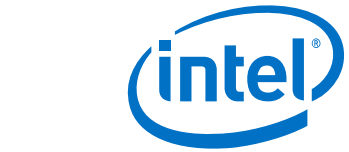

New 3rd Gen Intel® Xeon® Scalable Processor (Codename: Ice Lake-SP)

Irma Esmer Papazian

Intel Corporation, 2020

Notices and Disclaimers

All product plans and roadmaps are subject to change without notice.

Intel technologies may require enabled hardware, software or service activation.

Performance result estimates are as of dates shown in configurations and may not reflect all publicly available updates. No product or component can be absolutely secure.

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 complete information visit http://www.intel.com/performance.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Intel Advanced Vector Extensions (Intel AVX) provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at http://www.intel.com/go/turbo.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

Results have been estimated or simulated.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Ice Lake-SP is new 3rd Gen Intel Xeon Scalable Processor

- 10nm+ process technology
- 2-socket Whitley platform
- Incorporates Sunny Cove core
- Brings scalable and balanced architecture for increased throughput and per-core performance across all workloads in the datacenter

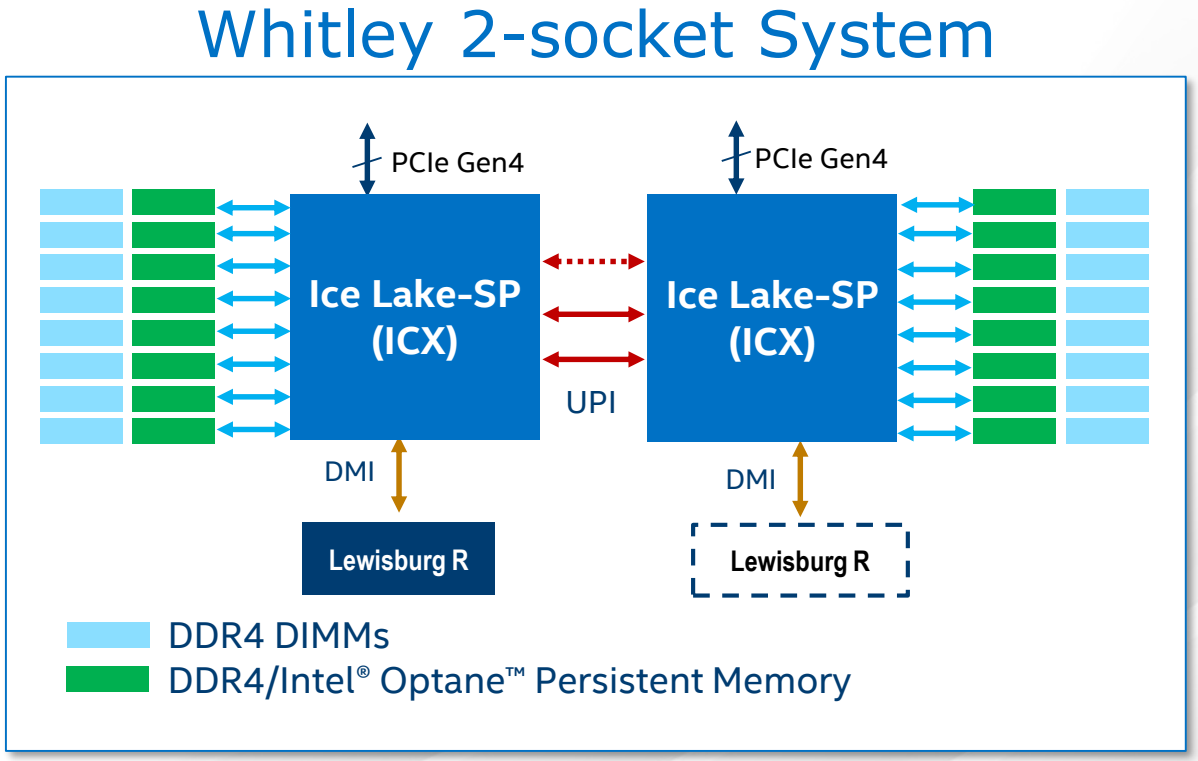

Agenda

Ice Lake Server CPU Core – Sunny Cove

- Core Microarchitecture
- New ISA Usages for Server

Ice Lake Server – The SoC

- Scalable Infrastructure and Architecture
- Power and Performance Optimizations

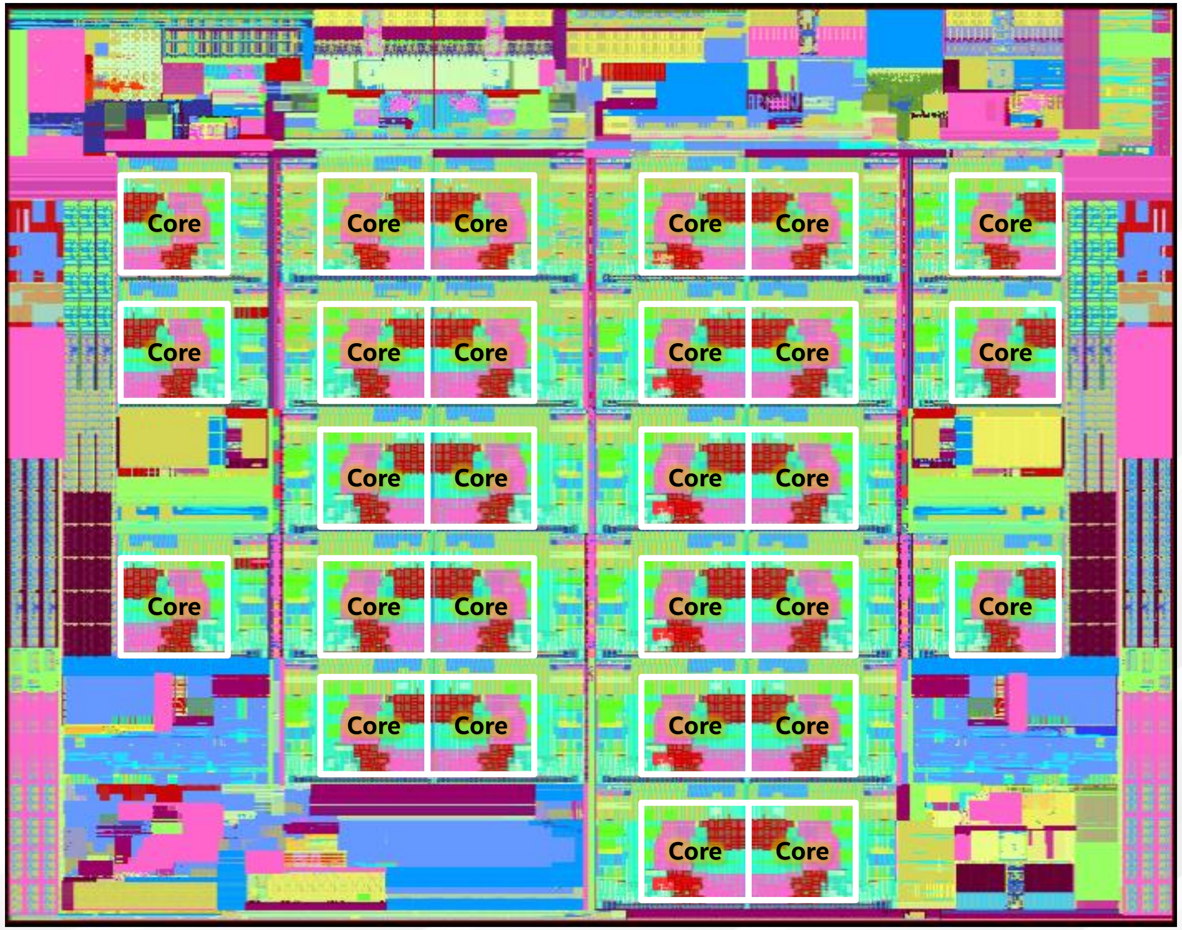

Die picture of a 28C Ice Lake-SP die

Sunny Cove Core Microarchitecture

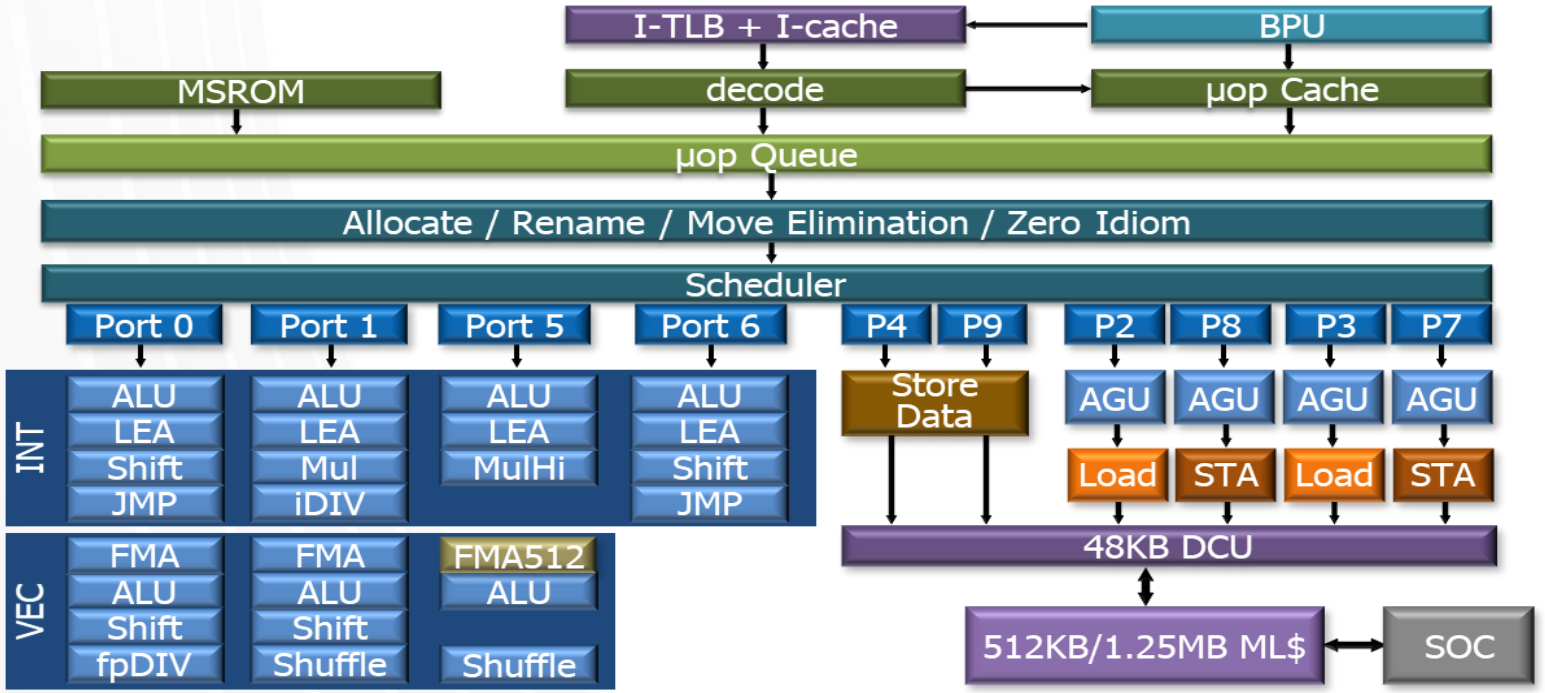

	Cascade Lake (per core)	Ice Lake (per core)
Out-of-order Window	224	352
In-flight Loads + Stores	72 + 56	128 + 72
Scheduler Entries	97	160
Register Files – Integer + FP	180 + 168	280 +224
Allocation Queue	64/thread	70/thread; 140/1 thread
L1D Cache (KB)	32	48
L1D BW (B/Cyc) – Load + Store	128 + 64	128 + 64
L2 Unified TLB	1.5K	2K
Mid-level Cache (MB)	1	1.25

- Improved Front-end: higher capacity and improved branch predictor
- Wider and deeper machine: wider allocation and execution resources + larger structures
- Enhancements in TLBs, single thread execution, prefetching
- Server enhancements – larger Mid-level Cache (L2) + second FMA

~18% Increase In IPC On Existing SPECcpu2017(est) Integer Rate Binaries

Results have been estimated based on pre-production tests at iso core count, frequency and memory BW per core as of July 2020. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Sunny Cove Core New Instructions

Cryptography

- Big-Number Arithmetic (AVX-512 Integer IFMA)
  - VPMADD52 – fused multiply add of 52-bit precision integer values for public key cryptography
- Vector AES and Vector Carry-less Multiply Instructions
  - AES (Advanced Encryption Standard) & AES-GCM (AES mode that allows parallel processing)
- Galois Field New Instructions (GFNI)
  - Encryption algorithms, error correction algorithms, bit matrix multiplications
- SHA-NI – Hardware acceleration of Secure Hash Algorithms

Compression/Decompression & Special SIMD

- Bit Algebra
  - VPOPCNT returns the #of bits set to 1 in byte/word/DW/QW
  - Bit Shuffle – shuffle bits from QW elements
- VBMI – Vector Bit Manipulation Instruction
  - Permutes, shifts, expand, and compress operations
  - Used for columnar database access, discrete mathematics, dictionary-based decompression, data-mining routines

New SIMD ISA Utilizing AVX512 on ICX

- Vector CLMUL
- Vector AES
- VPMADD52
- SHA Extensions
- GFNI

Software / Algorithms

Multi-Buffer

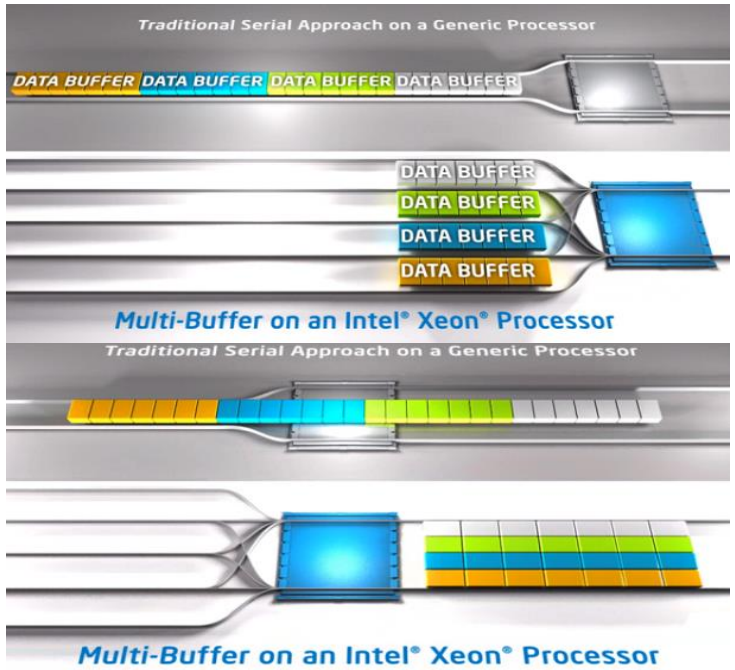

Function Stitching

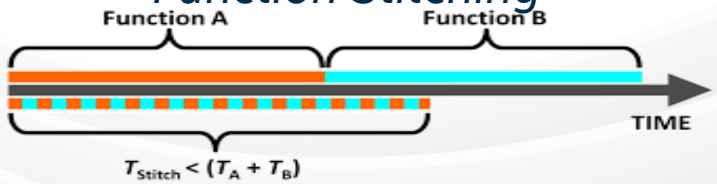

Ice Lake vs. Cascade Lake Per Core Performance

- ECDHE x25519 ≥ 8X
- RSA Sign 2048 ≥ 7.5X
- ECDHE p256 ≥ 6X
- AES-CTR ≥ 3.5X
- AES-CMAC ≥ 3.5X
- AES-XTS ≥ 3.5X
- AES-GCM ≥ 3X
- ECDSA Sign p256 ≥ 3.5X
- CRC ≥ 2X
- ZUC ≥ 1.5X

Algorithms, HW and SW Co-innovation Results In Unprecedented Performance Advances In Cryptography

Ice Lake SP CPU: SoC (all but the cores)

- New infrastructure architecture
- SoC architecture
- Memory Hierarchy and IOs
- Performance and Power enhancements

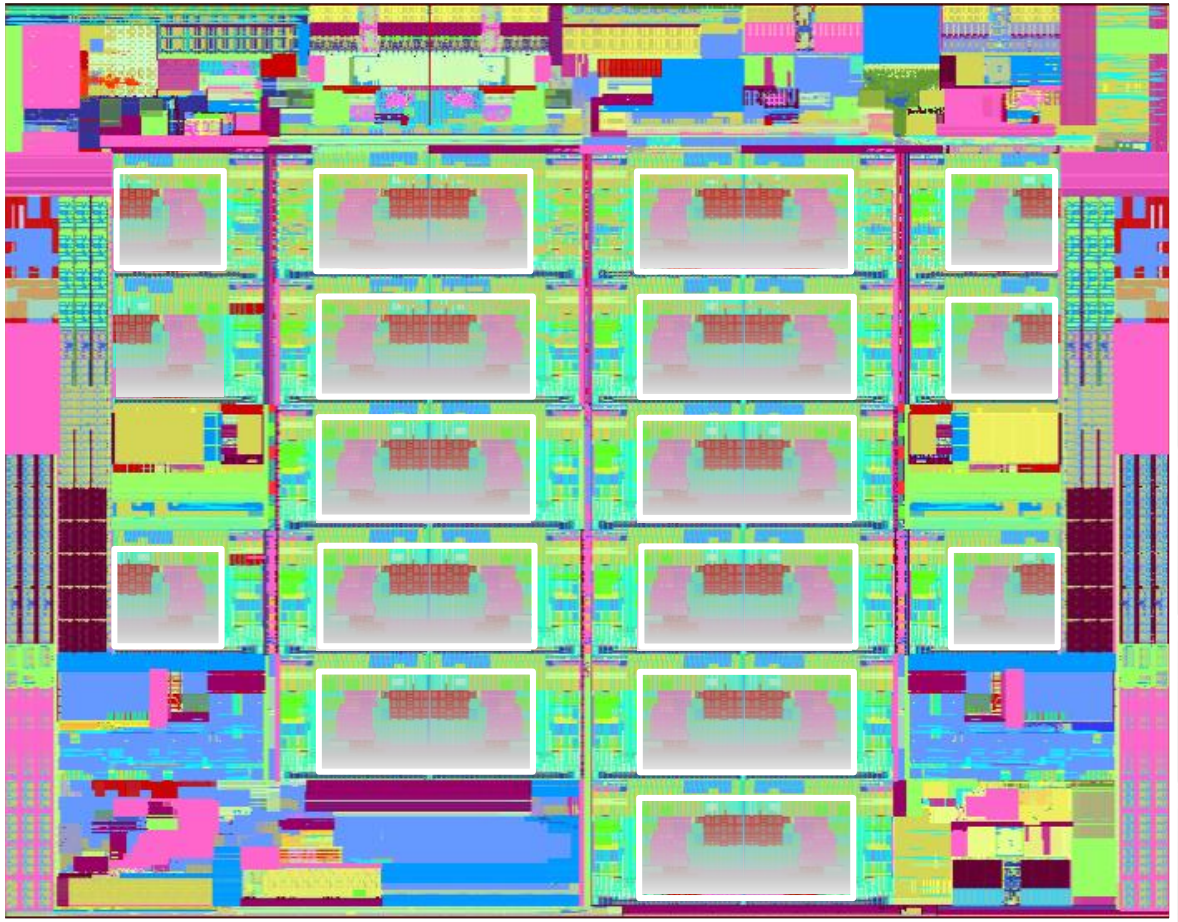

Die picture of a 28C Ice Lake-SP die

New Infrastructure and Control Structure for Improved Scalability and Responsiveness

New Micro-architecture for Infrastructure

- Consistent communication across all IP subsystems for control and telemetry

Dedicated Power Management Sideband

- Each subsystem has a PMA that abstracts IP-specific behavior implements HW FSMs for all PM control and telemetry

Separate General Purpose Sideband

- For all configuration, register access and error handling related communication

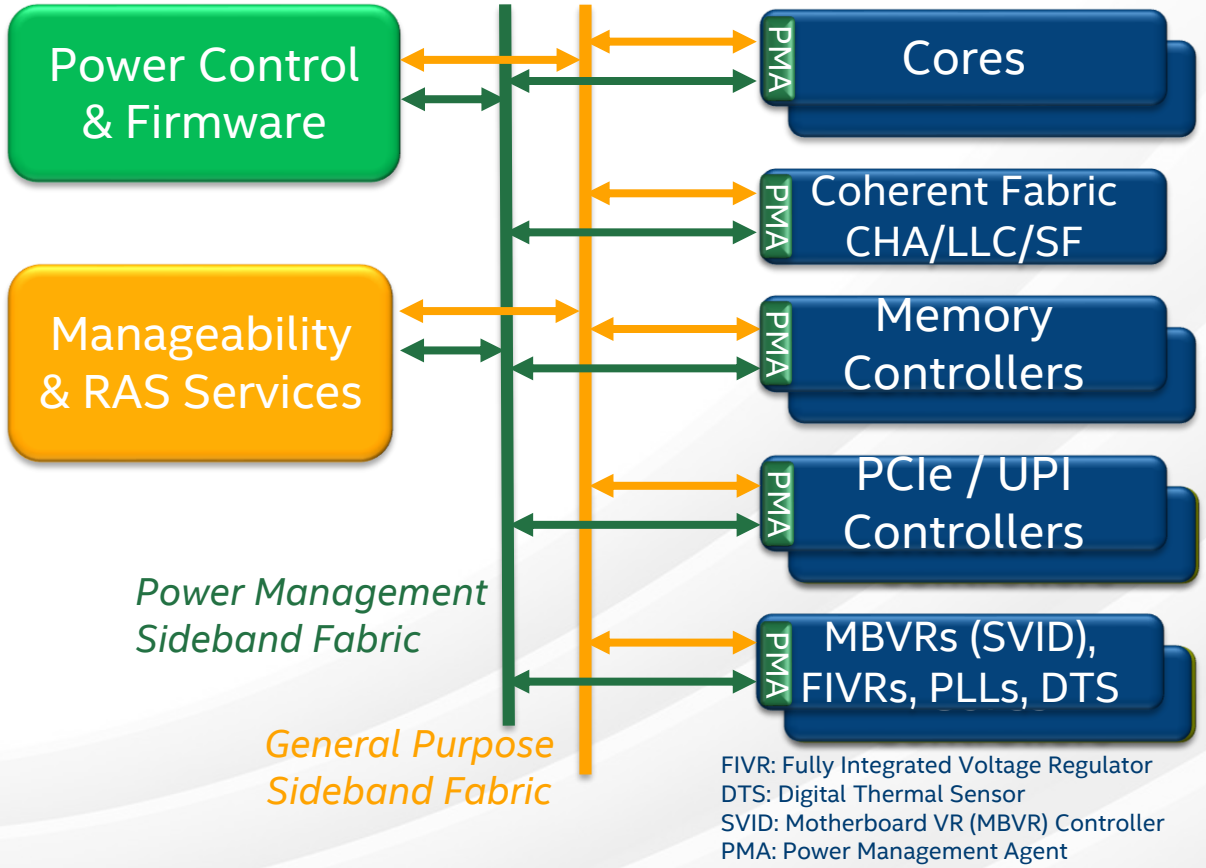

Reduces Complexity, Improves Response Time and SoC Management

SoC Architecture: Cascade Lake-SP to Ice Lake-SP

Cascade Lake SP - aka. CLX (28 cores)

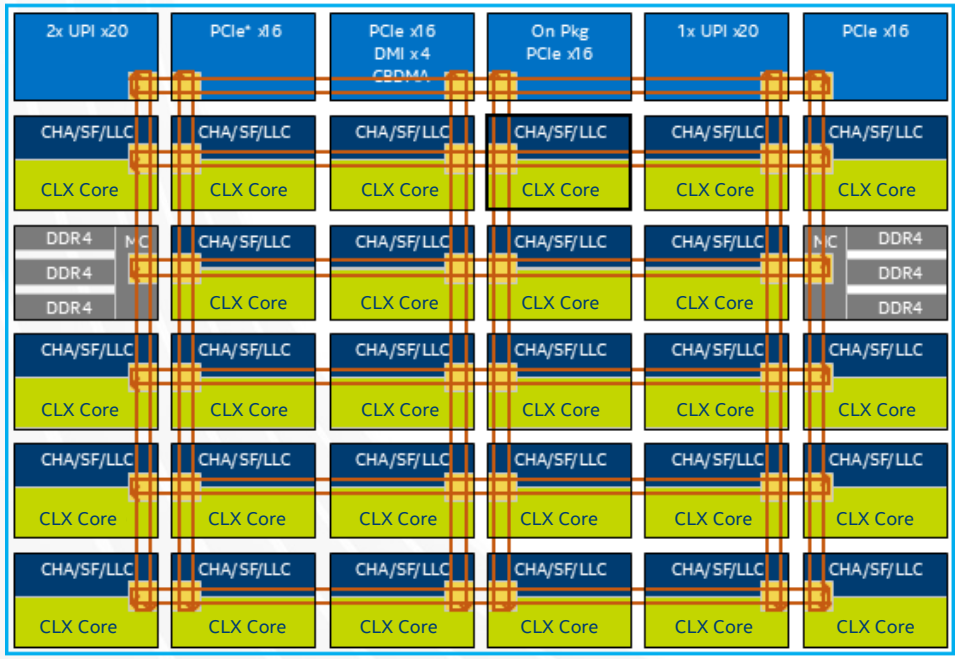

CHA – Caching and Home Agent SF – Snoop Filter, LLC – Last Level Cache
SKX Core – Skylake Server Core UPI – Intel® UltraPath Interconnect

Ice Lake SP (28 core example)

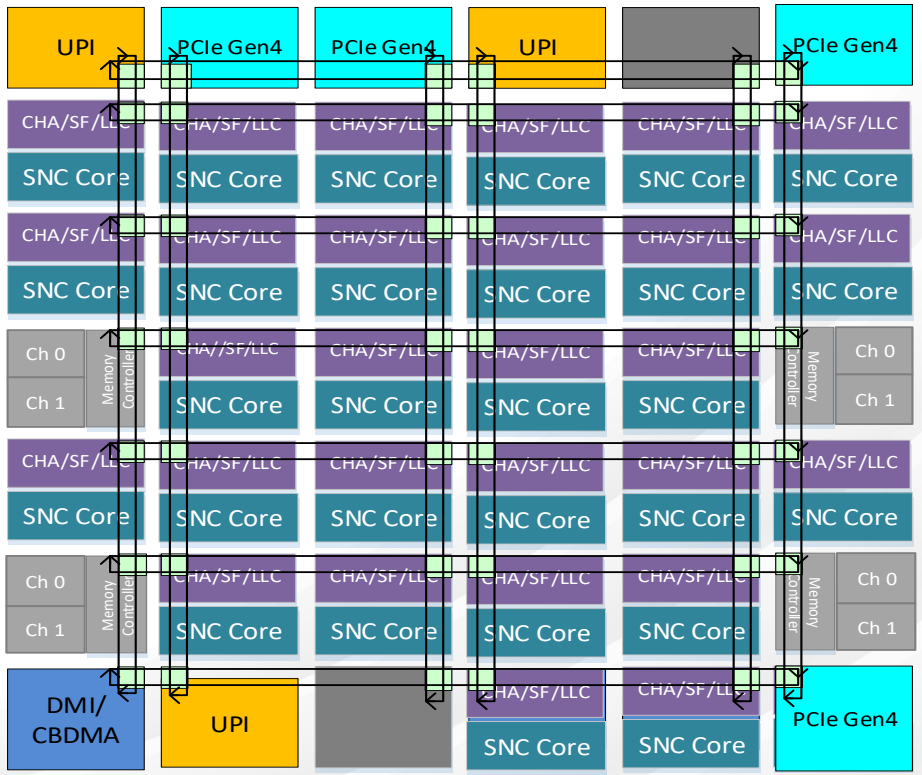

CHA: Caching and Home Agent, SF: Snoop Filter, LLC: Last Level Cache
SNC Core: Sunny cove Core, UPI: Intel® UltraPath Interconnect

Continued Emphasis On Modularity & Balanced Performance Scalability

Ice Lake-SP IO and Memory Hierarchy

Integrating PCIe Gen4 controllers

- New IO Virtualization design, enables up to 3x BW scaling on large payloads (2x frequency, larger TLB, supports 2M/1G pages for in translation requests)
- New P2P credit fabric implementation to reach top P2P BW targets

3 independently clocked UPI links

4 Memory Controllers with enhanced per channel schedulers

- New memory controller design w/ optimizations

Intel® Total Memory Encryption (TME)

- DRAM encrypted using AES-XTS 128bit

Intel Optane Persistent Memory 200 Series (Barlow Pass)

- Higher speed and better power profile

Ice Lake SP (28 core example)

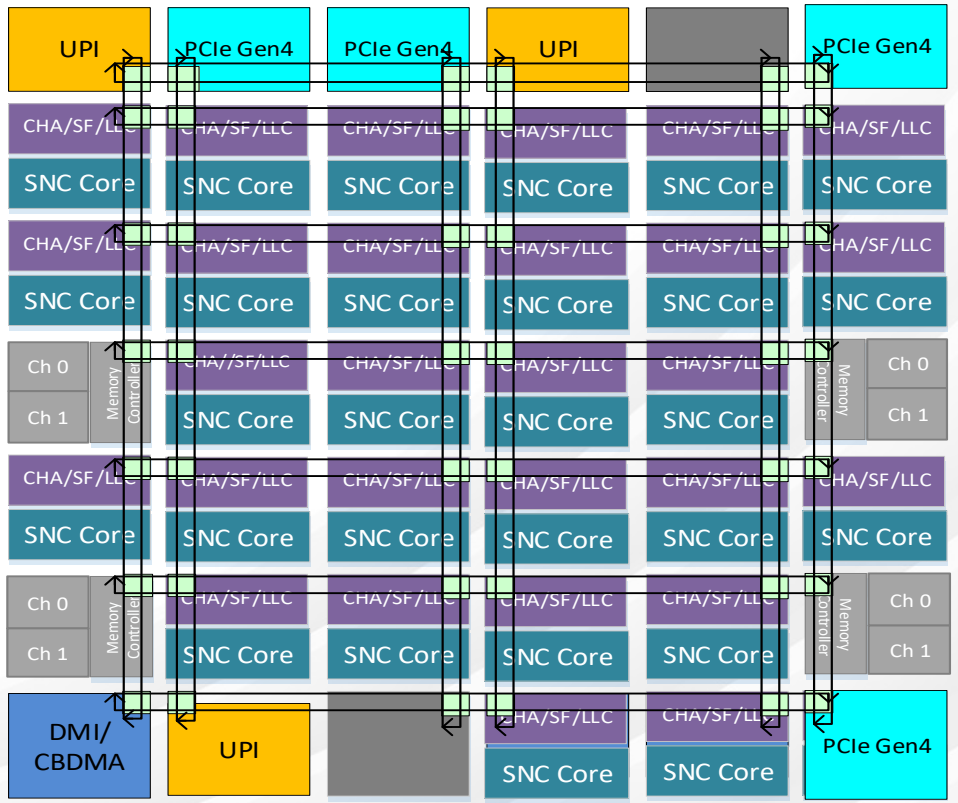

Emphasis On Performance Scalability

Results have been estimated based on pre-production tests at iso core count and frequency as of July 2020. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Latency & Coherence Optimizations

DPT: Dynamic Prefetch throttling

- New Core – SoC handshake; modulates prefetching under heavy memory bandwidth use cases (up to 7% increase on 100% Rd BW micros)

SpecI2M optimization: Convert RFO to specI2M when memory subsystem is heavily loaded

- Reduces mem bandwidth demand on streaming WLs that do full cache line writes (25% efficiency increase)

Non-temporal Write (NTW) Optimization

- Improves low core count active throughput (1 core 50% Rd BW increases >80%, 100% Wr BW increases >3.5X)

IODC: Directory Cache Extended

- Added support for remote core writes, reduces directory write overhead (High memory BW demanding non-NUMA micros benefit 10-90%)

OSB (Opportunistic Snoop Broadcast) Extended

- Added support for new opcodes (data reads and RFO), reduces remote socket cache to cache transfer latency on reads by ~70ns

Hemisphere mode

- Default mode for socket UMA interleaving, optimizes local socket UMA latency if memory is populated symmetrically (~2ns saving)

M2I vs. M2S on socket coherence flow update

- Improves lock scaling performance, expect 2-3% upside on OLTP DB workloads by speeding up compare-exchange operations and atomic updates to critical regions

Results have been estimated based on pre-production tests as of July 2020. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Memory Bandwidth-Latency: Ice Lake vs. Cascade Lake-SP

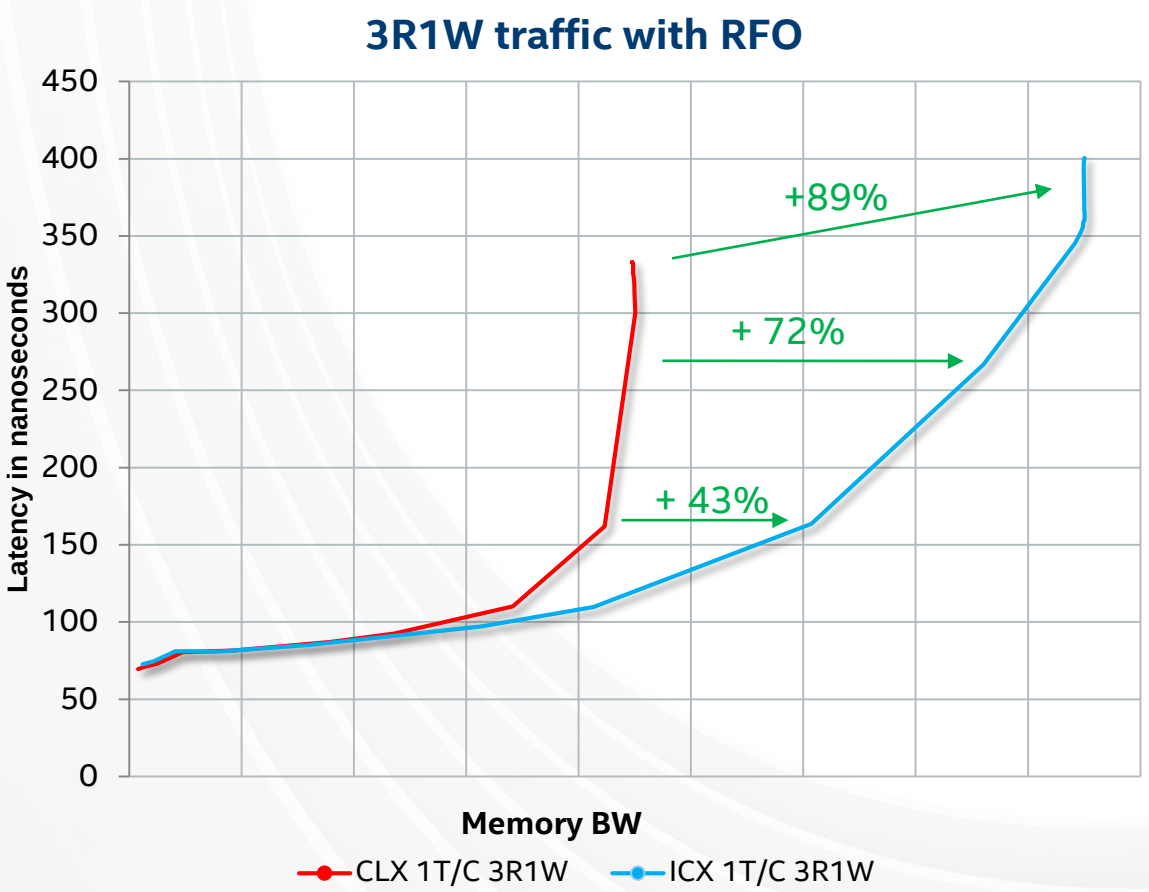

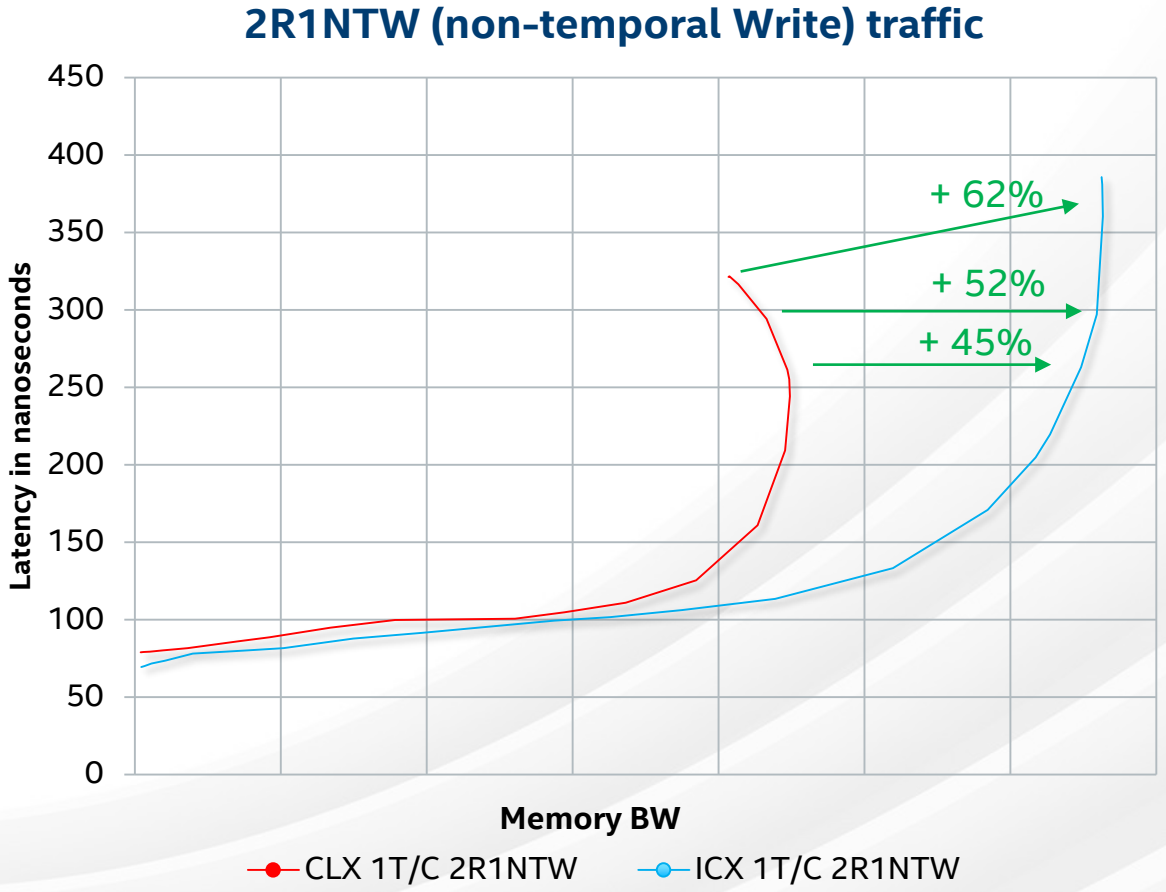

Improved Efficiency Beyond Memory Speed * Channel Count Increase

Power Management Latency Improvements

Goal: Minimize impact of PM latencies on performance

- Fast Core Frequency Change
  - Capability to move the PLL from current to target frequency in a continuous sweep without stopping the clocks
  - Allows fast P-state transitions optimizing power/perf without latency cost
- Coherent Fabric ("mesh") Drainless Frequency Change
  - Capability to move the PLL from current to target frequency without draining the buffers
  - Reduces frequency transition time by ~3x

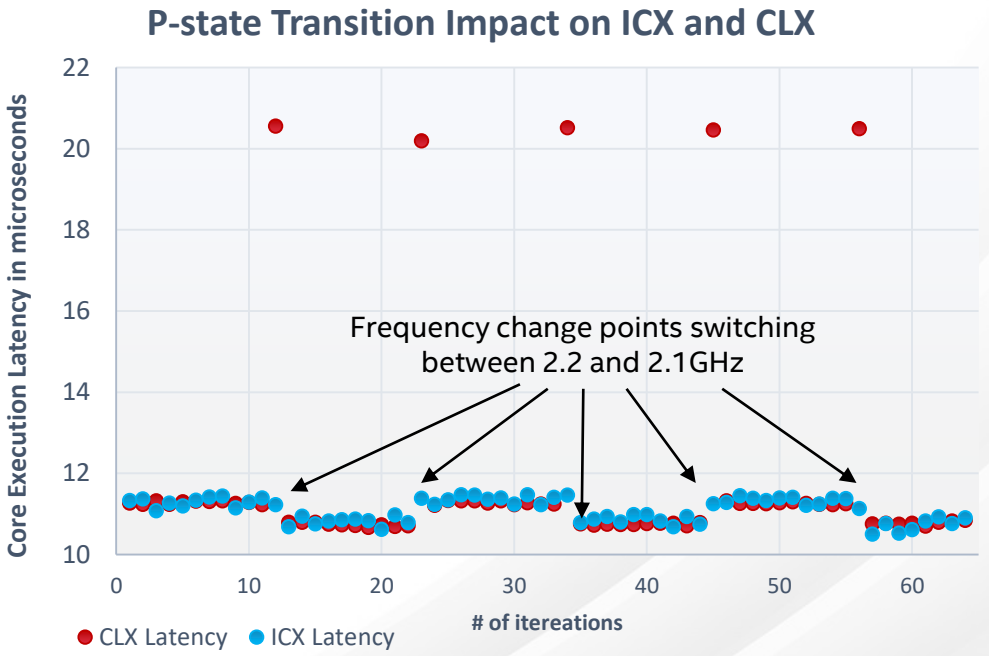

Time	CLX	ICX
Core Frequency Transition block time	12 us	~0 us
Mesh Frequency Transition – IO Block time	20us	7 us
Typical C6-State Exit Time	30 us	~20 us

Results have been estimated based on pre-production tests as of July 2020. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Power Management Latency Improvements

Goal: Minimize impact of PM latencies on performance

- Fast Core Frequency Change
  - Capability to move the PLL from current to target frequency in a continuous sweep without stopping the clocks
  - Allows fast P-state transitions optimizing power/perf without latency cost
- Coherent Fabric (“mesh”) Drainless Frequency Change
  - Capability to move the PLL from current to target frequency without draining the buffers
  - Reduces frequency transition time by ~3x

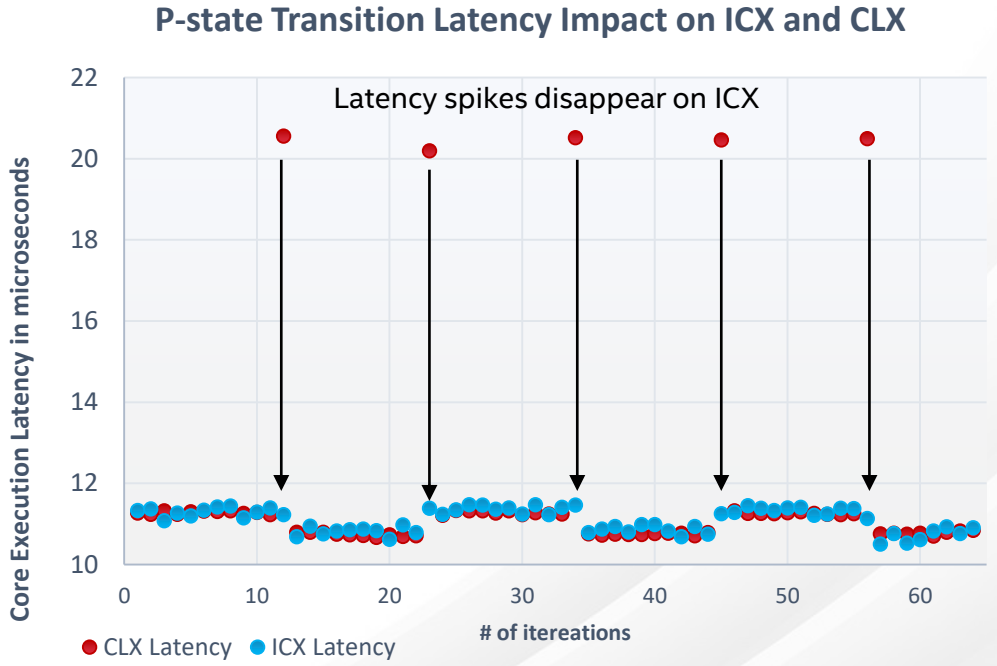

Time	CLX	ICX
Core Frequency Transition block time	12 us	~0 us
Mesh Frequency Transition – IO Block time	20us	7 us
Typical C6-State Exit Time	30 us	~20 us

Results have been estimated based on pre-production tests as of July 2020. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

AVX Frequency Improvements

Goal: Minimize frequency impact on AVX 512 bit operations when not bounded by physical limits

- Not all AVX512 instructions consume high power
  - 512-bit loads, 512-bit stores, 256-bit FP, integer multiply are a few examples
  - Smarter mapping between instructions and specific power levels

Provides software writers greater latitude when using these instructions to optimize their code for performance

Power level	Class of instructions	Instruction types per class (not the full list)
0	SSE/256L	all 64b and 128bit
	256H	FP Mul, INT Mul, VNNI, FMA 256b
1	512L	VPCLMUL, VAES, VBMI, Ld, St 512b
2	512H	FP Mul, INT Mul, VNNI, FMA, VPMADD52 512b

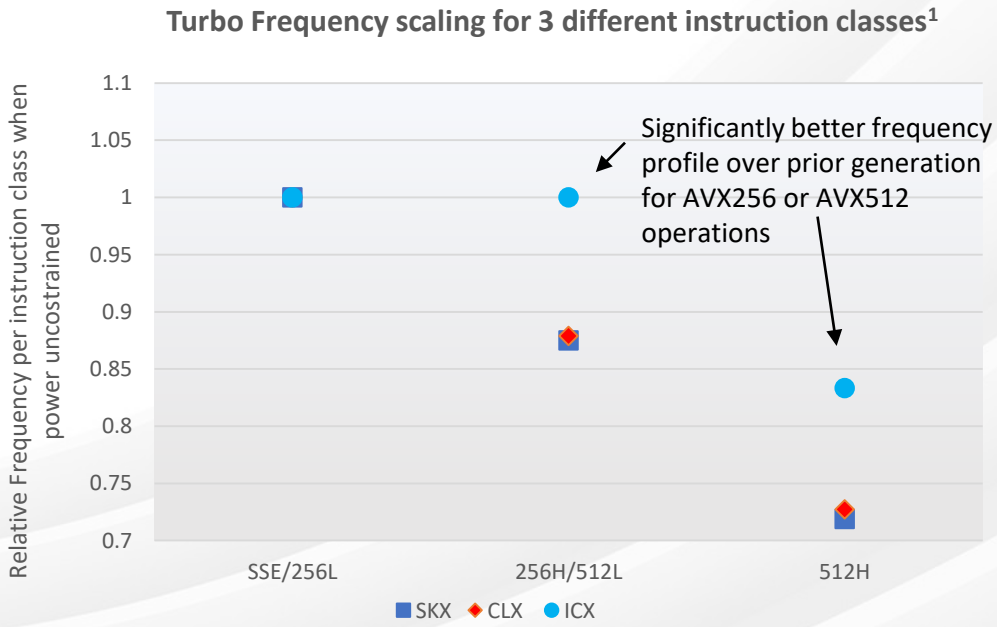

[1]Baseline: all cores active turbo frequency for SSE for each product

Intel® Speed Select Technology (Intel® SST) Features

Offers a suite of capabilities to allow users to re-configure the processor – dynamically, at runtime to match the usage / WL and maximize performance

Intel Speed Select Technology–**Performance Profile** (Intel SST-PP)

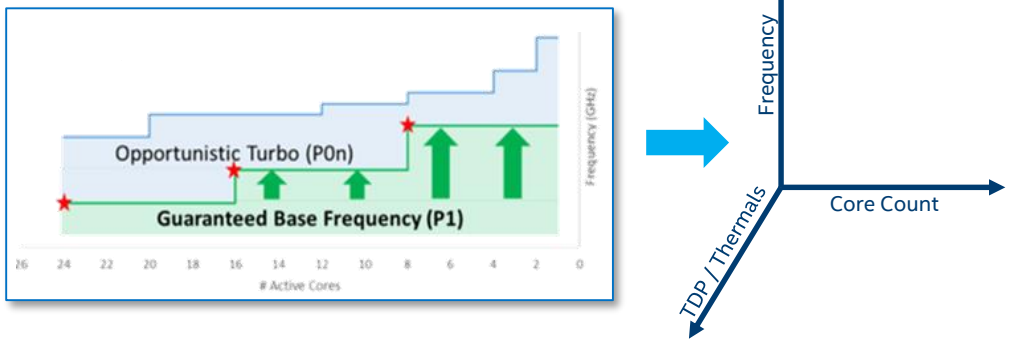

Intel Speed Select Technology–**Base Frequency** (Intel SST-BF)

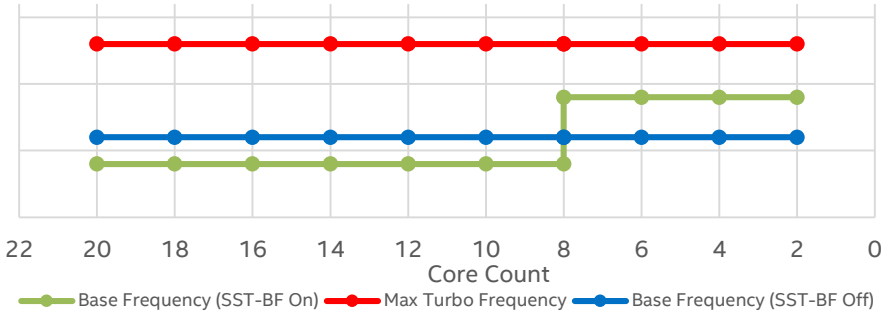

Intel Speed Select Technology–**Core Power** Intel SST-CP

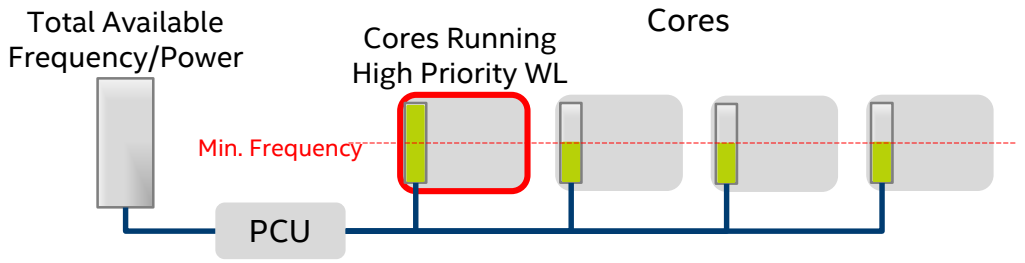

Intel Speed Select Technology–**Turbo Frequency** (Intel SST-TF)

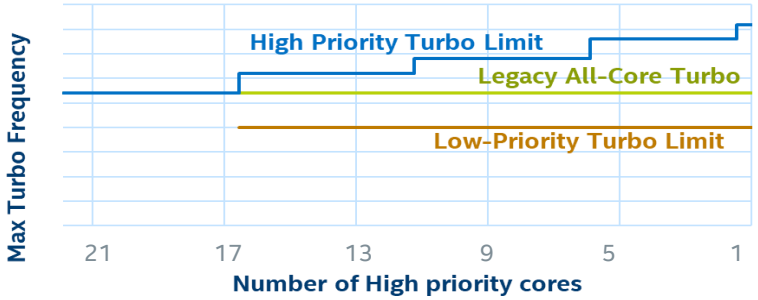

New Intel SST Capabilities Enable Prioritization Of Critical WLs with Ease Of Use & Deployment

Wrap Up

New 3rd Gen Intel Xeon Scalable Processor
(Codename: Ice Lake-SP)

- First Intel Xeon Scalable Processor using 10nm technology
- Sunny Cove core improves per core performance with existing binaries
- New software innovations and new ISA together lead to unprecedented per core performance increases
- With new scalable infrastructure, power and performance optimizations Ice Lake CPU achieves both high throughput and per core performance

Ice Lake: A Balanced CPU For All Server Usages

Results have been estimated based on pre-production tests as of July 2020. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Special Thanks

Wajdi Feghali and Crypto Team

Nilanjan Palit

Ritesh Parikh

Prateek Gupta

Rupin Vakharwala

Ice Lake-SP Architecture, Design,
Validation Teams and Program Management

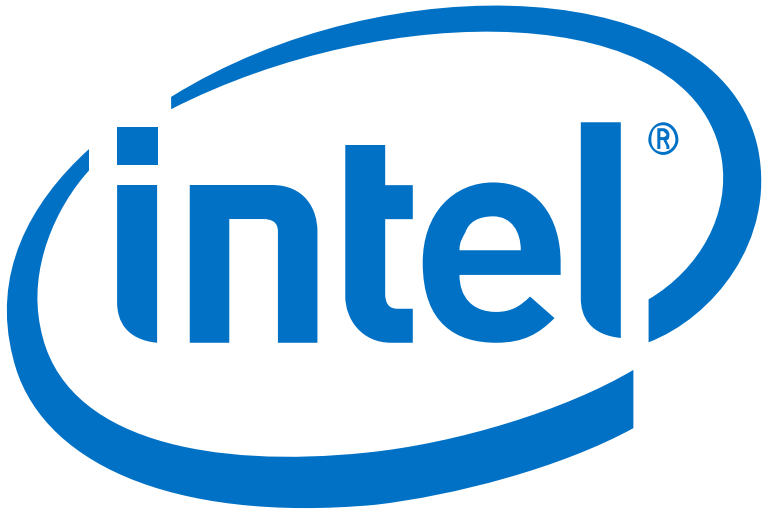
intel