# Performance:

Current and coming systems

# plan

Discuss currently installed hardware

Simpler theoretical numbers (cores, clocks, memory) and microbenchmarks

Consider a simple but real application benchmark

Some guesses about next-generation hardware

# AMD Glamour Shot

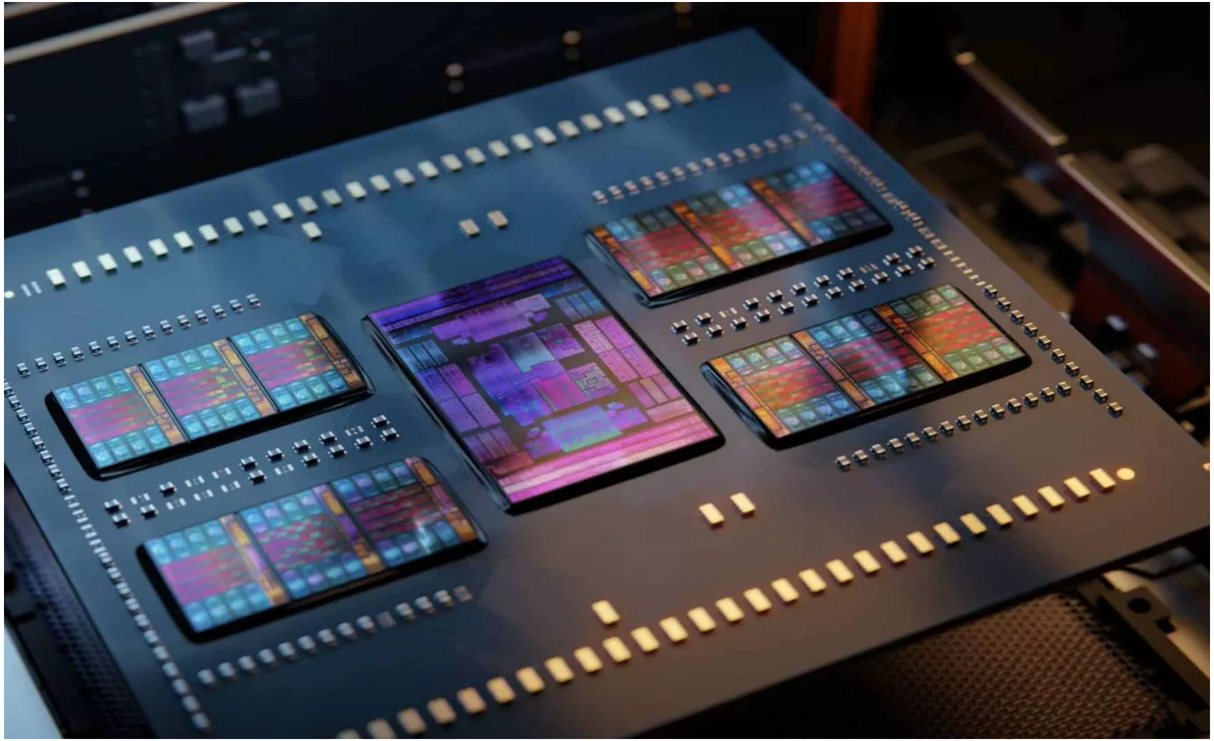

# Intel Glamour Shot

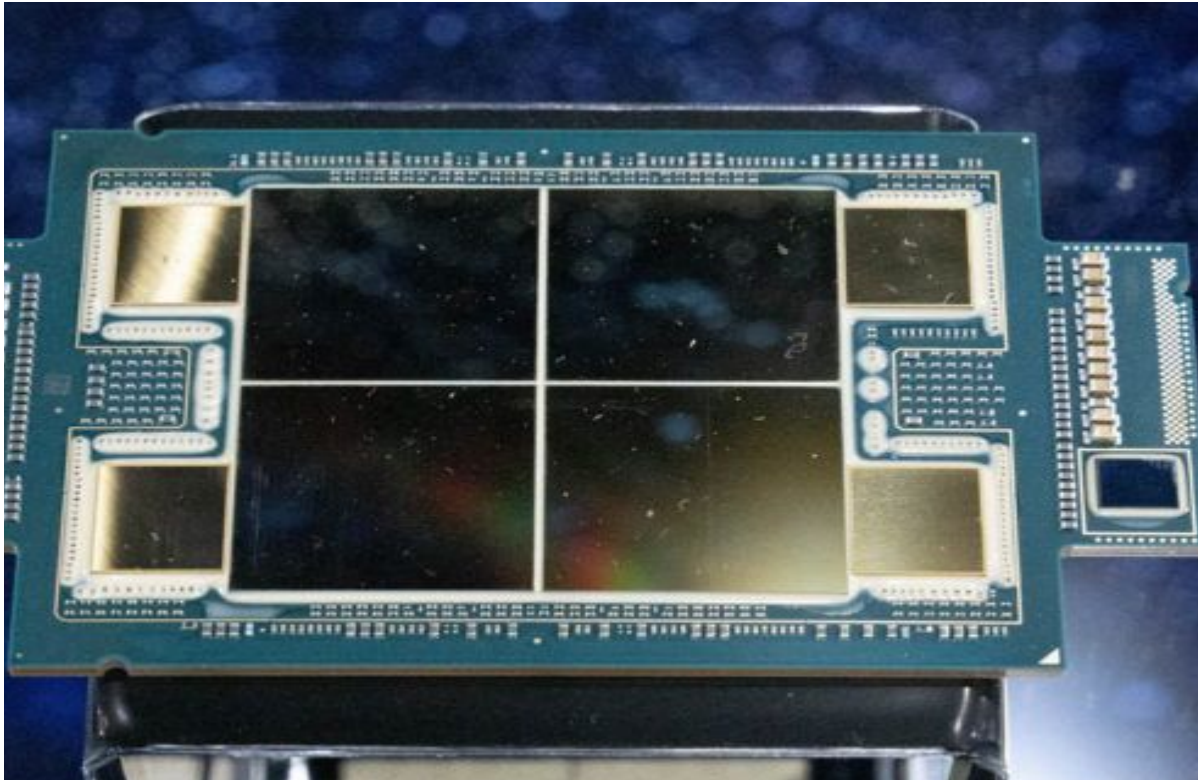

# Survey of currently used HW

| Cluster | model | codename | cores | clock | L3 cache | mem bw |
|---|---|---|---|---|---|---|
| Gra,Cdr | E5-2683 v4 | Broadwell | 16 | 2.2-2.9 | 40M | 76.8 |
| Cdr | Ptnm 8160 | Skylake | 24 | 2.1-3.7 | 33M | 128 |
| Nia | Gold 6148 | Skylake | 20 | 2.4-3.7 | 27.5M | 128 |
| Gra | Gold 6238 | Cascade | 22 | 2.1-3.7 | 30M | 140.7 |
| Cdr | Gold 6260 | Cascade | 24 | 2.1-3.7 | 30M | 140.7 |
| Nvl | 7532 | Rome | 32 | 2.4-3.3 GHz | 256M | 204.8 |
| Nvl | 7413 | Milan | 24 | 2.65-3.6 | 128M | 204.8 |

# More about scaling

| codename | cores | clock | L3 cache | mem bw | l3/core | bw/core | bw/(core*clock) |
|---|---|---|---|---|---|---|---|
| Broadwell | 16 | 2.2-2.9 | 40M | 76.8 | 2.5 | 4.8 | 2.2 |
| Skylake | 24 | 2.1-3.7 | 33M | 128 | 1.4 | 5.3 | 2.5 |
| Cascade | 22 | 2.1-3.7 | 30M | 140.7 | 1.4 | 6.4 | 3.0 |
| Rome | 32 | 2.4-3.3 | 256M | 204.8 | 8 | 6.4 | 2.7 |

# Some of why

| codename | date | cores | clock | L3 cache | Core process | power |
|---|---|---|---|---|---|---|
| Broadwell | 1q16 | 16 | 2.2-2.9 | 40M | 14 | 120 |
| Skylake | 3q17 | 24 | 2.1-3.7 | 33M | 14 | 140 |
| Cascade | 2q19 | 22 | 2.1-3.7 | 30M | 14 | 140 |
| Rome | 1q20 | 32 | 2.4-3.3 | 256M | 7 | 200 |

# More about scaling

| codename | theoretical mem bw | theoretical bw/core | Thrash (memset) | Stream triad | triad/core |
|---|---|---|---|---|---|
| Broadwell | 76.8 | 4.8 | 5.6 | 94 | 2.9 |
| Skylake | 128 | 5.3 | 6.7 | 140 | 2.9 |
| Cascade | 140.7 | 6.4 | 6.7 | 150 | 3.4 |
| Rome | 204.8 | 6.4 | 14 | 217 | 3.4 |

# More about scaling

| codename | theoretical mem bw | theoretical bw/core | cores | Thrash (memset) | Stream triad dual socket | triad/core |
|---|---|---|---|---|---|---|
| Broadwell | 76.8 | 4.8 | 16 | 5.6 | 94 | 2.9 |
| Skylake | 128 | 5.3 | 24 | 6.7 | 140 | 2.9 |
| Cascade | 140.7 | 6.4 | 22 | 6.7 | 150 | 3.4 |
| Rome | 204.8 | 6.4 | 32 | 14 | 217 | 3.4 |
| Intel next | 307.2 | 5.5 | 56 | ? | ? | ? |
| Intel next HBM | 1220 | 22 | 56 | ? | ? | ? |
| AMD next | 460.8 | 5.8 | 96 | ? | 2.3x Rome? | ? |

# Application benchmark

Standard GROMACS benchMEM

Run on a single whole-node

Run from 1..ncpus threads

Caveats:

- No MPI
- No competition, so low-thread performance is overestimated

# Elapsed time

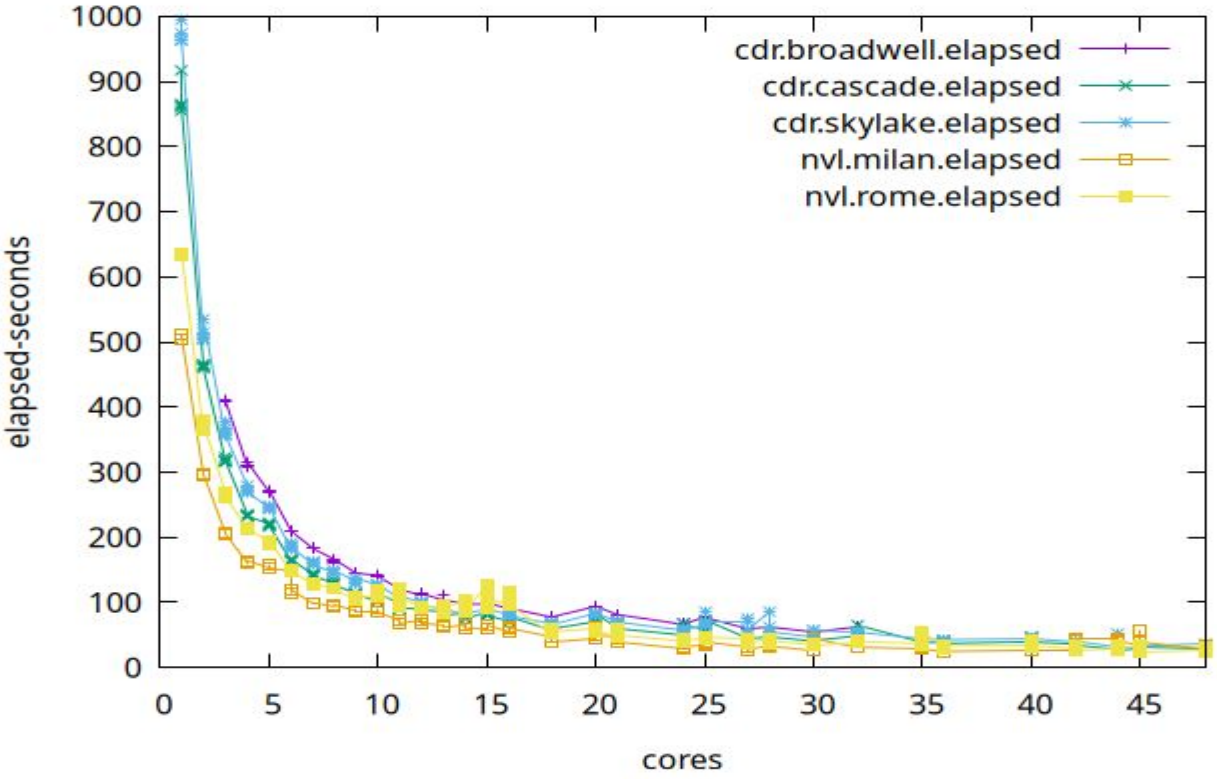

# Core-seconds

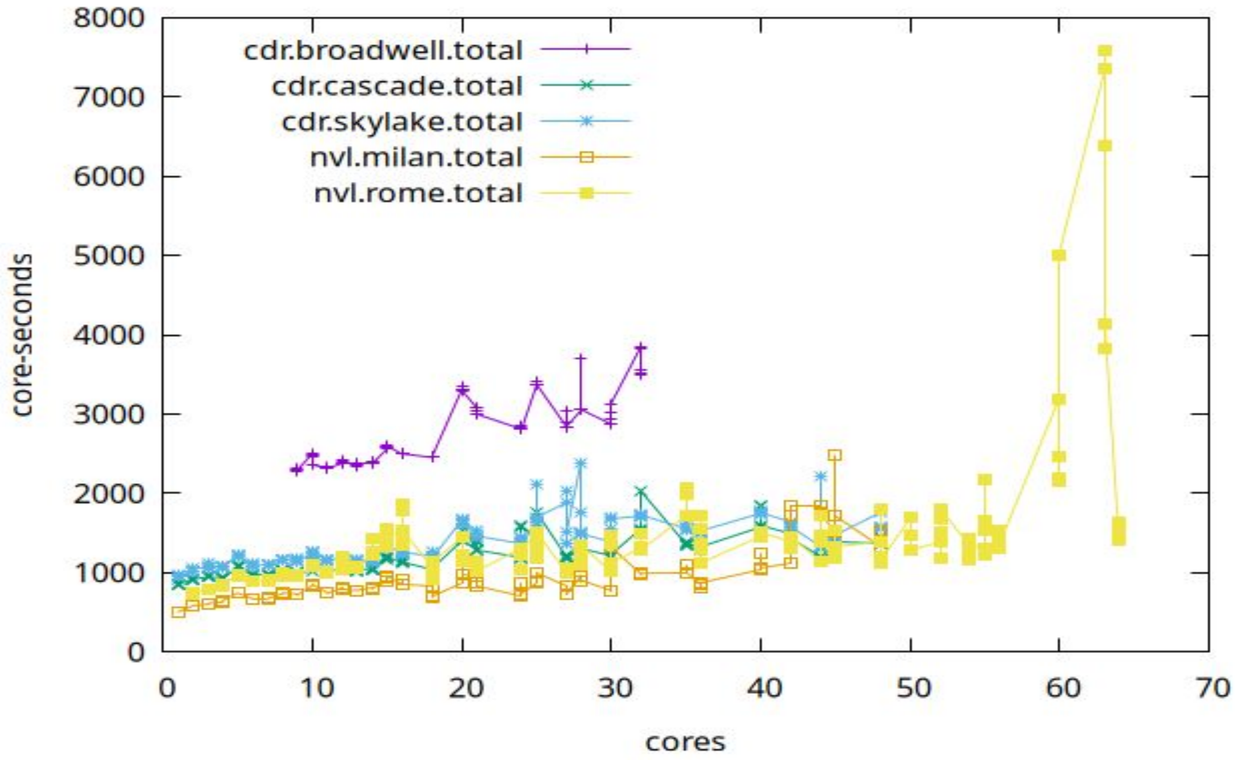

# Where are we

Current clusters were commissioned about 6 years ago.

There have been some upgrades since then.

We are on the verge of some important new changes.

Systems continue to be torn between CPU and GPU.

Also torn between compact and memory-intensive.

# Our CPUs haven’t changed a lot recently

Intel has been stuck

- Minor changes in core density
- Minor changes in memory clock
- Minor improvements in instruction set (eg avx512)
- Modest improvements in power/performance

AMD has reinvented itself

- Serious advantages from TSMC
- Chiplets
- Advanced memory

# AMD advantages

Consuming 5 nm fab capacity from TSMC

- Compared to 10nm from Intel
- Not literally a 2x difference (4x density)
- Intel bet against EUV, now playing catchup

Chiplets are a significant advantage

- Very large chips are a yield problem
- Ability to leverage multiple processes
- Stacking HBM and cache

# Intel definitely not out of the game

| 4th Gen Intel Xeon Scalable Launch SKUs | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Family | Model | Cores | Base (GHz) | All Core Turbo (GHz) | Max Turbo (GHz) | Cache (MB) | TDP (W) | Max Sockets | Max DDR5 | UPI | DSA | QAT | DLB | IAA | SGX Max (GB) | Price | On Demand | Die |
| Platinum | 8480+ | 56 | 2.0 | 3.0 | 3.8 | 105.0 | 350 | 2 | 4800 | 4 | 1 | 1 | 1 | 1 | 512 | 10710 | Yes | XCC |
| Platinum | 8470 | 52 | 2.0 | 3.0 | 3.8 | 105.0 | 350 | 2 | 4800 | 4 | 1 | 0 | 0 | 0 | 512 | 9359 | Yes | XCC |
| Platinum | 8468 | 48 | 2.1 | 3.1 | 3.8 | 105.0 | 350 | 2 | 4800 | 4 | 1 | 0 | 0 | 0 | 512 | 7214 | Yes | XCC |
| Platinum | 8460Y+ | 40 | 2.0 | 2.8 | 3.7 | 105.0 | 300 | 2 | 4800 | 4 | 1 | 1 | 1 | 1 | 128 | 5558 | Yes | XCC |
| Platinum | 8462Y+ | 32 | 2.8 | 3.6 | 4.1 | 60.0 | 300 | 2 | 4800 | 3 | 1 | 1 | 1 | 1 | 128 | 5945 | Yes | XCC |
| Gold | 6448Y | 32 | 2.1 | 3.0 | 4.0 | 60.0 | 225 | 2 | 4800 | 3 | 1 | 0 | 0 | 0 | 128 | 3583 | Yes | MCC |
| Gold | 6442Y | 24 | 2.6 | 3.3 | 4.1 | 60.0 | 225 | 2 | 4800 | 3 | 1 | 0 | 0 | 0 | 128 | 2878 | Yes | MCC |
| Gold | 6444Y | 16 | 3.6 | 4.0 | 4.1 | 45.0 | 270 | 2 | 4800 | 3 | 1 | 0 | 0 | 0 | 128 | 3622 | Yes | MCC |
| Gold | 6426Y | 16 | 2.5 | 3.3 | 4.1 | 37.5 | 185 | 2 | 4800 | 3 | 1 | 0 | 0 | 0 | 128 | 1517 | Yes | MCC |
| Gold | 6434 | 8 | 3.7 | 4.1 | 4.1 | 22.5 | 195 | 2 | 4800 | 3 | 1 | 0 | 0 | 0 | 128 | 2607 | Yes | MCC |
| Gold | 5415+ | 8 | 2.9 | 3.6 | 4.1 | 22.5 | 150 | 2 | 4800 | 3 | 1 | 1 | 1 | 1 | 128 | 1066 | Yes | MCC |
| Platinum | 8452Y | 36 | 2.0 | 2.8 | 3.2 | 67.5 | 300 | 2 | 4800 | 4 | 1 | 0 | 0 | 0 | 128 | 3995 | Yes | MCC |
| Gold | 6438Y+ | 32 | 2.1 | 2.8 | 4.0 | 60.0 | 205 | 2 | 4800 | 3 | 1 | 1 | 1 | 1 | 128 | 3141 | Yes | MCC |
| Gold | 6430 | 32 | 2.0 | 2.6 | 3.4 | 60.0 | 270 | 2 | 4400 | 3 | 1 | 0 | 0 | 0 | 128 | 2128 | Yes | MCC |
| Gold | 5420+ | 28 | 2.0 | 2.7 | 4.1 | 52.5 | 205 | 2 | 4400 | 3 | 1 | 1 | 1 | 1 | 128 | 1848 | Yes | MCC |
| Gold | 5418Y | 24 | 2.0 | 2.8 | 3.8 | 45.0 | 185 | 2 | 4400 | 3 | 1 | 0 | 0 | 0 | 128 | 1483 | Yes | MCC |
| Silver | 4416+ | 20 | 2.0 | 2.9 | 3.9 | 37.5 | 165 | 2 | 4400 | 2 | 1 | 1 | 1 | 1 | 64 | 1176 | Yes | MCC |
| Silver | 4410Y | 12 | 2.0 | 2.8 | 3.9 | 30.0 | 150 | 2 | 4400 | 2 | 1 | 0 | 0 | 0 | 64 | 563 | Yes | MCC |
| Platinum | 8470Q | 52 | 2.1 | 3.2 | 3.8 | 105.0 | 350 | 2 | 4800 | 4 | 1 | 0 | 0 | 0 | 512 | 9410 | Yes | XCC |
| Gold | 6458Q | 32 | 3.1 | 4.0 | 4.0 | 60.0 | 350 | 2 | 4800 | 3 | 1 | 0 | 0 | 0 | 128 | 6416 | Yes | MCC |
| Gold | 6414U | 32 | 2.0 | 2.6 | 3.4 | 60.0 | 250 | 1 | 4800 | 0 | 1 | 0 | 0 | 0 | 512 | 2296 | Yes | XCC |
| Gold | 5412U | 24 | 2.1 | 2.9 | 3.9 | 45.0 | 185 | 1 | 4400 | 0 | 1 | 0 | 0 | 0 | 128 | 1113 | Yes | MCC |
| Bronze | 3408U | 8 | 1.8 | 1.9 | 1.9 | 22.5 | 125 | 1 | 4000 | 0 | 1 | 0 | 0 | 0 | 64 | 415 | No | MCC |
| Silver | 4410T | 10 | 2.7 | 3.4 | 4.0 | 26.3 | 150 | 2 | 4000 | 2 | 1 | 0 | 0 | 0 | 64 | 624 | Yes | MCC |
| Platinum | 8490H | 60 | 1.9 | 2.9 | 3.5 | 112.5 | 350 | 8 | 4800 | 4 | 4 | 4 | 4 | 4 | 512 | 17000 | No | XCC |
| Platinum | 8468H | 48 | 2.1 | 3.0 | 3.8 | 105.0 | 330 | 8 | 4800 | 4 | 4 | 4 | 4 | 4 | 512 | 13923 | No | XCC |
| Platinum | 8460H | 40 | 2.2 | 3.1 | 3.8 | 105.0 | 330 | 8 | 4800 | 4 | 4 | 0 | 0 | 4 | 512 | 10710 | No | XCC |
| Platinum | 8454H | 32 | 2.1 | 2.7 | 3.4 | 82.5 | 270 | 8 | 4800 | 4 | 4 | 4 | 4 | 4 | 512 | 6540 | No | XCC |
| Platinum | 8450H | 28 | 2.0 | 2.6 | 3.5 | 75.0 | 250 | 8 | 4800 | 4 | 4 | 0 | 0 | 4 | 512 | 4708 | No | XCC |
| Platinum | 8444H | 16 | 2.9 | 3.2 | 4.0 | 45.0 | 270 | 8 | 4800 | 4 | 4 | 0 | 0 | 4 | 512 | 4234 | No | XCC |
| Gold | 6448H | 32 | 2.4 | 3.2 | 4.1 | 60.0 | 250 | 4 | 4800 | 3 | 1 | 2 | 2 | 1 | 512 | 3658 | No | MCC |
| Gold | 6418H | 24 | 2.1 | 2.9 | 4.0 | 60.0 | 185 | 4 | 4800 | 3 | 1 | 0 | 0 | 1 | 512 | 2065 | No | MCC |
| Gold | 6416H | 18 | 2.2 | 2.9 | 4.2 | 45.0 | 165 | 4 | 4800 | 3 | 1 | 0 | 0 | 1 | 512 | 1444 | No | MCC |
| Gold | 6434H | 8 | 3.7 | 4.1 | 4.1 | 22.5 | 195 | 4 | 4800 | 3 | 1 | 0 | 0 | 1 | 512 | 3070 | No | MCC |
| Platinum | 8470N | 52 | 1.7 | 2.7 | 3.6 | 97.5 | 300 | 2 | 4800 | 4 | 4 | 4 | 4 | 0 | 128 | 9520 | Yes | XCC |
| Platinum | 8471N | 52 | 1.8 | 2.8 | 3.6 | 97.5 | 300 | 1 | 4800 | 4 | 4 | 4 | 4 | 0 | 128 | 5171 | Yes | XCC |
| Gold | 6438N | 32 | 2.0 | 2.7 | 3.6 | 60.0 | 205 | 2 | 4800 | 3 | 1 | 2 | 2 | 0 | 128 | 3351 | Yes | MCC |
| Gold | 6428N | 32 | 1.8 | 2.5 | 3.8 | 60.0 | 185 | 2 | 4000 | 3 | 1 | 2 | 2 | 0 | 128 | 3200 | Yes | MCC |
| Gold | 6421N | 32 | 1.8 | 2.6 | 3.6 | 60.0 | 185 | 1 | 4400 | 3 | 1 | 2 | 2 | 0 | 128 | 2368 | Yes | MCC |
| Gold | 5418N | 24 | 1.8 | 2.6 | 3.8 | 45.0 | 165 | 2 | 4000 | 3 | 1 | 2 | 2 | 0 | 128 | 1664 | Yes | MCC |
| Gold | 5411N | 24 | 1.9 | 2.8 | 3.9 | 45.0 | 165 | 1 | 4400 | 3 | 1 | 2 | 2 | 0 | 128 | 1232 | Yes | MCC |
| Platinum | 8468V | 48 | 2.4 | 2.9 | 3.8 | 97.5 | 330 | 2 | 4800 | 3 | 1 | 1 | 1 | 1 | 128 | 7121 | Yes | XCC |
| Platinum | 8458P | 44 | 2.7 | 3.2 | 3.8 | 82.5 | 350 | 2 | 4800 | 3 | 1 | 1 | 1 | 1 | 512 | 6759 | Yes | XCC |
| Platinum | 8461V | 48 | 2.2 | 2.8 | 3.7 | 97.5 | 300 | 1 | 4800 | 0 | 1 | 1 | 1 | 1 | 128 | 4491 | Yes | XCC |
| Gold | 6438M | 32 | 2.2 | 2.8 | 3.9 | 60.0 | 205 | 2 | 4800 | 3 | 1 | 0 | 0 | 1 | 128 | 3273 | Yes | MCC |
| Gold | 6454S | 32 | 2.2 | 2.8 | 3.4 | 60.0 | 270 | 2 | 4800 | 4 | 4 | 4 | 4 | 0 | 128 | 3157 | Yes | XCC |
| Gold | 5416S | 16 | 2.0 | 2.8 | 4.0 | 30.0 | 150 | 2 | 4400 | 3 | 1 | 2 | 2 | 0 | 128 | 944 | Yes | MCC |
| Max | 9480 | 56 | 1.9 | 2.6 | 3.5 | 112.5 | 350 | 2 | 4800 | 4 | 4 | 0 | 0 | 0 | 512 | 12980 | No | XCC |
| Max | 9470 | 52 | 2.0 | 2.7 | 3.5 | 105.0 | 350 | 2 | 4800 | 4 | 4 | 0 | 0 | 0 | 512 | 11590 | No | XCC |
| Max | 9468 | 48 | 2.1 | 2.6 | 3.5 | 105.0 | 350 | 2 | 4800 | 4 | 4 | 0 | 0 | 0 | 512 | 9900 | No | XCC |
| Max | 9460 | 40 | 2.2 | 2.7 | 3.5 | 97.5 | 350 | 2 | 4800 | 3 | 4 | 0 | 0 | 0 | 128 | 8750 | No | XCC |
| Max | 9462 | 32 | 2.7 | 3.1 | 3.5 | 75.0 | 350 | 2 | 4800 | 3 | 4 | 0 | 0 | 0 | 128 | 7995 | No | XCC |

# Intel Sapphire Rapids vs AMD Genoa

- Intel: Up to 56 cores, 112.4M cache, 350W TDP 2.0 / 3.0 / 3.8 GHz (base, allcores, fewcores)
- Intel “Max”: HBM version similar but lower clock: 2.7 / 3.1 / 3.5, but 20% higher price
- AMD: Genoa up to 96 cores, 384MB cache, 360W, 2.4 - 3.7
- Genoa-X maybe 1152MB stacked cache
- Bergamo (Zen4c): up to 128c/socket (probably reduced L3)

# How’s our timing?

Let’s assume we’re close to getting federal approval

Then we need provincial match

Eventually there would be some call to vendors (RFP, etc)

These have taken 1-2 years in the past (writing RFP to running jobs)

Perhaps we can expect less delay since we’re not starting from scratch

Still hard to tell whether Genoa and Sapphire Rapids are the target

# Intel cores


I-TLB + I-Cache
Predict
MSROM
Decode
µop Cache
µop Queue
Allocate / Rename / Move Elimination / Zero Idiom
Scheduler
INT
VEC
Store Data
48KB Data Cache
1.25MB/2MB ML Cache
New
Performance
x86 Core
A Step Function in CPU Architecture Performance For the Next Decade of Compute
A significant IPC boost at high power efficiency
Wider
Deeper
Smarter
Better supports large data set and large code footprint applications
Enhanced power management improves frequency and power
Machine Learning Technology: Intel® AMX – Tile Multiplication
All in a tailored scalable architecture to serve the full range of Laptops to Desktops to Data Centers
Under Embargo until November 9, 2022, 6 am PT
intel

# Zen4 architecture

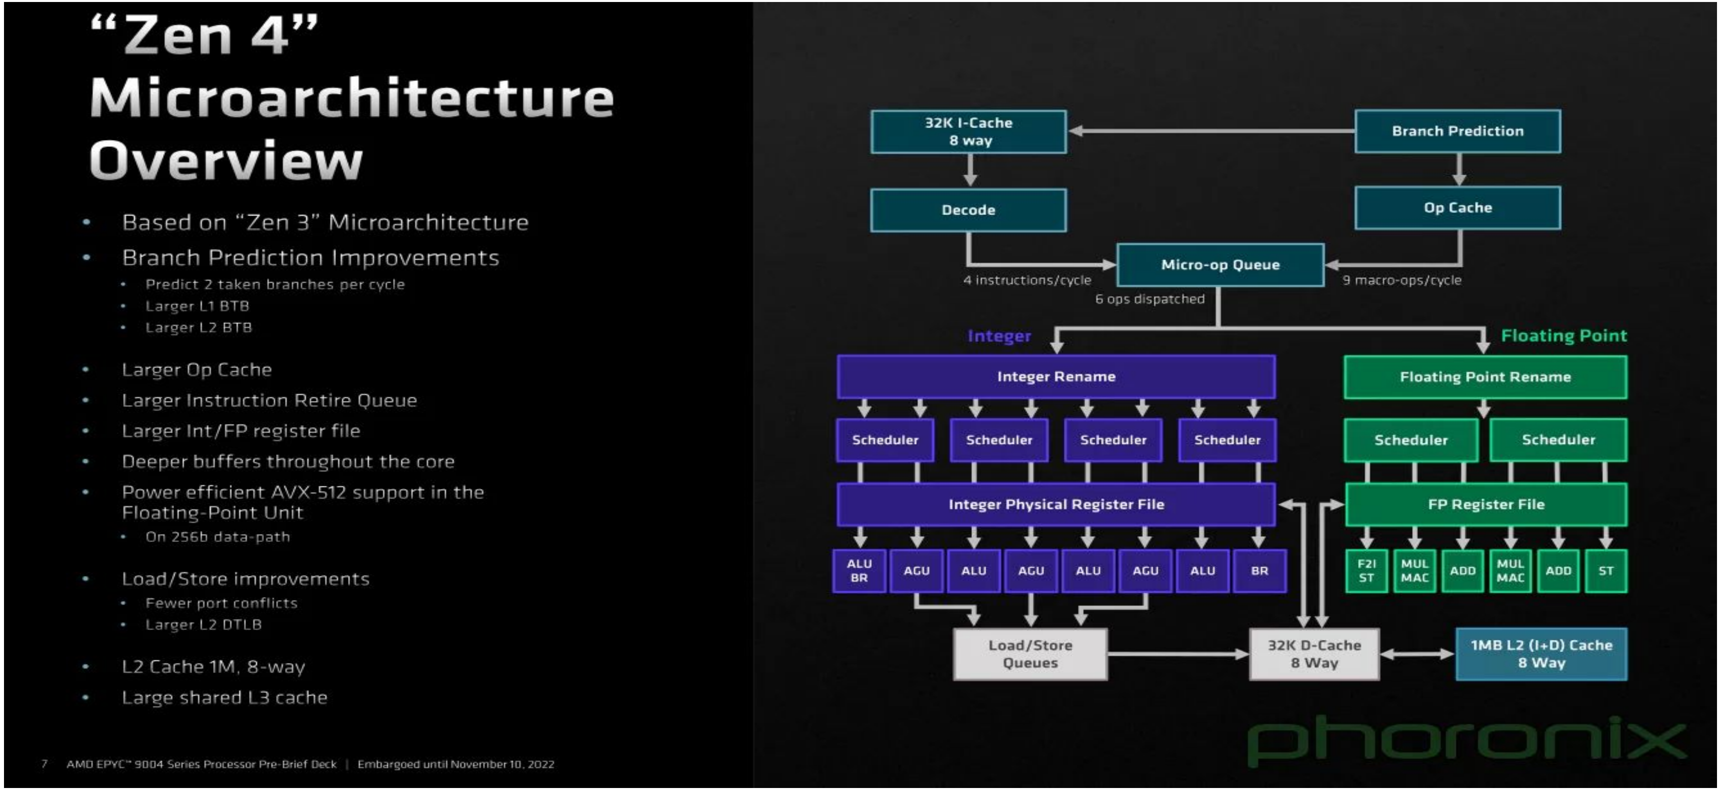

# Does "general purpose" exist anymore?

- We have quite a few GPU consumers
- We also have quote a few low-memory users (100M/core)
- We also have at least some very high memory users
  - Big shared-memory/parallel users
  - Unruly/dynamic pipelines
- Do we need some non-MPI-tuned clusters?
  - 256 cores/node oughta be good enough for anyone!
  - Different world from our original 32c/node installs
- Still need IO
  - Only some users care about node-local though

# Face off

https://wccftech.com/intel-4th-gen-xeon-cpus-official-sapphire-rapids-up-to-60-cores-8-socket-scalability-350w-tdp-17000-usd/

https://www.tomshardware.com/news/intel-sapphire-rapids-with-hbm-is-2x-faster-than-amds-milan-x

Hbm part about 1.2 GB/s for 64g vs 307.2 for 8x8x4.8

307.2/56=5.5

Amd 12x8x5.2=500

5.2 gb/s/core (96)

# How timing effects product

We’re not going to buy current products (AMD Milan, Intel Sapphire Ridge introduced this week).

- Intel in-package HBM?
  - But AMD has 12-channel memory
- AMD stacked cache (-X chips)
  - Which workloads benefit?
- Emphasis on cores?  (AMD 128c/package)
  - Advantage to more scalable codes?