

Data 2024 and Beyond

AFCEA – Enabling Distributed C2 Symposium

CW5(R) Jason K. Dunn-Potter
January 31, 2024



intel®

Agenda

- It starts with Intel
- Extreme Search – Neuromorphic
- Artificial Intelligence (AI) advancements
- CPU Roadmap (4th Gen / 5th Gen)
- CXL Intro
- Graphics Processing Units (GPU's - Arc, Flex, Max)
- Cloud Cybersecurity (SGX/TDX)

What is Intel?

Autonomous Vehicles



Aerial Art



Foundry



5G Farming



Graphics Cards



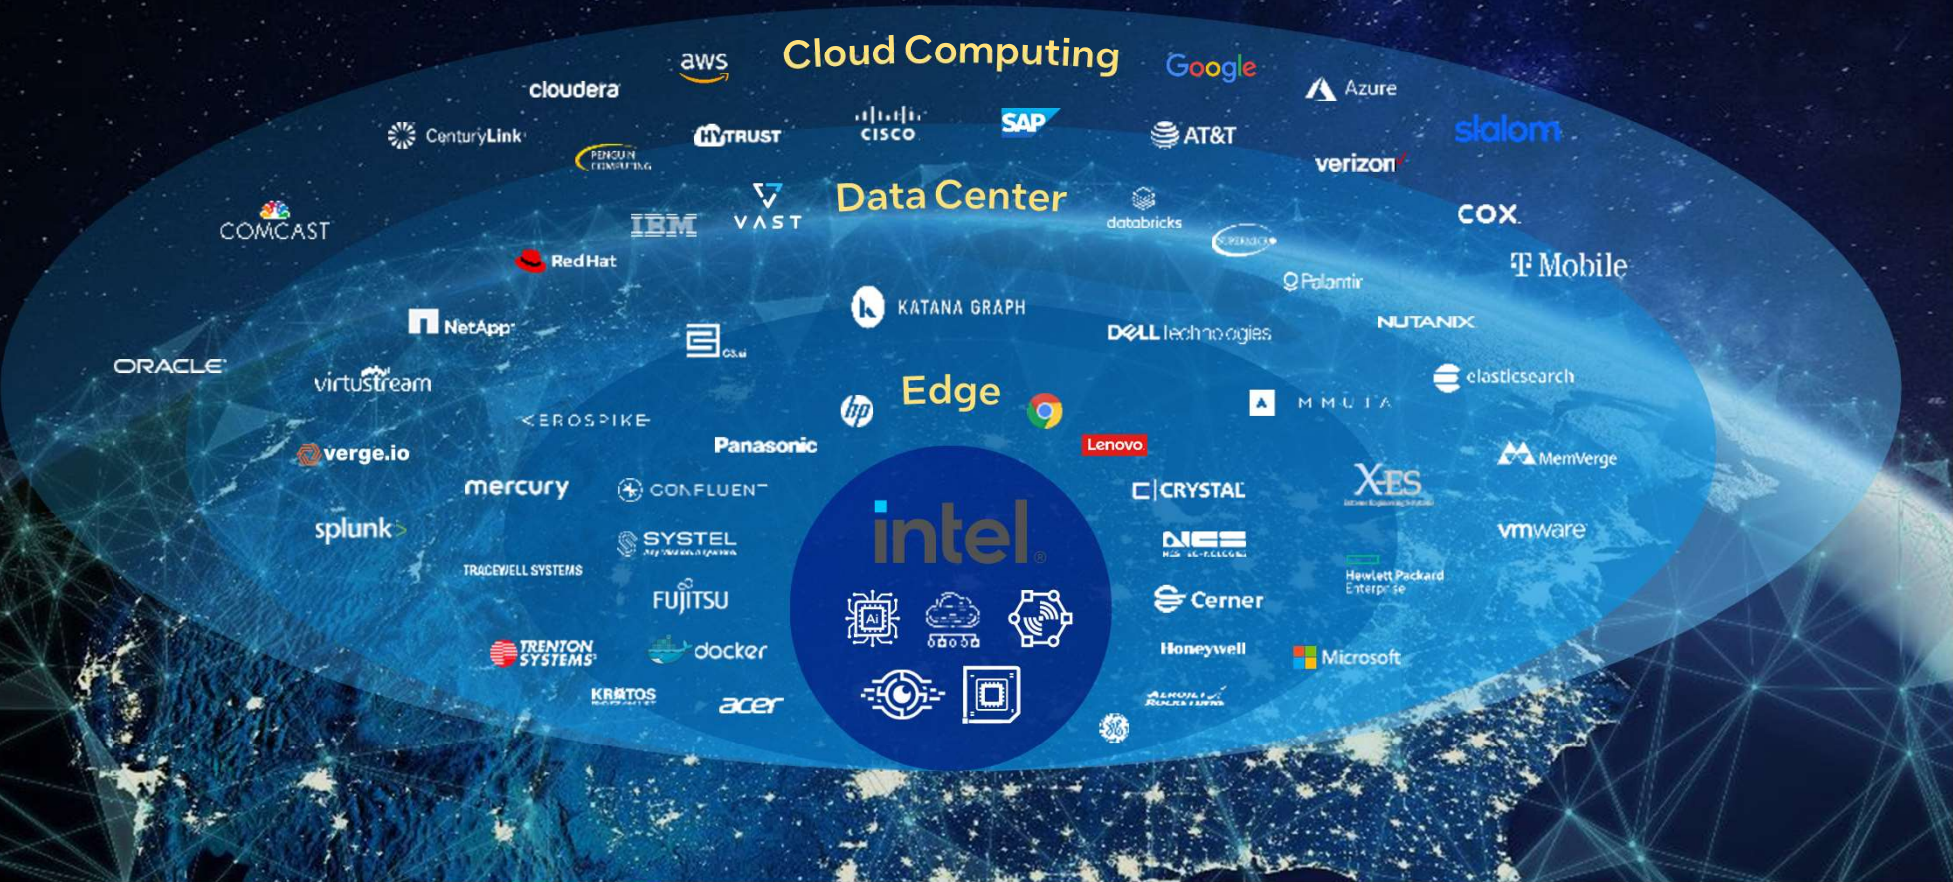
Cloud



GenAI



Ecosystem Collaborators



Extreme Search – Neuromorphic Processors



intel®

How Lewis Rhodes Labs (LRL) created an extreme search capability...

In 2021, Lewis Rhodes Labs in conjunction with several industry partners delivered a new Extreme Search capability. Intel & Dell has since refined the innovation into a deliverable capability...Extreme Search Server.

Follow the story how they went from solving TBI's in children to changing edge technologies forever...
From Neurology to Neuromorphic Computing #141 | Embracing Digital Transformation
<https://www.youtube.com/watch?v=Dv4AdtHNQio>




© Lewis Rhodes Labs, Inc

Intel Property



Unleashing the Potential of Data with Search in Storage



4x cards
per server



Dell R750XA

Dual Slot PCIe Gen4
NPUs on Intel® Agilex™

Each card contains: 16 TB NVME SSD
Search Bandwidth: 20 GB/s per card

64 TB NVME SSD

Total Search

Bandwidth: 80 GB/sec/server

Search time: <12 minutes*



© Lewis Rhodes Labs, Inc

intel.

NPU Search Card for Extreme Search

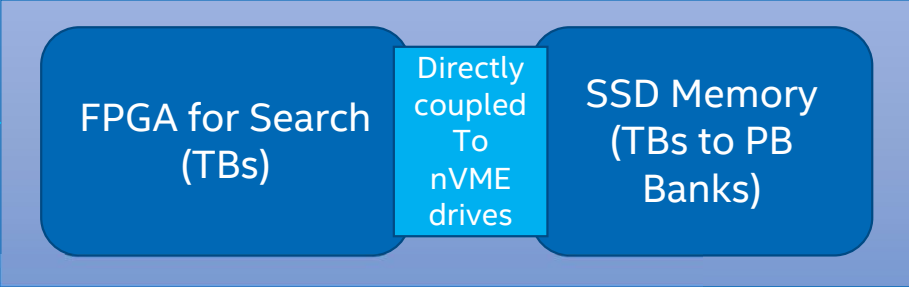
What is it?

- Full Sized GPU type card including with:
 - 2x Intel Agilex FPGA (Programmable CPU's)
 - 16 TB of M.2 SSD's (Storage)*
 - PCIe Gen4 compatible
 - Gluster File System (Scalable)
 - Running off of Python (Jupyter Notebook) API Interface

What it's NOT?

- AI
- CPU
- GPU
- Replacing your current software applications
- Replacing your Archive Storage solution

Acts as 2 devices:
1) Search Storage Results
2) Search device



Indicates where the matches are located



*Each System contains 4x Cards
8x FPGA's
64x TB total

Extreme Search Characteristics

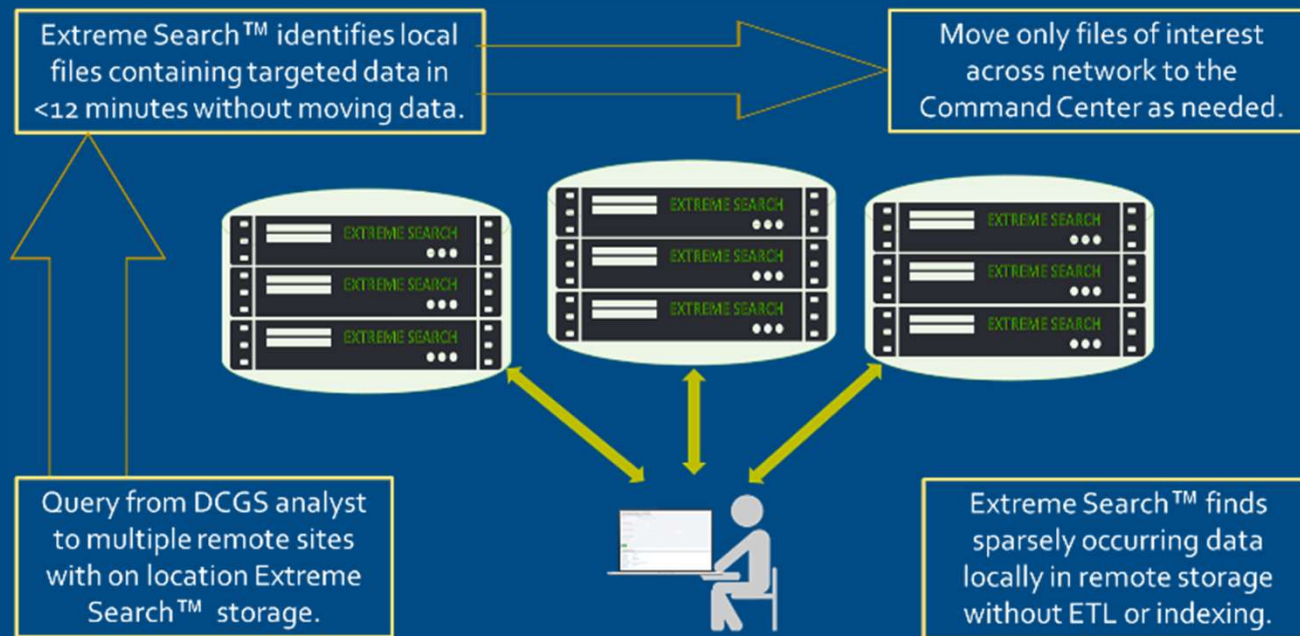
- Raw data analysis - No indexing or other preconditioning required
- Byte level data analysis and results in near real time (analyze as you collect data)
- Brings compute to data – not data to compute – data stays where it resides
- Queries are based on 2x lines of Python and can be run by any software that already queries data
- Operates under 50W, <1% of equal performance server architectures
- **200x improvement in performance over existing architectures**
- Field Tested - Deployed & in contiguous operations to a DOE Lab for over 24 months



"We integrated the technology in a matter of days and optimized some cybersecurity search analysis from 600 minutes (10 hours) to 3 minutes total. It allowed us to find needles in haystacks and build better haystacks".

Cyber Security R&D, DOE Lab, Aug 2023

Centralized Visibility on Distributed Data Sets



There are various options that are simultaneously available:

- 1: Many - One user to multiple data stacks
- Many: 1 - Multiple users to a single stack
- Many: Many – Multiples users accessing multiple data stacks

Optimal Use Cases

1. Cybersecurity (Splunk/Elastic/etc)

- a) Firewall - Comb 100% of all PCAP (Packet Capture) collected data to ID known issues / concerns (SolarWinds)
- b) Foot printing - ID Cyber foot printing across all networks based on uncommon patterns over time
- c) Heuristics - Analyze and ID unknown patterns (Log4j)
- d) Zero Day Defeat - Establish zero-day defense based on known vulnerability issues while mitigation strategy is underway. Note: Patch development and implementation is underway – Which can take years for some systems to actually get fully patched.

2. AI Advancement

- a) Retraining - Pull needles of key data points & build a better model faster, import only the changes in data behaviors

3. Electronic Warfare

- a) ID activity while still airborne and relay to second wave
- b) Triangulate transmission locations using GPS and radio telemetry (Build a 3d model map of Electromagnetic Spectrum by AOR)

4. Content detection

- a) Spam - Content filtering of spam and junk mail across the entire email server / active directory of the several million users across the DOD in 7 continents from any of the HQ's.
- b) Patterns - Content detection of nefarious content based on known hash's, point of origin, VPN traffic that pushed past border firewall, but was picked up inside the network via Zero Trust (Exploited children, Rootkits, Logic Bombs, Polymorphic, etc)
- c) Content filtering - Content detection of policy violations (hate speech, known bad actors, collusion of suspected criminal activity)

Optimal Use Cases

5. Extreme Data analysis

- a) Generic - Hot data v. Cold Data (accelerate data via digital twins that can be dumped after analysis)
- b) Medical – HIPAA – ID data trends while maintaining data policies – How many O-Neg patients does the community have?
- c) Supply – How much of a critical resource did we use/have over the last 10 days at all sites
- d) Legal - Legal review of all case files in the firm (or in the Judicial reviews)
- e) Law Enforcement - Digital crimes unit combing through decades of cold case files looking for patterns aka “MO’s”
- f) Chain of Custody – Query data while maintaining its integrity supporting both Crimes units & Attorney’s
- g) Finance - Bank audits of internal archives looking for low yield scammers
- h) Advanced Trend Analysis - Telco – ID Span callers and tracing point of origin’s (especially if they shift numbers but not techniques) based on profiles (ie. If 99 out of 100 callers hang up in the first 10 seconds).
- i) Telco – DDOS attacks reduction – Defeating a zombie attack (ie. DDOS – Distributed Denial of Service)....stop it at the roots everywhere.

6. Extreme Edge - Query remote edge devices in remote and austere environments

- a) Oceanic & Deep sea - USV (Unmanned Surface Vehicles), Submersibles, Buoys, etc.
- b) Aerial - UAS aka “Drones” / Blimps / Solar Sails
- c) Space - Comms, ISR, etc
- d) Mountain tops (5G antenna’s, GSR’s (Ground Surveillance Radar’s))
- e) Austere Environments – (Tundra / Desert / Tropical)

Optimal Use Cases

7. IC Community

- a) Large Scale Open-source scrub – Identify keywords utilized on web content or mobile devices across multiple mediums (Websites, Social Media, Logs, Blogs, Databases, Email servers, etc)
- b) Log analysis – Answering the 5W's of log data across weeks or months of logs (who accessed a system, when, and what changed)
- c) Data Integrity - Run remote queries in segmented environments which allows the content owner to provide data sets without providing write access to material.
- d) Create isolated “triggers” that activate third party effects across multiple classifications without moving or modifying data. Thus, creating a “Shrek Effect” where one site does not understand why your site asked for a query. This creates an environment where a sensor sweep for an unclass item then triggers a CDS (Cross Domain Solution) activity on the “high side”.
- e) Reverse “Shrek Effect” – High side detects an issue based on a finding, thus pushing out queries to each remote site at different classification levels. This obfuscates who is asking the site to make changes or the why.
 - Example: Creditable intelligence indicates a threat to a key CVE (Known Vulnerability) or Zero Day (that should already be patched) that is actively being exploited. This would create a query to push to all sites at every classification to actively scan for this CVE and report to the local elements as a finding and to take priority action to mitigate.
- f) Content filtering - Content detection of policy violations (hate speech, known bad actors, collusion of suspected criminal activity)



CXL Board of Directors



Jim Pappas
CXL Chairman and President
Director of Technology Innovation
Nearly 30 years at Intel – 14 at HPE
*Led the Creation of USB
*Led the Creation of PCIe

Industry Open Standard for
High Speed Communications

200+ Member Companies

Compute Express Link™ and CXL™ Consortium are trademarks of the Compute Express Link Consortium.



Reducing Complexity with Intel® Optimized AI Software



Data Analytics at Scale*

Machine & Deep Learning Frameworks, Optimization and Deployment Tools*



Get Started Here → <https://www.intel.com/content/www/us/en/developer/topic-technology/artificial-intelligence/overview.html>

Note: components at each layer of the stack are optimized for targeted components at other layers based on expected AI usage models, and not every component is utilized by the solutions in the rightmost column

*This list includes popular open source frameworks that are optimized for Intel hardware



Accelerate End-to-End Data Science and AI



Intel Developer Cloud and Intel Developer Catalog

Try the latest Intel tools and hardware, and access optimized AI Models



Full stack ML operating system

Intel® Geti

Annotation/Training/Optimization Platform



Intel optimizations and fine-tuning recipes, optimized inference models, and model serving



Open, Standards-Based unified programming model

One Programming Model for Multiple Architectures & Vendors – **CODE reuse, no need to reprogram or start from scratch.**

Freedom to Make Your Best Choice

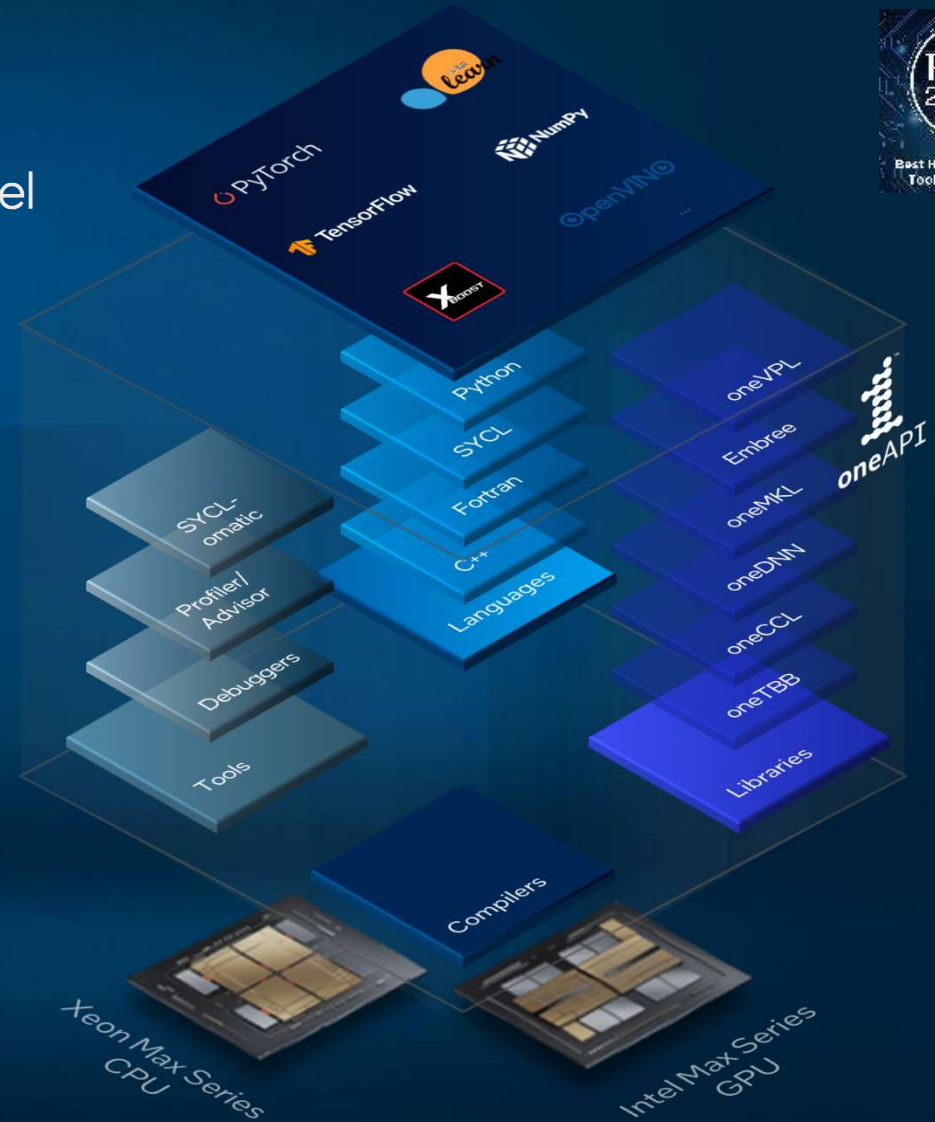
- Choose the best accelerated technology the software doesn't decide for you

Realize all the Hardware Value

- Performance across CPU, GPUs, FPGAs, and other accelerators

Develop & Deploy Software with Peace of Mind

- Open industry standards provide a safe, clear path to the future
- Compatible with existing languages and programming models including C++, Python, SYCL, OpenMP, Fortran, and MPI



Intel® Hardware Shield

Industry leading PC security

SECURE THE PLATFORM



		2019 WHL	2020 CML	2021 TGL	2021 RKL	2022 ADL
Advanced Threat Detection	GPU Offloading w/ TDT-AMS	✓	✓	✓	✓	✓
	Crypto-mining detection w/ TDT	✓	✓	✓	✓	✓
	Ransomware detection w/TDT	✓	✓	✓	✓	✓
Application & OS Protection	Intel® Control-Flow Execution Technology (Intel® CET)*			✓		✓
	Intel® Total Memory Encryption (Intel® TME)			✓		✓
	Intel® Virtualization Technology (VT-X)*	✓	✓	✓	✓	✓
	Intel® Virtualization Technology (VT-D)*	✓	✓	✓	✓	✓
	Mode-Based Execution Control (MBEC)*	✓	✓	✓	✓	✓
	Intel® AES-NI	✓	✓	✓	✓	✓
	Intel® Secure Key	✓	✓	✓	✓	✓
	Virtualized Application Interrupt Controller (vAPIC)			✓	✓	✓
	Multi-Key Total Memory Encryption(MK-TME)+					✓
	Intel Virtualization Technology (VT-d w/ Posted Interrupts)					✓
	Hypervisor Linear Address Translation(HLAT)+					✓
Below-OS Security	Intel® BIOS Guard	✓	✓	✓	✓	✓
	Intel® Boot Guard*	✓	✓	✓	✓	✓
	Intel® Platform Trust Technology (PTT)	✓	✓	✓	✓	✓
	Intel® Trusted Execution Technology (TXT)*	✓	✓	✓	✓	✓
	Intel® Runtime BIOS Resilience (IRBR)	✓	✓	✓	✓	✓
	Intel® System Security Report (ISSR)	✓	✓	✓	✓	✓
	Intel® System Resources Defense (ISRd)			✓	✓	✓
	Firmware Update Restart (CF)			✓	✓	✓
	Firmware Update Recovery(CF)			✓	✓	✓
	Firmware Telemetry (CF)					✓

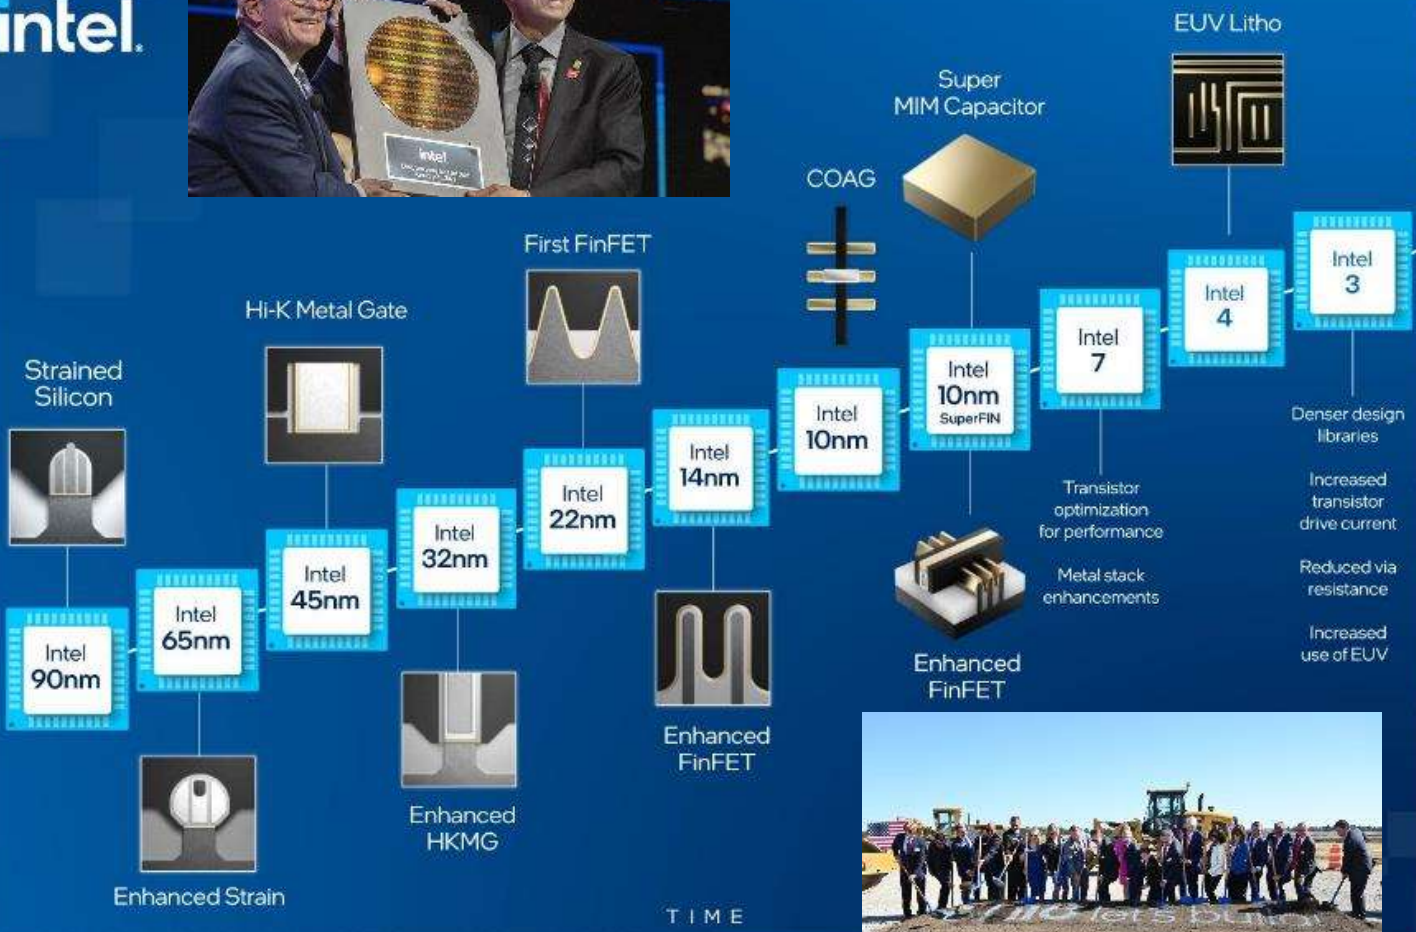
vPro exclusive | MSFT Secure Core PC req* | ADL vPro only +

Enables full-stack PC protection

Security Starts with Intel



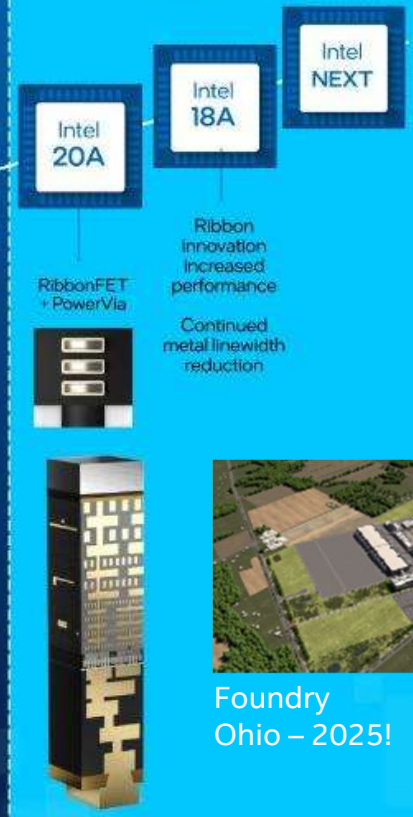
PERFORMANCE PER WATT



EUV Litho



Angstrom era



Foundry Ohio - 2025!

Developer Tools for 4th Gen Intel[®] Xeon[®] Scalable Processors

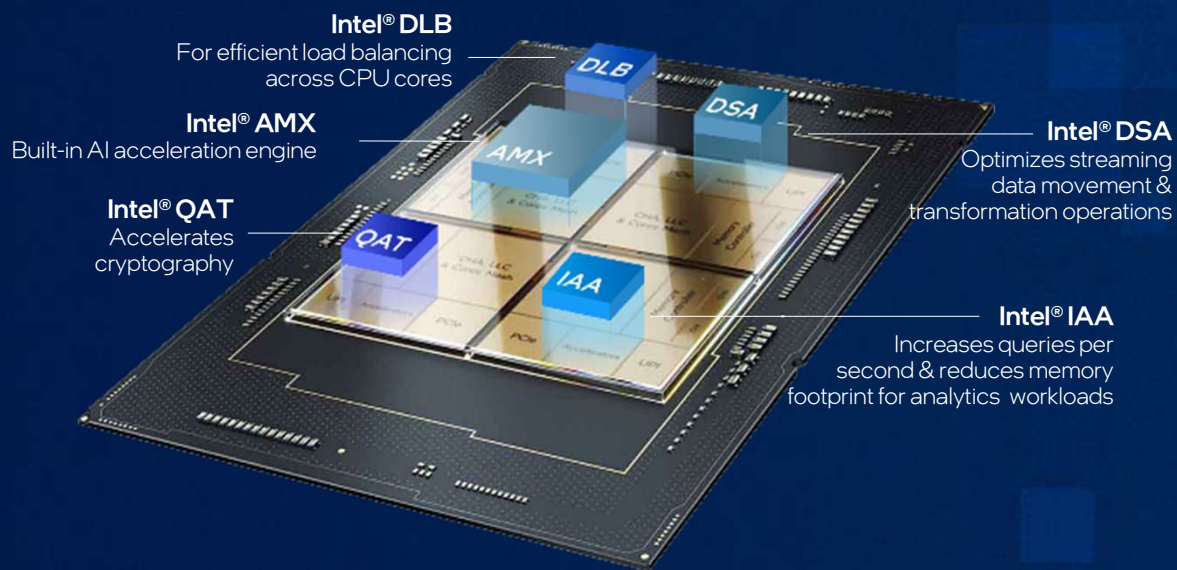
Intel oneAPI, AI tools and optimized AI frameworks help developers maximize application performance by activating advanced capabilities of 4th Gen Intel[®] Xeon[®] Scalable processors and Intel[®] Max Series processors. In multiarchitecture systems with Intel Xeon processors and Intel GPUs, using a single codebase through [oneAPI](#) delivers productivity and performance.

[Compilers, libraries & analysis tools](#) support built-in accelerators to unleash performance, and fast training and inference for AI workloads.

- **Intel[®] oneAPI Math Kernel Library**
for HPC and technical compute
- **Intel[®] oneAPI Deep Neural Network Library**
for deep learning training + inference
- **Intel[®] Query Processing & Intel[®] Data Mover Library***
for query processing, compression and data movement
- **Intel[®] VTune™ Profiler**
helps locate time-consuming parts of code and identify significant issues affecting application performance

Learn more: [Software for 4th Gen Intel Xeon & Max Series Processors](#)

*Intel[®] QPL is open source. Open source Intel[®] DML in beta, v1 coming soon



OpenVINO

Powered by oneAPI

Client Security



intel®

The World's Clouds Run on Intel



25 REGIONS
81 AVAILABILITY ZONES
245 COUNTRIES/TERRITORIES



60+ REGIONS
67 AVAILABILITY ZONES
140 COUNTRIES/TERRITORIES



28 REGIONS
85 AVAILABILITY ZONES
200+ COUNTRIES/TERRITORIES



100% REGIONS
100% AVAILABILITY ZONES
100% COUNTRIES/TERRITORIES

Intel Optimizes, Accelerates and Operationalizes Public Multi-Cloud



Discover

Workload Strategy

Cloud Workshops

- Strategy
- Foundational - Organizational Change & Operations
- Cloud Security
- DevOps
- Cloud Native Development
- Application Modernization & Placement

Assess

Application Strategy

- Rehost, Replatform, Refactor, Retain, Replace, Retire

Cloud Tools

- Intel® Migration Advisor
- Benchmarking – Intel® DC Perf Kit
- Edge to Cloud Reference Architecture

Operationalize

Plan, Build, Modernize, Deploy

Service Recommendation

- IaaS / PaaS identification & configuration
- Benchmarking

Modernize

- Application Modernization
- Orchestration

Monitor & Optimize

Monitor

- Telemetry & automation
- Bottleneck Analysis

Optimize

- Integration with Intel® acceleration technology
- Cost & Performance

Cloud Tools


- Intel® Cloud Optimizer
- Intel® Workload Optimizer
- Intel Platform Monitoring Units
- Intel Cloud DevOps Toolkit

Isolation via Confidential computing

Hardware enforced isolation for data in memory

Any malicious insiders and software across the stack could have access to data in use

Without Confidential Computing




VM Security enhanced with Intel TDX



Intel TDX limits what you have to trust

App security enhanced with Intel SGX



Intel SGX has the least software in trust boundary

 Attack Surface

It starts with Intel...

Industry Enhancements

Angstrom Era



PowerVIA
"Backside Power"

RibbonFET
"Faster Transistor Switching Speeds"

ASML 2.0
"EUV Light"

Glass Substrate by 2030


5G Network/ Advanced Communications



vRAN 3.0



IoT/Edge




Deep Link




Cloud Computing

Confidential Compute
"Project Amber"




HPC and AI/ML



≥2 Exaflop
"Project Aurora"

Foundry Services



Multi-Chip Packaging (MCP)



Ohio in 2025!

Software



Fakecatcher

Client Computing



Jason Dunn-Potter
Chief Warrant Officer 5, U.S. Army (ret)
Solution Architect
Defense / National Security | Intel Corp.
210-868-2243 (mobile)
jason.dunn-potter@intel.com



intel®