AMD Naples

EPYC Family

Alif R Rochman (alifrr2)

Robert A Ruester (ruester2)

William Sentosa (sentosa2)

Overview

- General Information
- Layout
- Security
- Memory Access
- Infinity Fabric
- Pipeline
- Caching
- Power
- Branch Prediction
- Multithreading
- Performance (vs predecessor)

General Information (1)

- Highest-performance enterprise-level server based on Zen
- Branded as EPYC 7000 series
- Released in 20 June 2017
- 14nm process
- Socket SP3 (also released in June 2017)
- Support dual socket
- Opteron 6k -> **EPYC 7k** -> Milan (2019)

Sources: [10][14][15][7]

General Information (2)

List of Naples Processors										
Model	Family	Price	Launched	Cores	Threads	TDP	L2$	L3$	Base	Turbo
7251	EPYC	$ 574.00	20 June 2017	8	16	120 W	4 MiB	32 MiB	2.1 GHz	2.9 GHz
7261	EPYC		14 June 2018	8	16	155 W, 170 W	4 MiB	64 MiB	2.5 GHz	2.9 GHz
7281	EPYC	$ 650.00	20 June 2017	16	32	155 W, 170 W	8 MiB	32 MiB	2.1 GHz	2.7 GHz
7301	EPYC	$ 825.00	20 June 2017	16	32	155 W, 170 W	8 MiB	64 MiB	2.2 GHz	2.7 GHz
7351	EPYC	$ 1,100.00	20 June 2017	16	32	155 W, 170 W	8 MiB	64 MiB	2.4 GHz	2.9 GHz
7351P	EPYC	$ 750.00	20 June 2017	16	32	155 W, 170 W	8 MiB	64 MiB	2.4 GHz	2.9 GHz
7371	EPYC		2019	16	32		8 MiB	64 MiB	3.1 GHz	3.8 GHz
7401	EPYC	$ 1,850.00	20 June 2017	24	48	155 W, 170 W	12 MiB	64 MiB	2 GHz	3 GHz
7401P	EPYC	$ 1,075.00	20 June 2017	24	48	155 W, 170 W	12 MiB	64 MiB	2 GHz	3 GHz
7451	EPYC	$ 2,400.00	20 June 2017	24	48	180 W	12 MiB	64 MiB	2.3 GHz	3.2 GHz
7501	EPYC	$ 3,400.00	20 June 2017	32	64	155 W, 170 W	16 MiB	64 MiB	2 GHz	3 GHz
7551	EPYC	$ 3,400.00	20 June 2017	32	64	180 W	16 MiB	64 MiB	2 GHz	3 GHz
7551P	EPYC	$ 2,100.00	20 June 2017	32	64	180 W	16 MiB	64 MiB	2 GHz	3 GHz
7601	EPYC	$ 4,200.00	20 June 2017	32	64	180 W	16 MiB	64 MiB	2.2 GHz	3.2 GHz

General Information (3)

- 32 cores System-On-Chip (SOC) design
- 8-channels of memory per “Naples” device.
- Support up to 16 DIMMS of DDR4 on 8 memory channels, delivering up to 2TB of memory.
- High-speed I/O supporting 128 lanes of PCIe 3
- Die size: 23nm
- Transistors: 19.2 billion+

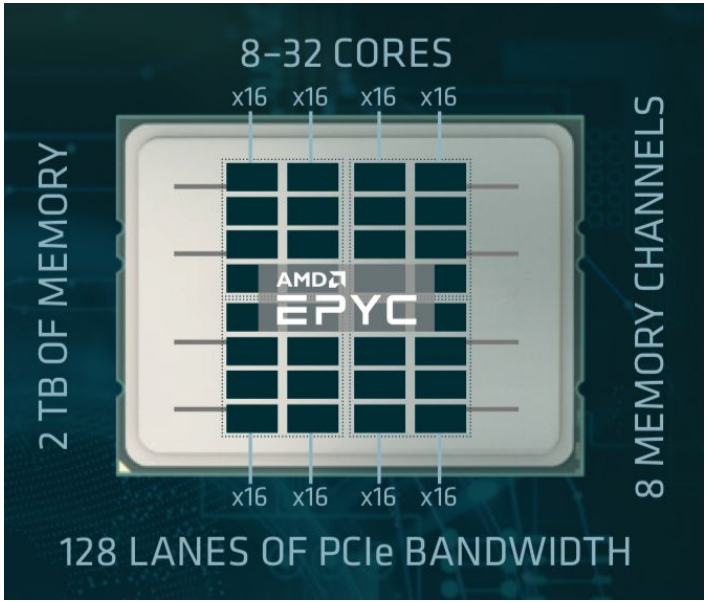

EYPC Processor Layout

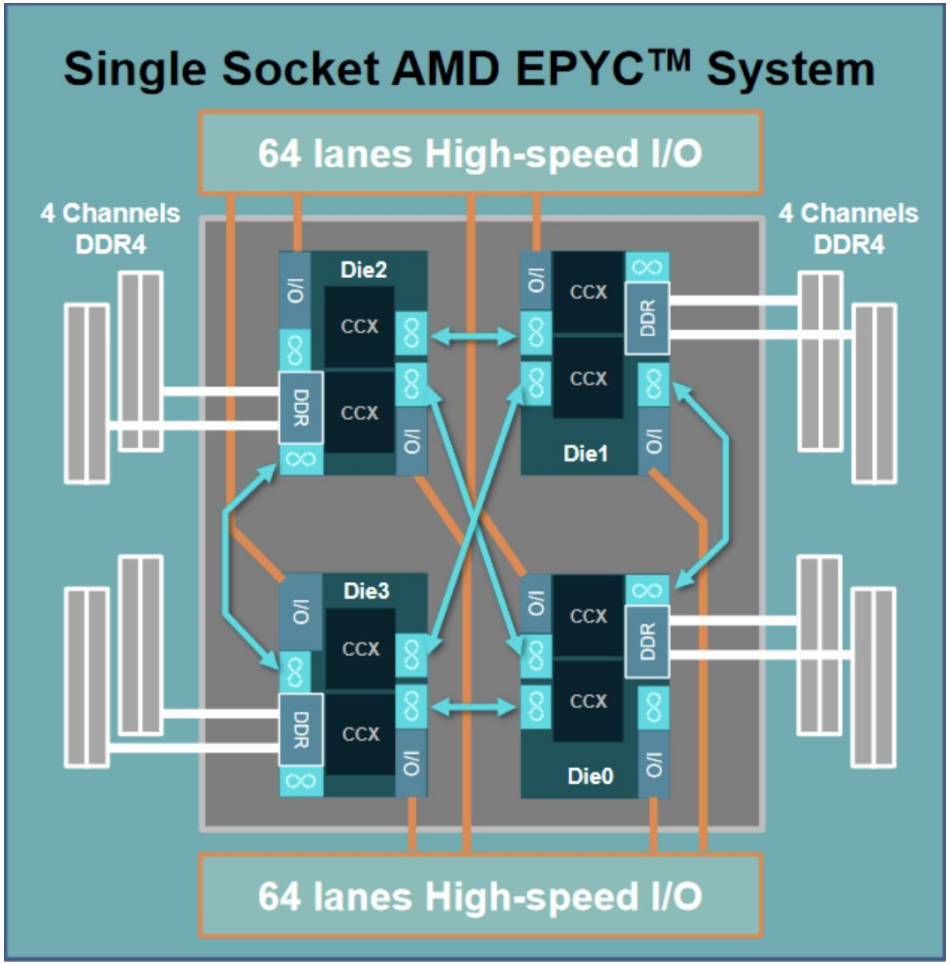

Retrieved from [1]

Multichip Module (MCM) technology used in EYPC

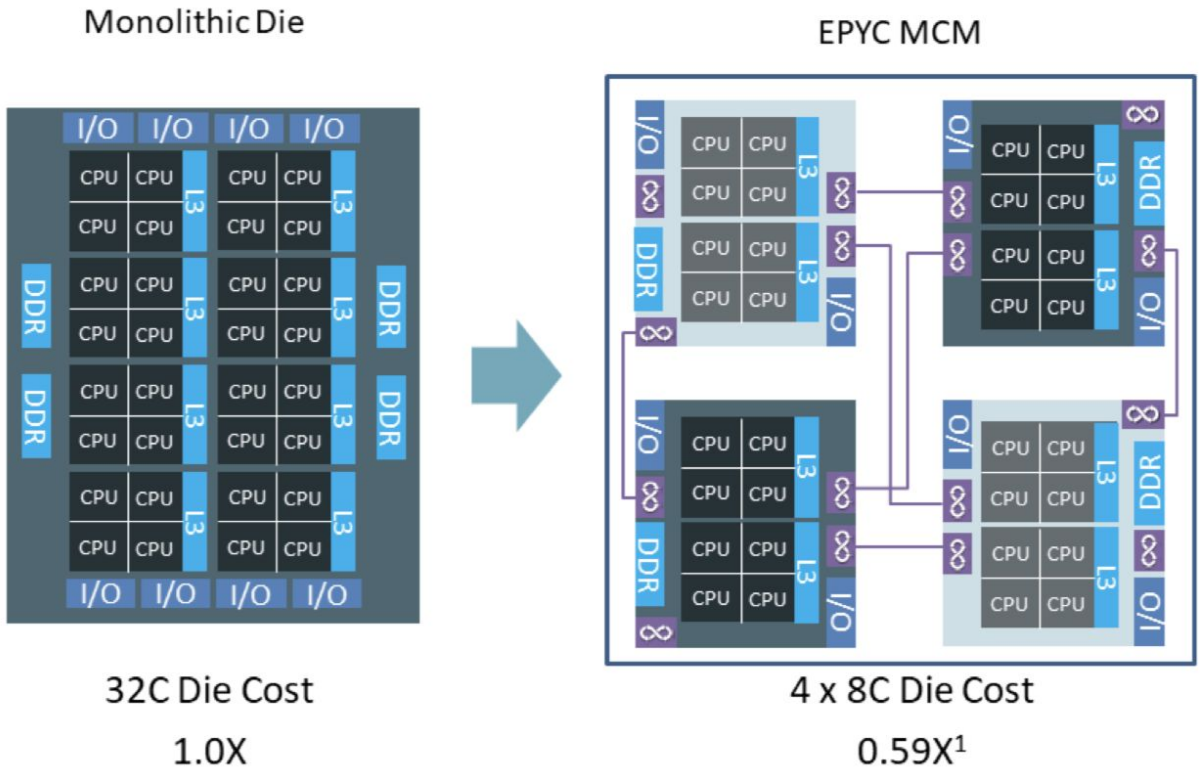

Retrieved from [2]

Security features on EYPC

- **Secure Root-of-Trust**
  AMD Secure Processor and Secure Boot
- **Secure Run**
  Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEV)
- **Secure Move**
  SEV-enabled servers, API, and third-party key management

Reference [3]

EPYC Memory

- Adopt the NUMA model.
- Infinity Fabric, as a scalable interconnect for on-die, on-package, and multi-package communication.

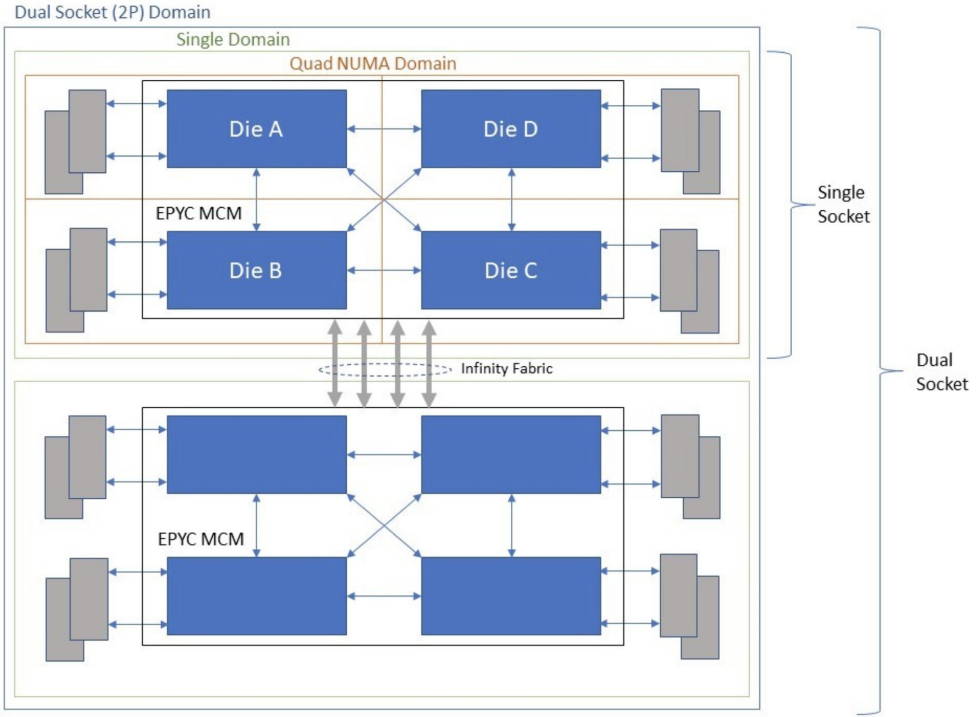

Sources: AMD & TIRIAS Research

Infinity Fabric

- High speed, low latency, power efficient
- Scalable Data Fabric connects CCX, memory, I/O
- Scalable Control Fabric connects controllers
- Core communication using MDOEFSI protocol

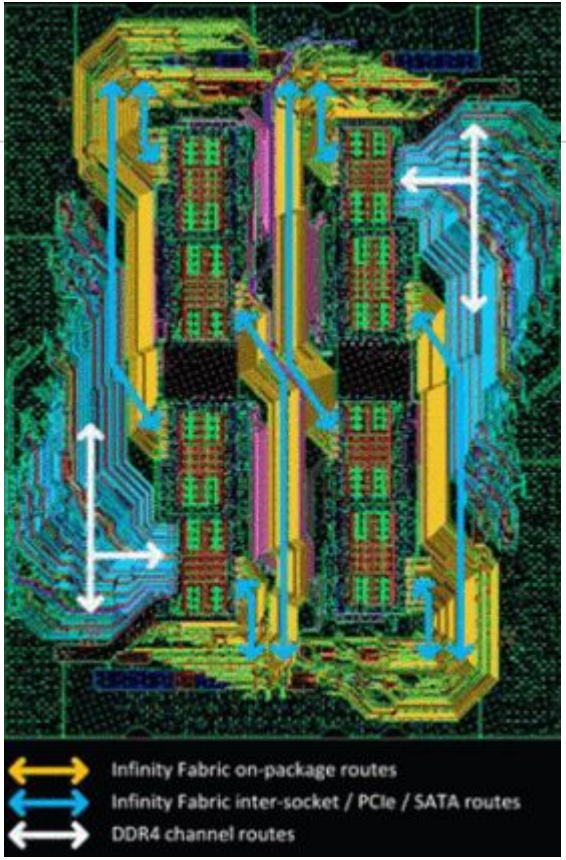

Source: [8]

Pipeline - Fetch

- Four instructions per cycle
- Instructions from instruction cache or branch prediction
- TLB translates PC for instruction prefetch

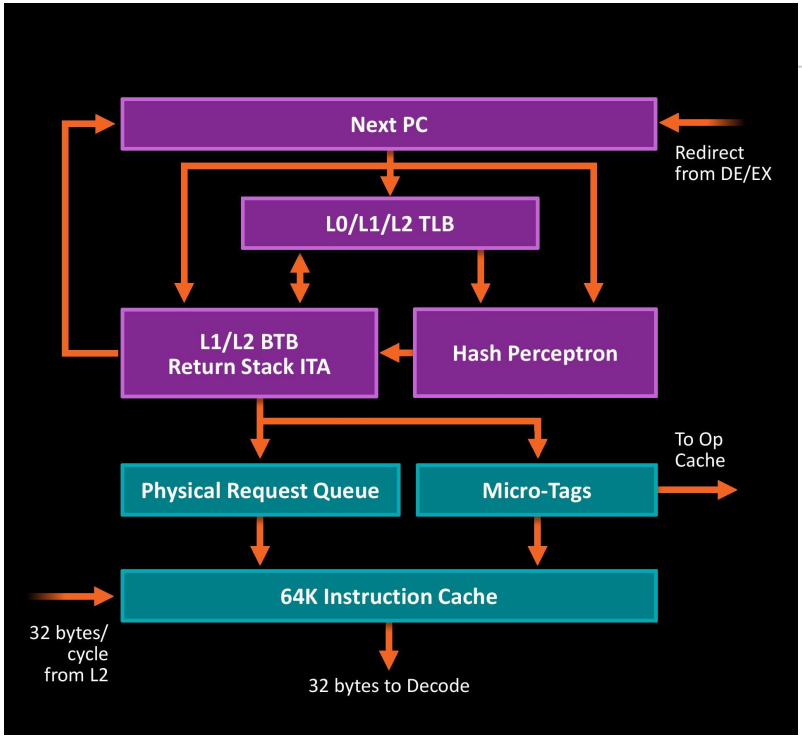

Pipeline - Decode

- Op cache can partially skip decode
- Memfile allows store-to-load forwarding
- Can dispatch 6 µops/cycle to int units and 4/cycle to FP units

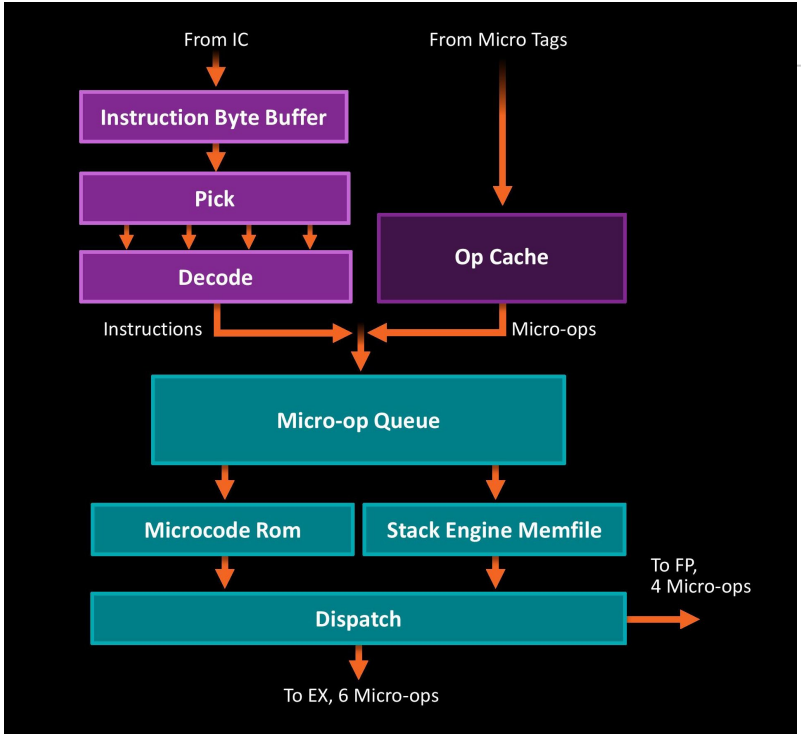

Pipeline - Execute

- 168 integer registers, 160 FP registers
- 4 ALUs, 2 AGUs, 4 FPUs
- Retire buffer has 192 entries
- 8-wide retire

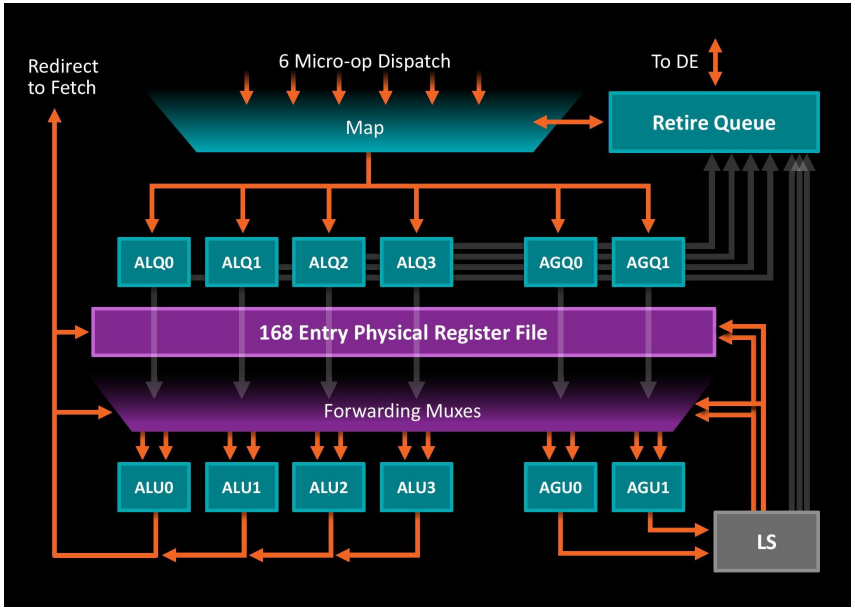

Sources: [5][8]

Caching

- L1 and L2 cache per core
- Shared L3 victim cache per four cores

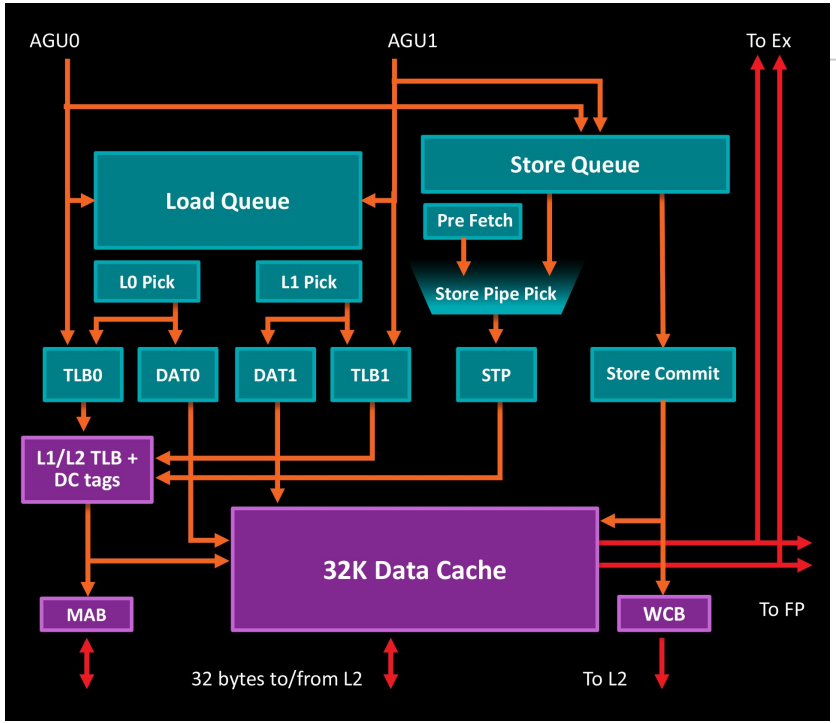

Sources: [5][8]

Power

- Consistent power draw/performance is key for servers
- Power/performance deterministic modes
- Per-core voltage regulation allows power gating
- Workload-Aware Power Management
- System Management Unit uses PID to determine optimal frequency

Sources: [4][6][8]

Branch Prediction (1)

- The key idea is to use a perceptron (the basic compute element of neural networks, essentially a simple model of a neuron) as an alternative to commonly used two-bit counter.
- The branch misprediction penalty for “Zen” is in the range from 12 to 19 cycles, depending on the type of mispredicted branch and whether or not the instructions are being fed from the microop cache

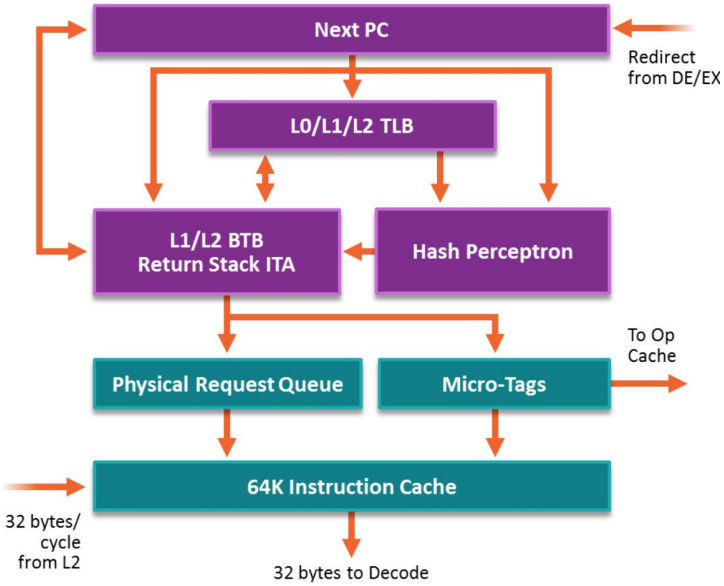

Sources: [16][20]

Branch Prediction (2)

- BP is now decoupled from fetch stage
- Next Address Logic: When no branches are identified in the current fetch block, the next-address logic calculates the starting address of the next sequential 64-byte fetch block. This calculation is performed every cycle to support the 64 byte per cycle fetch bandwidth of the op cache.
- 2 Branches per BTB entry (if branches in same 64-byte line).
- L0 BTB : 4 forward taken branches and 4 backward taken branches, and predicts with 0 bubbles. (No CALLs / RETs)
- L1 BTB : 256 entries, it creates 1 bubble if prediction differs from L0BTB
- L2 BTB :4096 entries, it creates 4 bubbles if prediction differs from L1BTB
- Return stack: 31 entry (2 * 15-entry in dual-thread)

Branch Prediction (3)

- The conditional branch predictor uses a global history scheme that keeps track of the previously executed branches.
- Global history is not updated for not-taken branches.
- Conditional branches not-taken always: not marked in the BTBs.
- Conditional branches after first-taken: predicted as always-taken.

Multi Threading (1)

- Goal: Utilize the resources in the processor core by allocating resources to two concurrent threads
- Running SMT on a single physical core will not double the performance because of resource conflicts, but SMT can get up to 40% improved performance on the multithreaded software

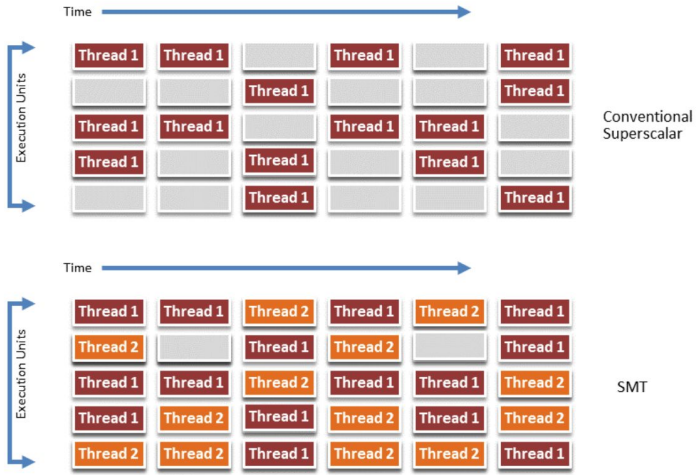

Multi Threading (2)

- Use Simultaneous MultiThreading (SMT) instead of the previous Clustered Multithreading (CMT)
- Competitively shared structures
- Competitively shared and SMT tagged
- Competitively shared /w Algorithmic Priority
- Statically Partitioned

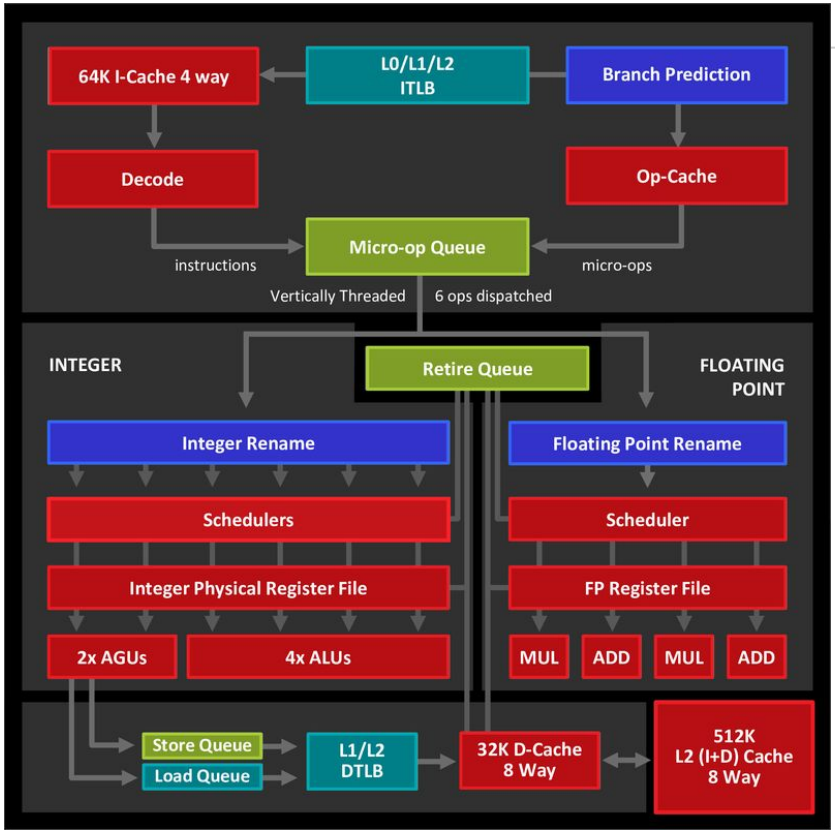

Sources: [17][11][20]

Performance (1)

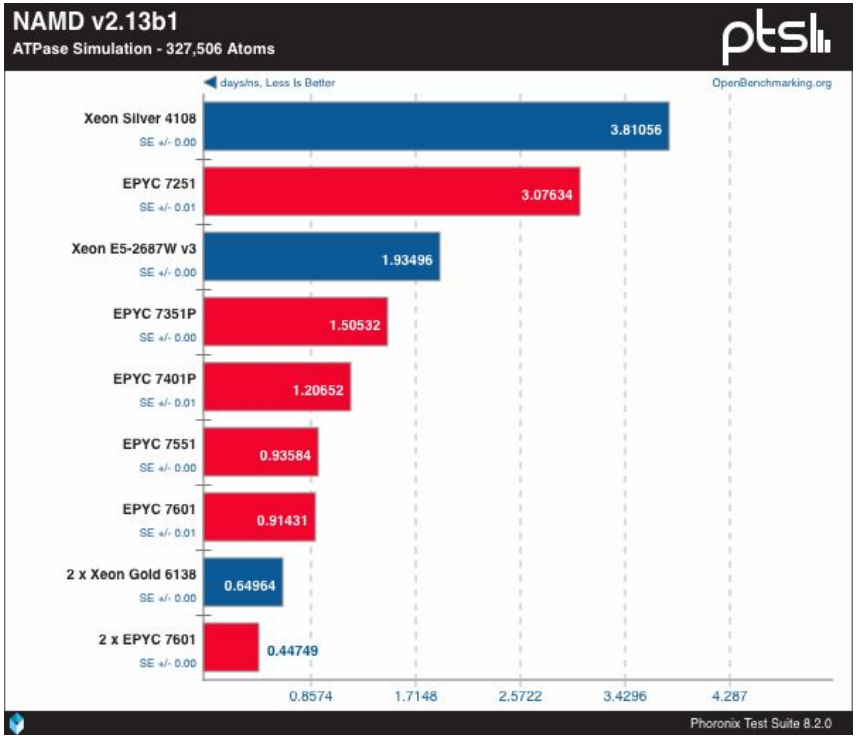

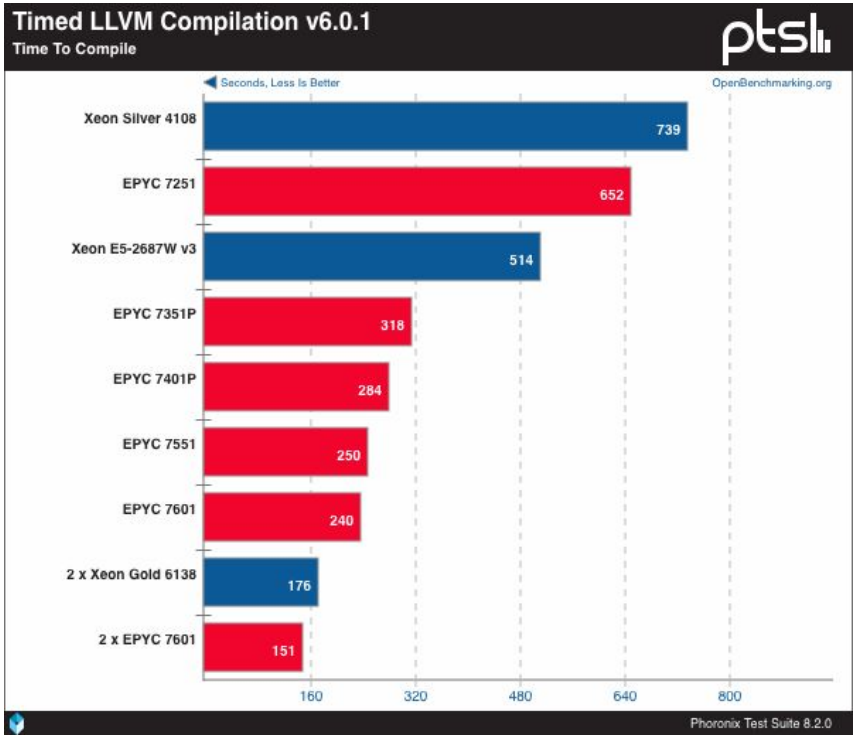

Performance (2)

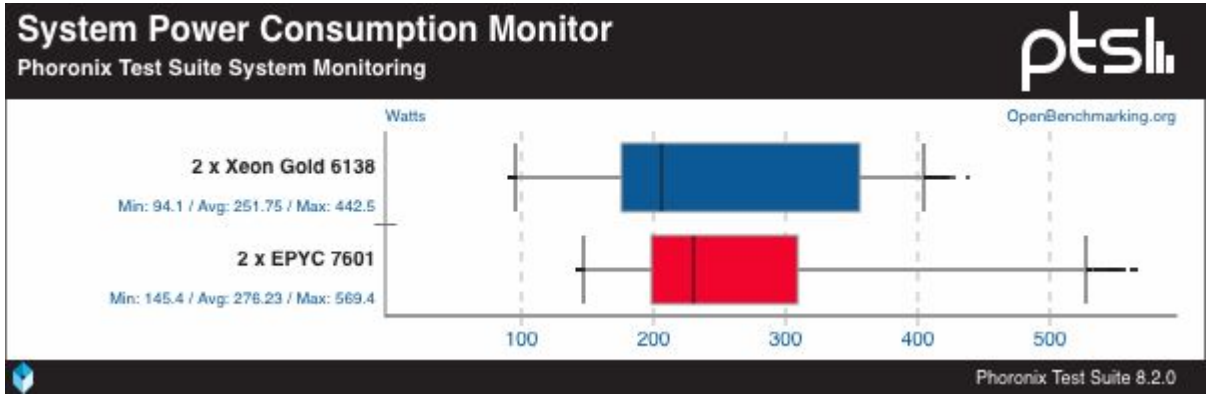

Sources: [18]

Performance (3)

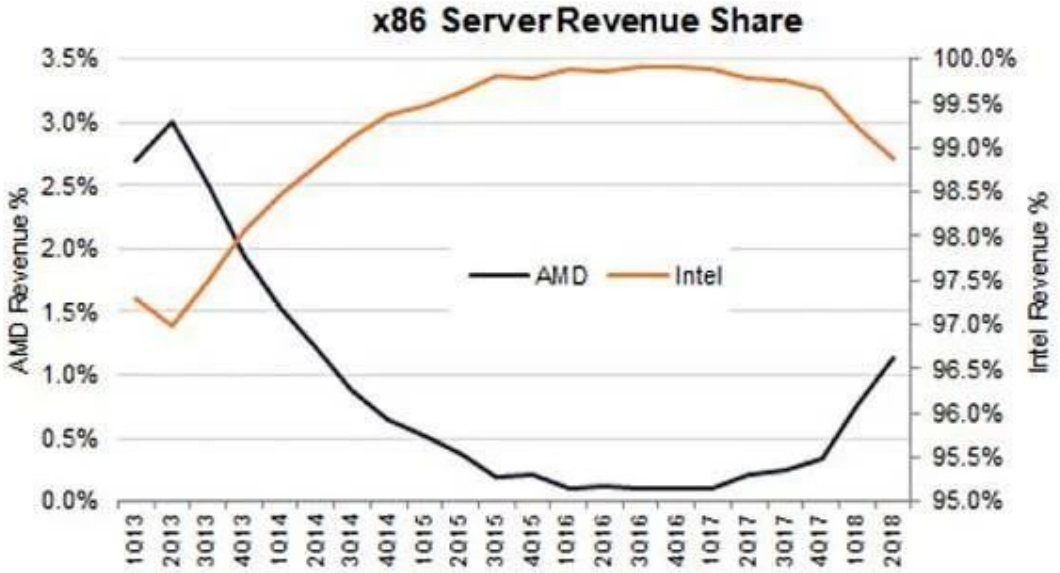

Source: Mercury Research; Wells Fargo Securities, LLC

Sources: [19]

References (1)

[1] https://www.extremetech.com/extreme/264415-epyc-secrets-amd-explains-why-it-believes-new-epyc-cpus-can-outmaneuver-outscale-the-competitio
[2] https://www.amd.com/system/files/2018-03/AMD-Optimizes-EPYC-Memory-With-NUMA.pdf
[3] http://developer.amd.com/wordpress/media/2013/12/AMD-EPYC-Security-White-Paper-Final-Jun-2017.pdf
[4] https://www.amd.com/system/files/documents/The-Energy-Efficient-AMD-EPYC-Design.pdf
[5] https://www.youtube.com/watch?v=Ln9WKPEHm4w (Hot Chips Conference 28)
[6] https://www.amd.com/system/files/2017-06/Power-Performance-Determinism.pdf
[7] https://www.amd.com/en/press-releases/amd-previews-naples-2017mar07
[8] Burd, Thomas, Noah Beck, Sean White, Milam Paraschou, Nathan Kalyanasundharam, Gregg Donley, Alan Smith, Larry Hewitt, and Samuel Naffziger. "``Zeppelin'': An SoC for Multichip Architectures." IEEE Journal of Solid-State Circuits (2018).
[9] Singh, Teja, Sundar Rangarajan, Deepesh John, Carson Henrion, Shane Southard, Hugh McIntyre, Amy Novak et al. "3.2 Zen: A next-generation high-performance× 86 core." In 2017 IEEE International Solid-State Circuits Conference (ISSCC), pp. 52-53. IEEE, 2017.

References (2)

[10] https://www.youtube.com/watch?v=NoelgG8JoyQ (EPYC™ Tech Day: Scott Aylor)
[11] https://www.youtube.com/watch?v=B65uPAFGse8&t=556s (EPYC™ Tech Day: Kevin Lepak)
[12] https://www.youtube.com/watch?v=W5IhEit6NqY (EPYC™ Tech Day: Gerry Talbot)
[13] https://en.wikichip.org/wiki/amd/epyc/7601
[14] https://www.anandtech.com/show/11551/amds-future-in-servers-new-7000-series-cpus-launched-and-epyc-analysis
[15] https://en.wikichip.org/wiki/amd/cores/naples
[16] https://www.7-cpu.com/cpu/Zen.html
[17] https://en.wikichip.org/wiki/amd/microarchitectures/zen
[18] https://www.phoronix.com/scan.php?page=article&item=amd-epyc7601-2p&num=1
[19] https://www.theregister.co.uk/2018/08/06/amd_epyc_intel_xeon_x86_server_revenue_share/
[20] https://www.tiriasresearch.com/downloads/amds-ryzen-performance-smt-and-branch-prediction-logic/

Thanks!

Any questions?