



Ampere™ Altra™ Processor

Datacenter Efficiency

Cloud computing requirements have dramatically diverged from legacy enterprise requirements

Shift to
containers and
microservices

Growth in
Edge
data centers


Demand for
density and
power efficiency

Increased
capabilities of
ML and AI

Customized
heterogenous
infrastructures

Current solutions do not meet Cloud needs

Performance and Power Efficiency CAGR have slowed




Refer to end notes for details

The Cloud needs High Performance and Power Efficient Processors

Embargo: March 3, 2020 (6:00 AM Pacific time)

Power Consumption and Sustainability are critical global issues



Data center power consumption as percentage of worldwide energy consumption continues to grow



Data center contribution to worldwide greenhouse gas emissions already equals that of the airline industry

Power Efficient processors are required for the Cloud

The requirements for the modern Cloud



Predictable and
High Performance

Scalability across
the platform

Power Efficiency
and Density

Ampere™ Altra™ processor complex

80 64-bit Arm CPU cores @ 3.0 GHz Turbo

- 4-Wide superscalar aggressive out-of-order execution
- Single threaded cores for performance and security isolation


Arm v8.2+ features

Large Cache, all with ECC Protection

- 64 KB L1 I/D-cache per core
- 1 MB L2 cache per core
- 32 MB system level cache

2x 128-bit SIMD Units

int8 and fp16 for ML Inference performance

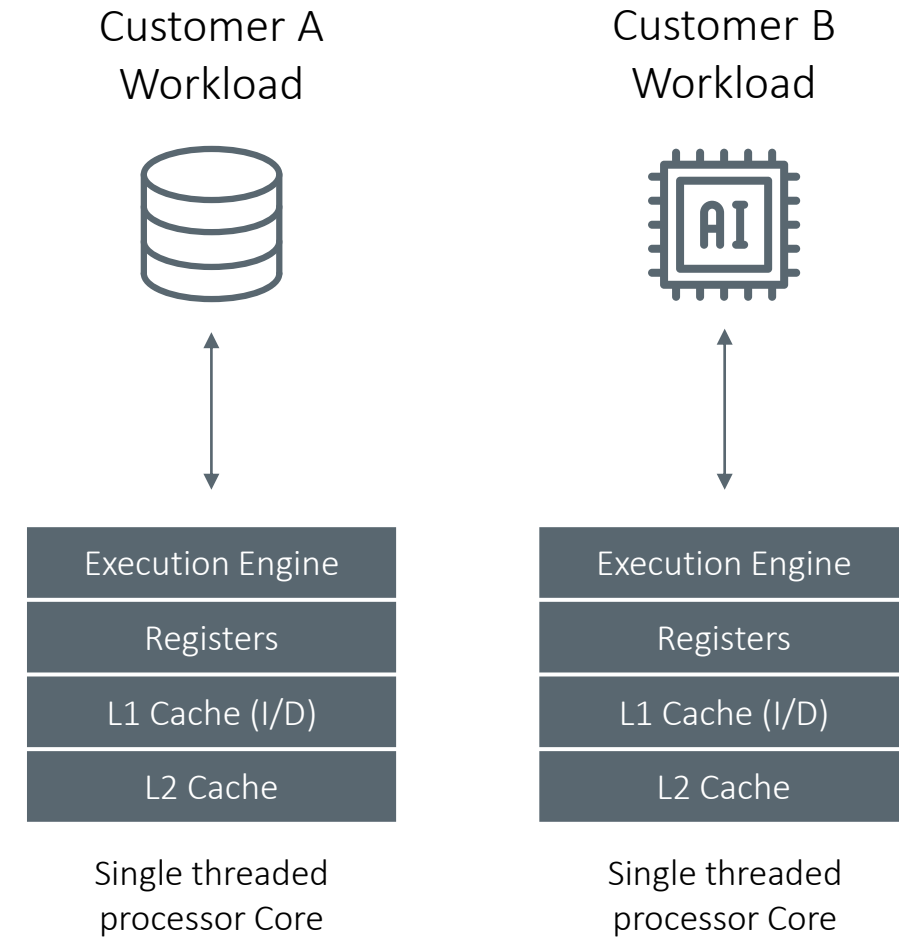


Embargo: March 3, 2020 (6:00 AM Pacific time)

Processor cores built to reduce “noisy neighbor” effects

Dedicated resources per core for:

- Excellent customer isolation
- Lower vulnerability to side channel attacks
- Consistent, predictable performance



Processor built from the ground up for the Cloud

Single threaded cores deliver for the cloud specific needs:

- Consistent predictable performance
- No *noisy neighbor* issues within core
- No resource contention within core

Instruction Execution Illustration

Single threaded core



Cache miss

Branch mis-predict

10 Clocks

Multithreaded core



Cache miss

Branch mis-predict

15 Clocks



Leadership Performance

Run Cloud workloads fastest on the Ampere™ Altra™ Processor:

- High Performance for Cloud Native workloads
- Predictable virtual machine and container performance



Search



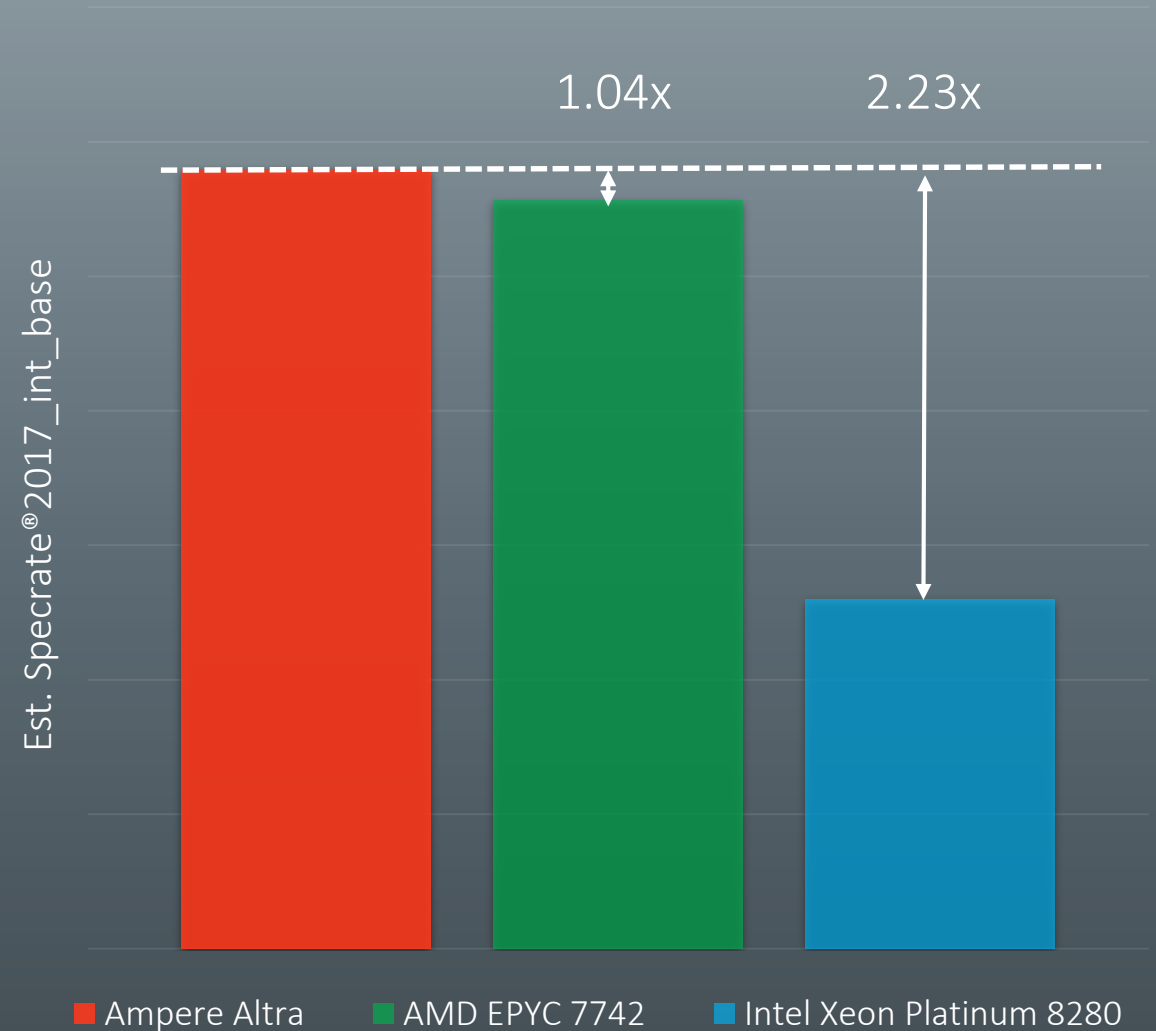
AI Inference



Media Transcoding



Database



Leadership Power Efficiency

Run Cloud workloads on the most power efficient Ampere™ Altra™ Processor:

- Scalability from Hyperscale to Edge
- Maximize data-center capacity within the same power envelope



Web Applications




Storage



Cloud Gaming




Edge Cloud



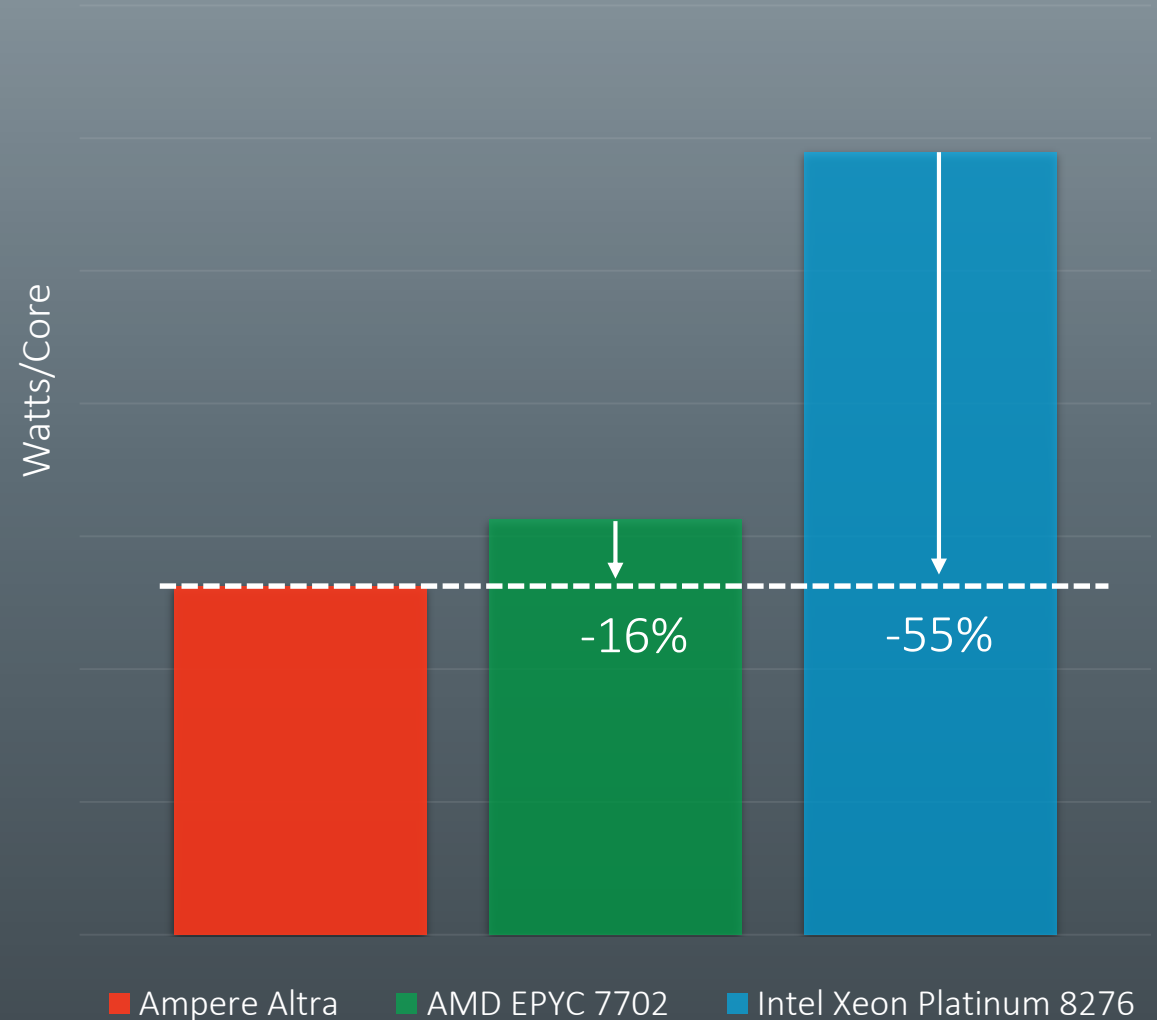
Leadership Scalability

Run Cloud workloads on the densest racks with the Ampere™ Altra™ Processor:

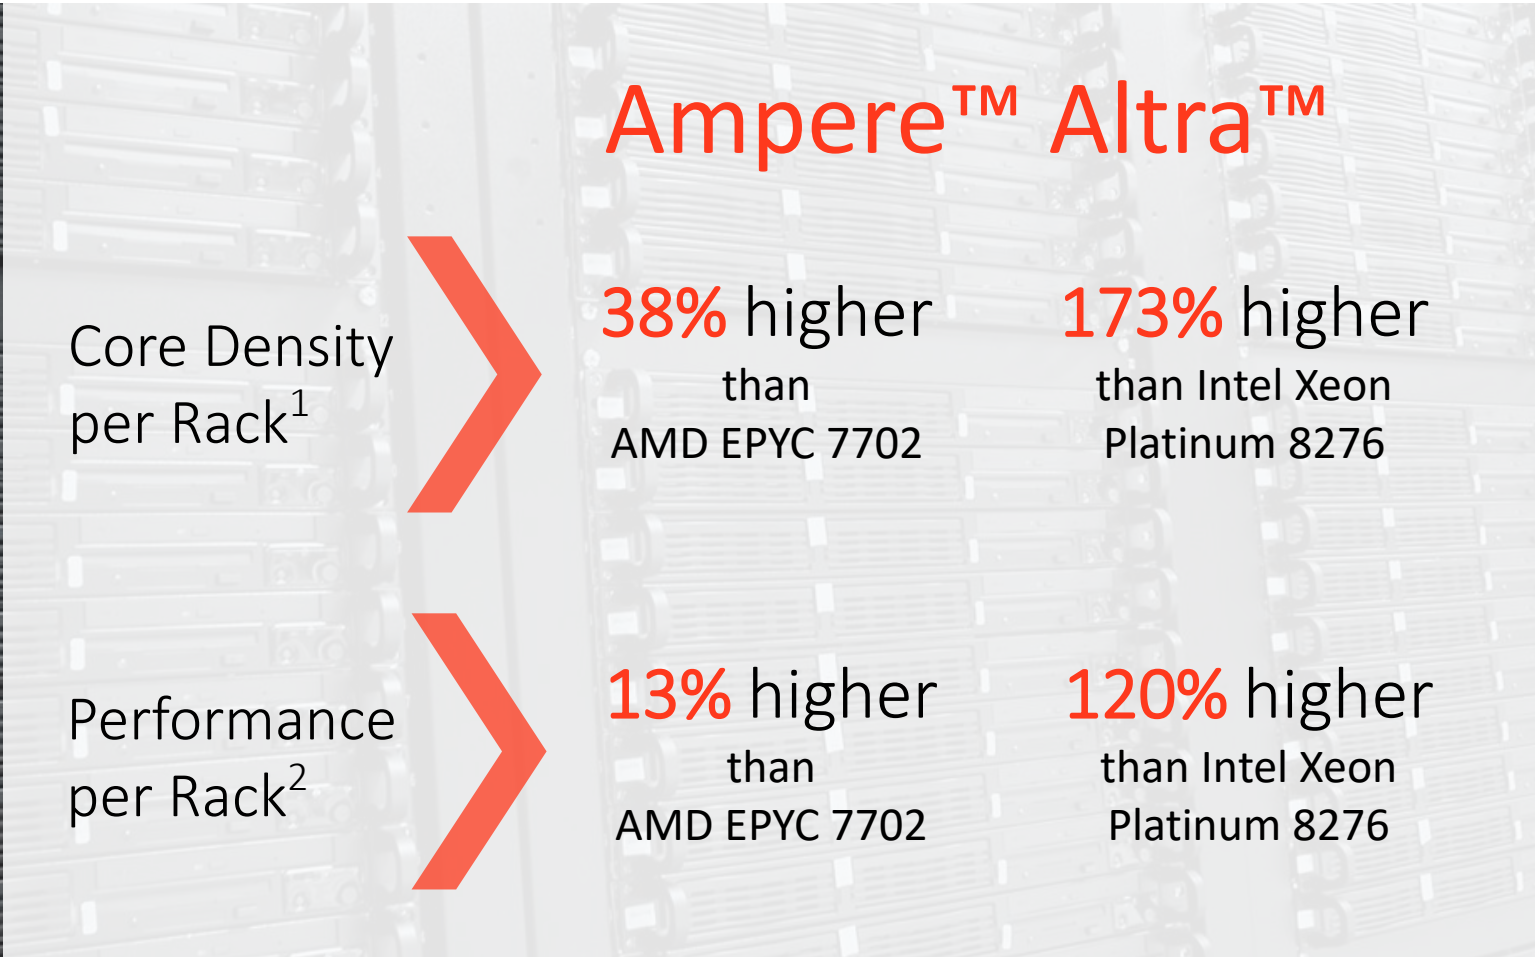
- 3500+ cores per Rack¹
- 11,000+ Est. SPECrate® 2017_int_base per Rack¹



¹Rack Config: 42U w/ 12.5kw rack power



Extreme density for maximum rack capacity



¹Rack Config: 42U w/ 12.5kw rack power

²Est. SPECrate®2017_int_base

End Notes

Embargo: March 3, 2020 (6:00 AM Pacific time)

