
AMD
RYZEN PROCESSOR MICROARCHITECTURE
ANASUA BHOWMIK
PRINCIPAL MEMBER OF TECHNICAL STAFF
MICROPROCESSOR PERFORMANCE MODELING AND ARCHITECTURE TEAM
AMD, BANGALORE

AGENDA

- Road to RYZEN – AMD's new X86 CPU
- High level architecture of Zen
  - Core engine
  - Floating point unit
  - Cache hierarchy
  - Simultaneous Multithreading
- Zen based SoCs
- Design Challenges
- Contributions from India Design Center
- Next step

RYZEN CPU

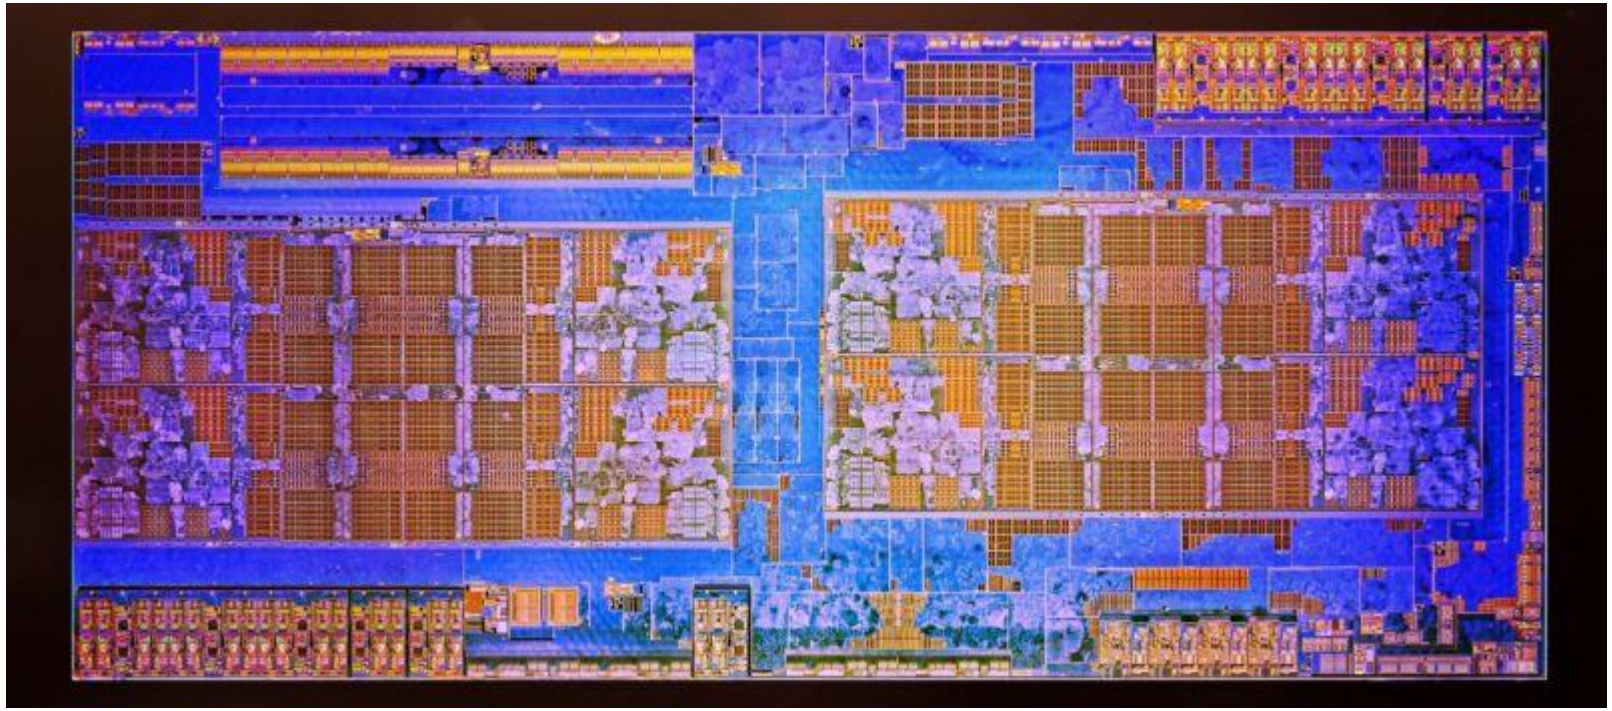

- High performance CPU
- 8 cores 16 threads
  - **Zen** core
- Over 4.8 billion transistors
- Over 2000 meters of wiring
- 14 nm process technology

CPU COMPLEX

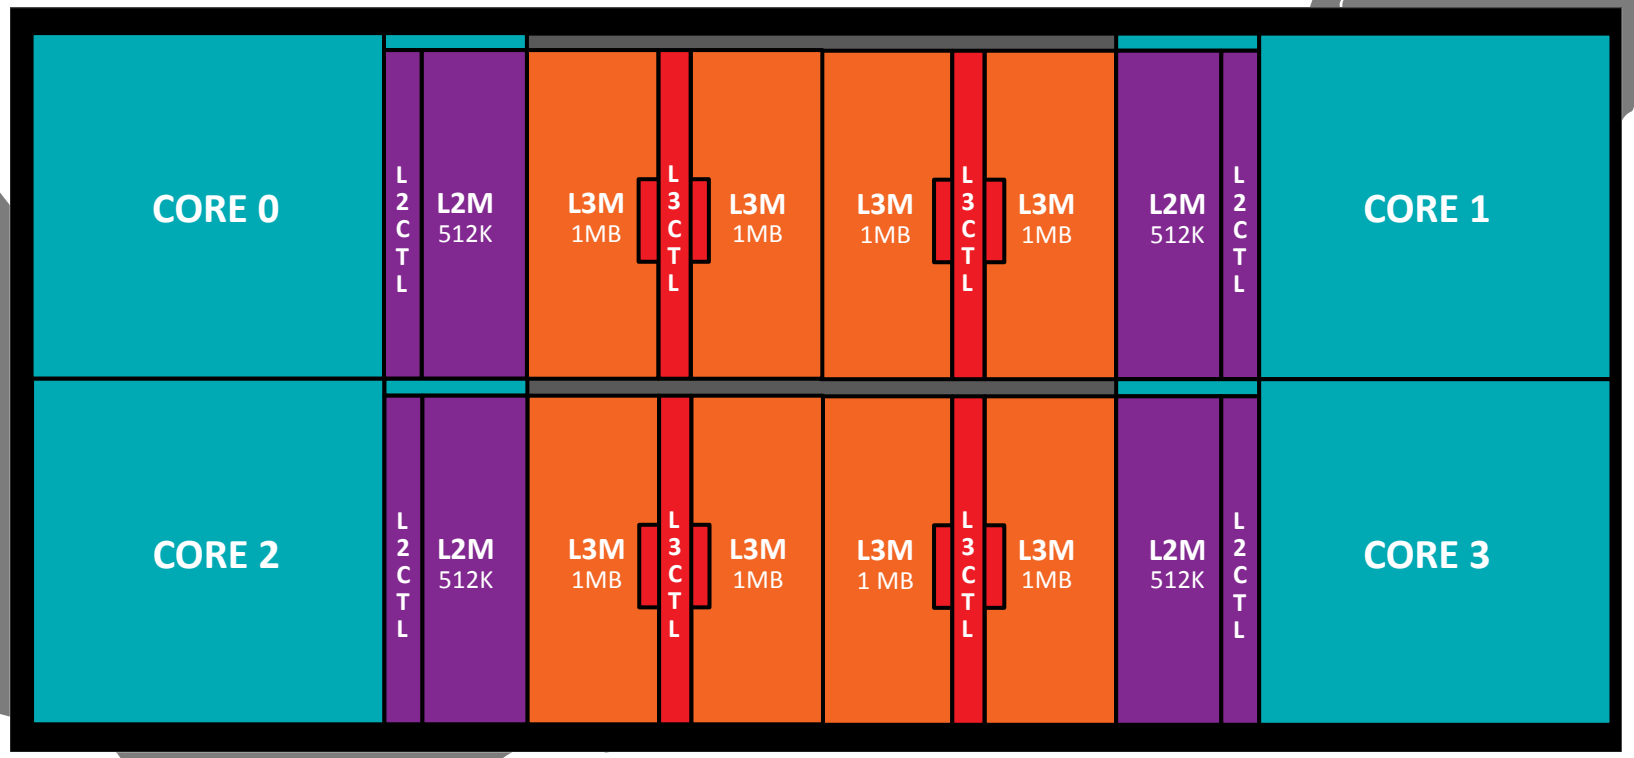

- A CPU complex (CCX) is four cores connected to an L3 Cache.
- The L3 Cache is 16-way associative, 8MB, mostly exclusive of L2.
- The L3 Cache is made of 4 slices, by low-order address interleave.
- Every core can access every cache with same average latency
- Ryzen CPU in Thread Ripper has 2 CCXs
- Epyc Server has 8 CCXs
- Ryzen mobile has 1 CCX

AMD CPU DESIGN OPTIMIZATION POINTS

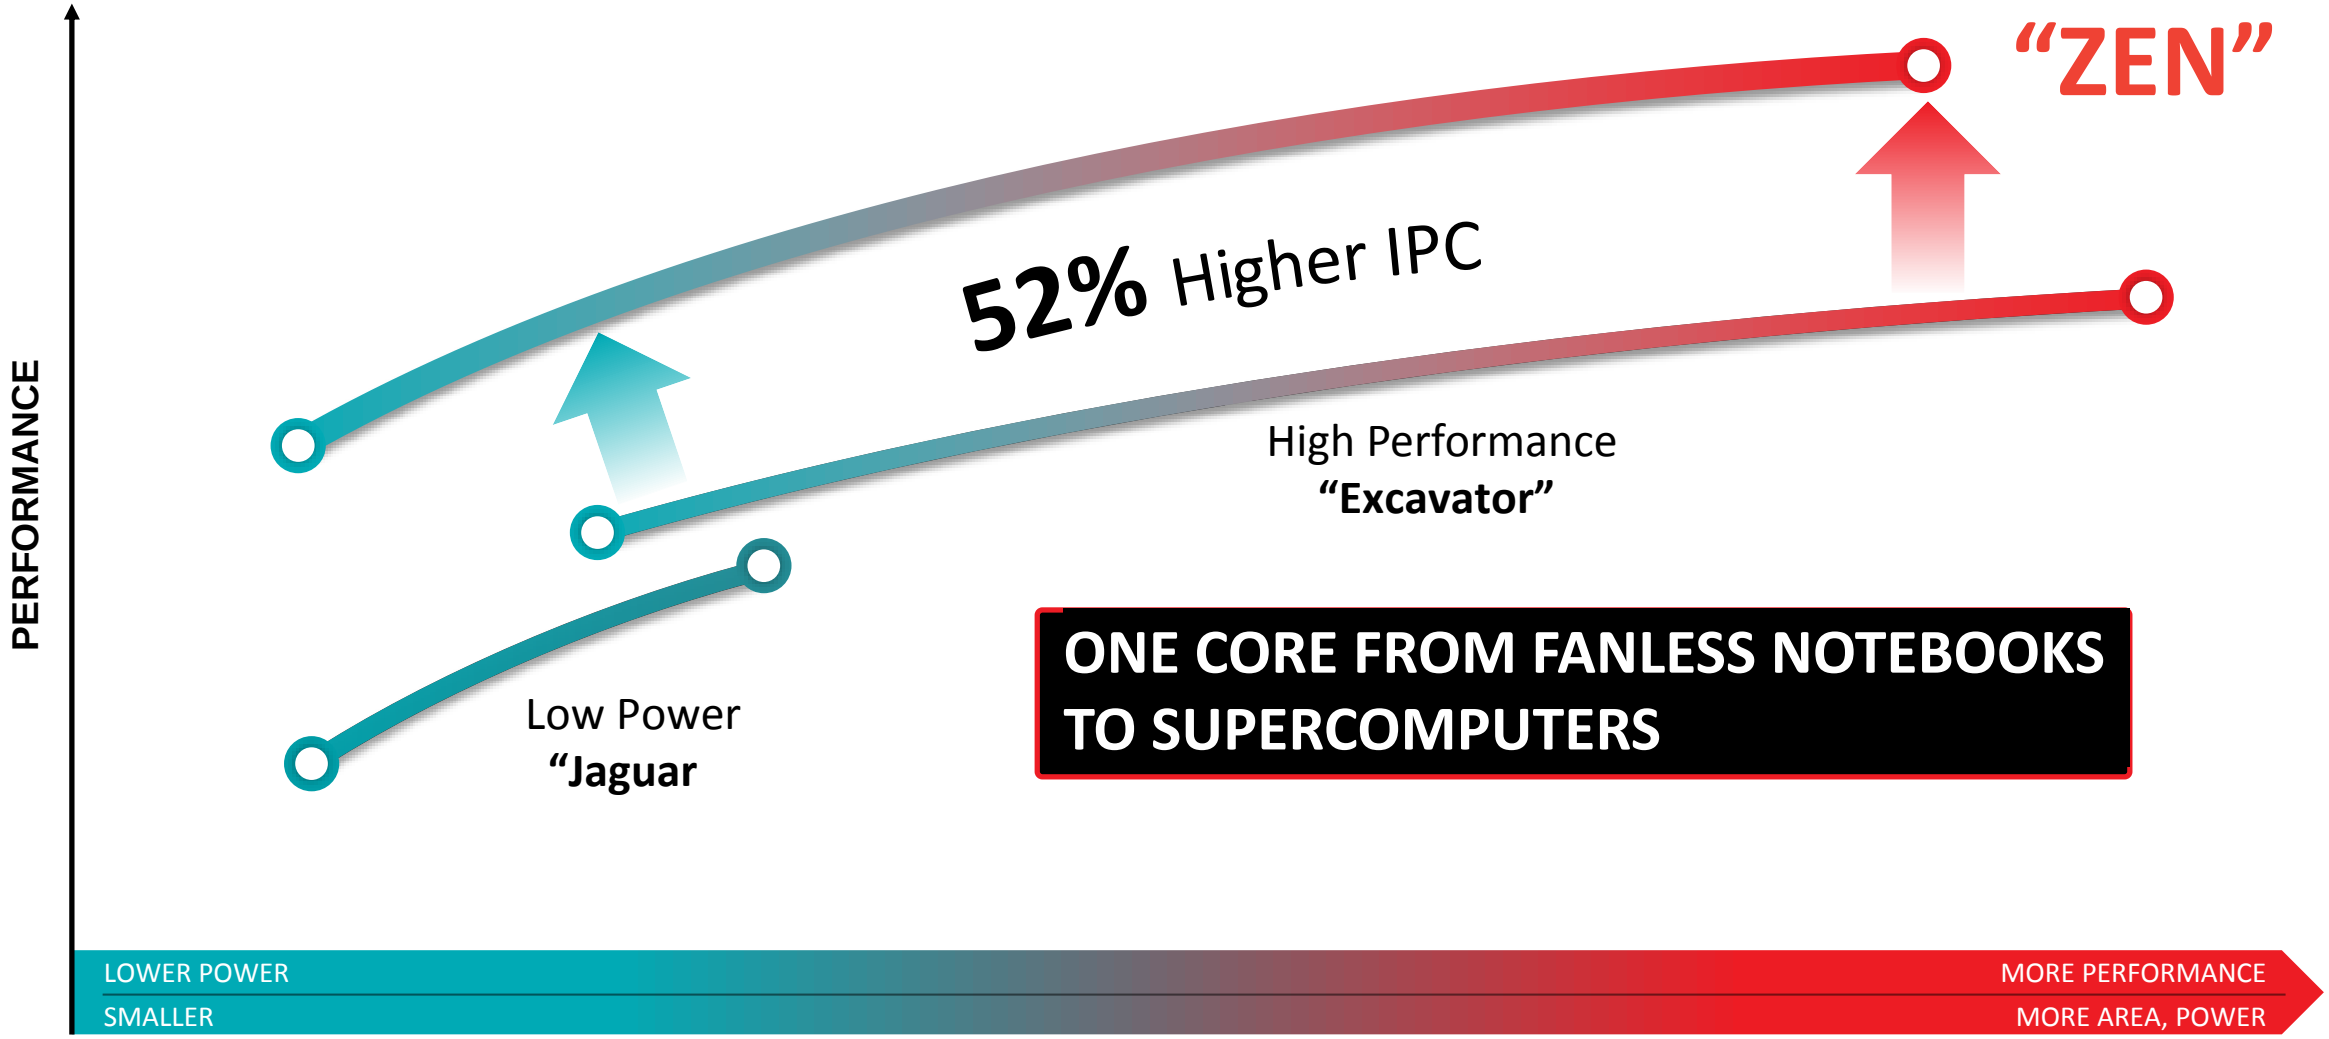

DEFYING CONVENTION: A WIDE, HIGH PERFORMANCE, EFFICIENT CORE

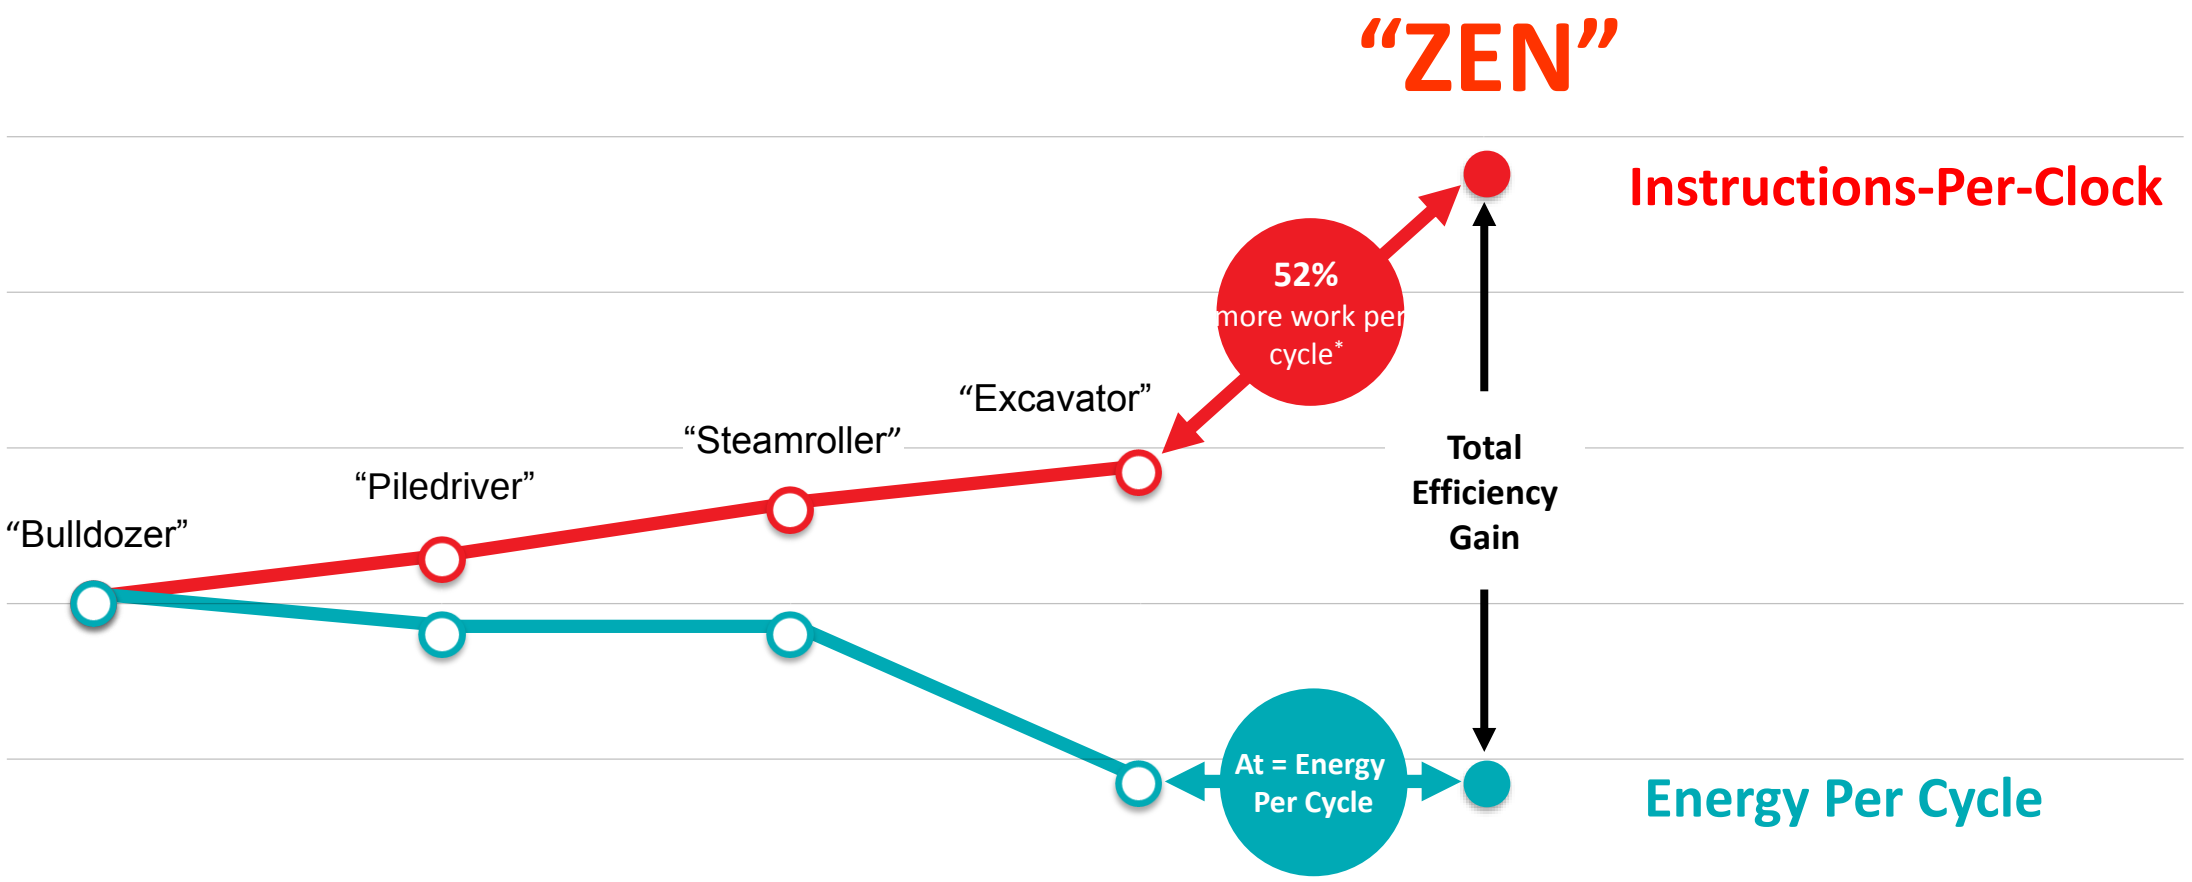

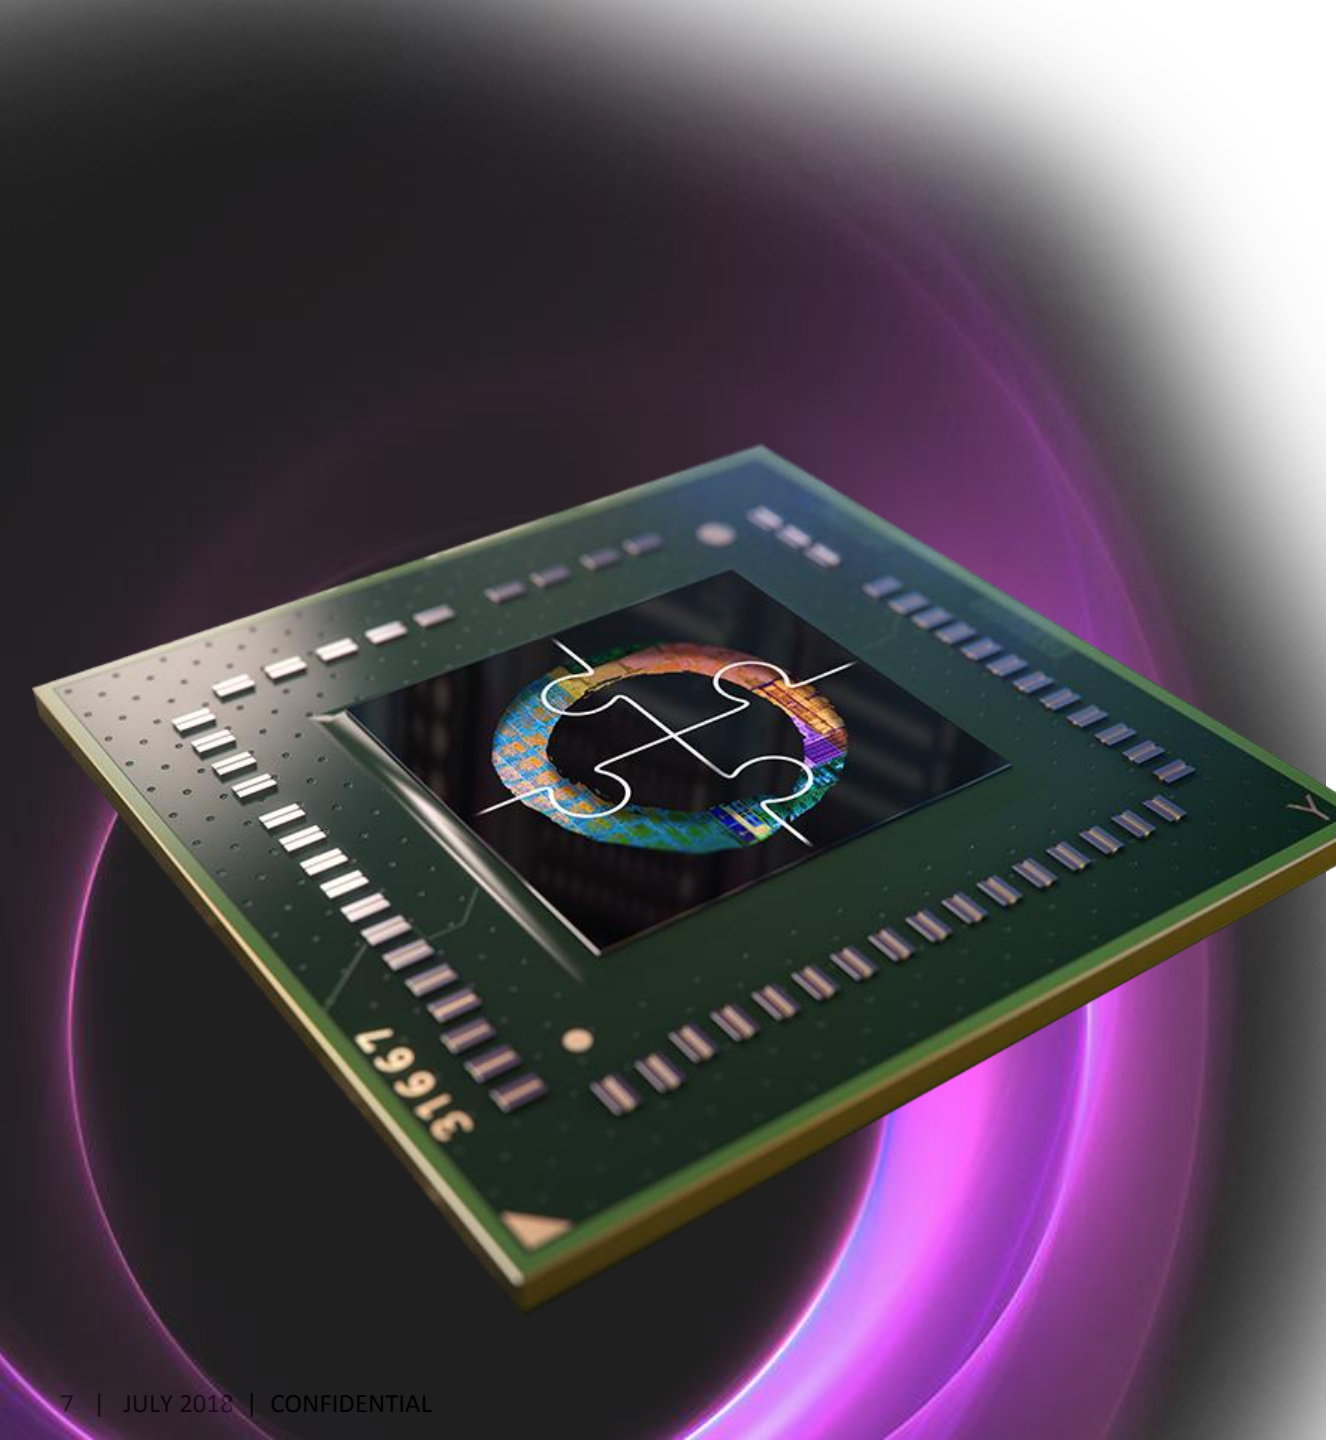

“ZEN”

DESIGNED FROM THE GROUND UP FOR OPTIMAL BALANCE OF PERFORMANCE AND POWER

Totally New High-performance Core Design	New High-Bandwidth, Low Latency Cache System
Simultaneous Multithreading (SMT) for High Throughput	Energy-efficient FinFET Design Scales from Enterprise to Client Products

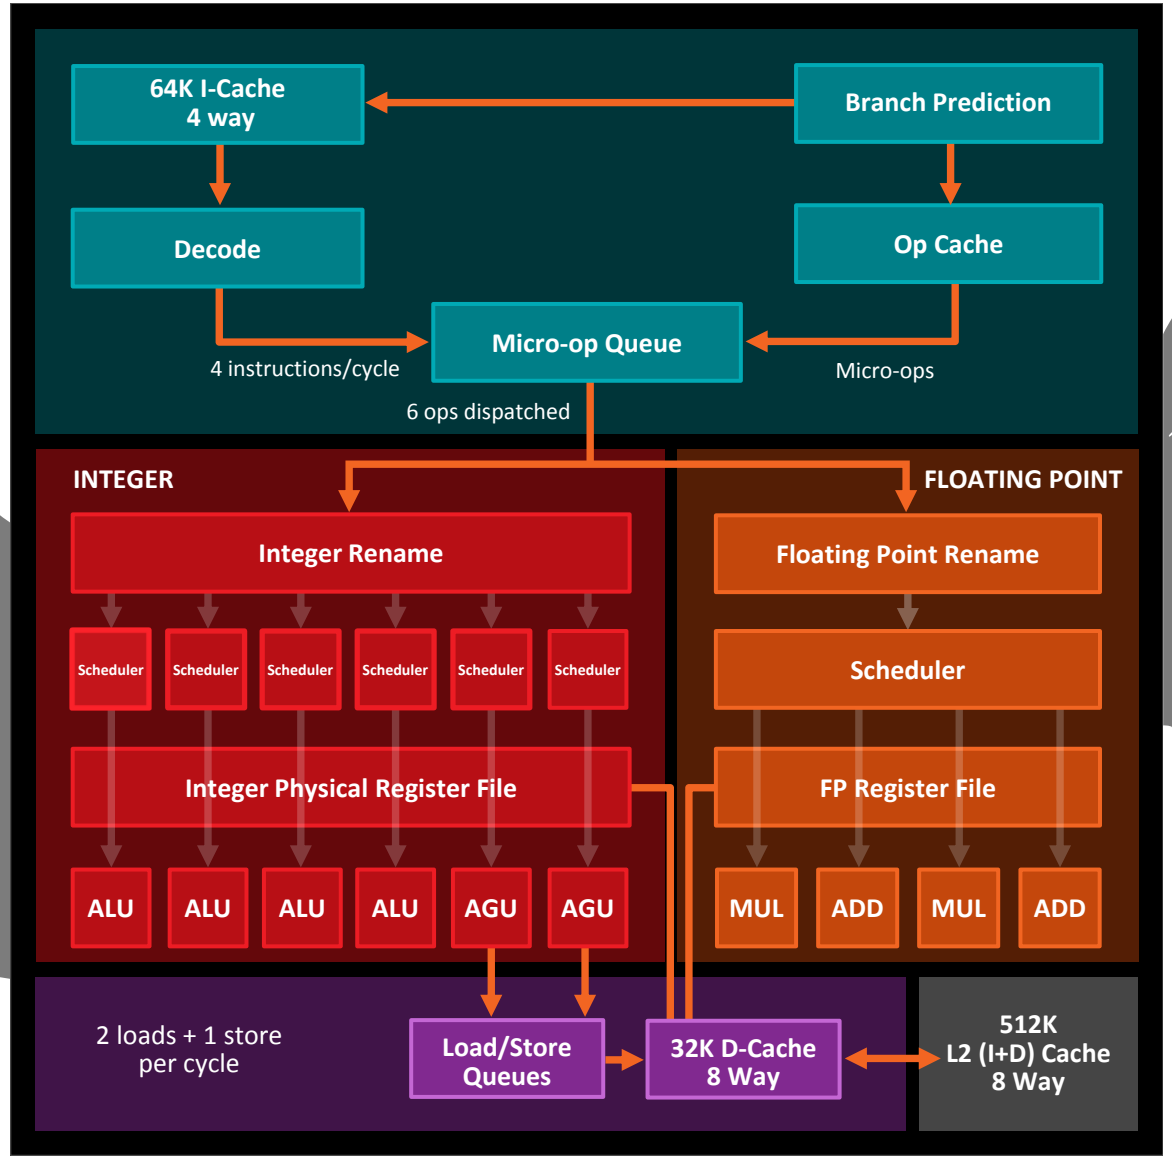

ZEN MICROARCHITECTURE

- Fetch Four x86 instructions
- Op Cache instructions
- Integer units
  - Large rename space – 168 Registers
  - 192 instructions in flight/8 wide retire
- Load/Store units
  - 72 Out-of-Order Loads supported
- Floating Point units
  - built as 4 pipes, 2 Fadd, 2 Fmul
- I-Cache 64K, 4-way
- D-Cache 32K, 8-way
- L2 Cache 512K, 8-way
- Large shared L3 cache – 8M 16 way
- 2 threads per core

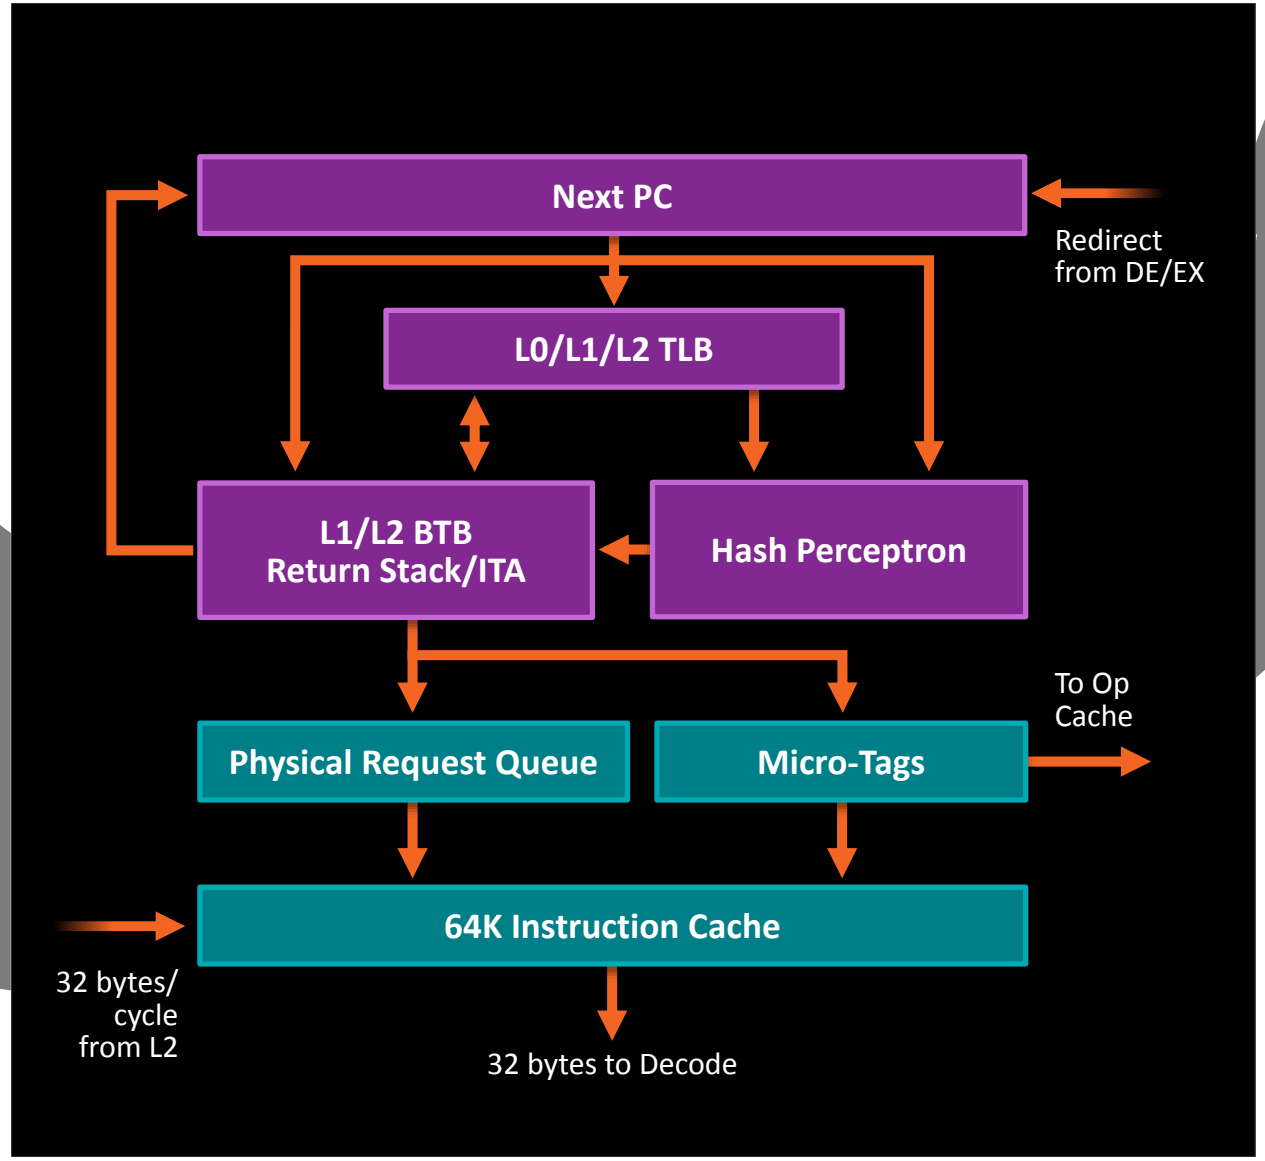

FETCH

- Decoupled Branch Prediction
- TLB in the BP pipe
  - 8 entry L0 TLB, all page sizes
  - 64 entry L1 TLB, all page sizes
  - 512 entry L2 TLB, no 1G pages
- 2 branches per BTB entry
- Large L1 / L2 BTB
- 32 entry return stack
- Indirect Target Array (ITA)
- 64K, 4-way Instruction cache
- Micro-tags for IC & Op cache
- 32 byte fetch

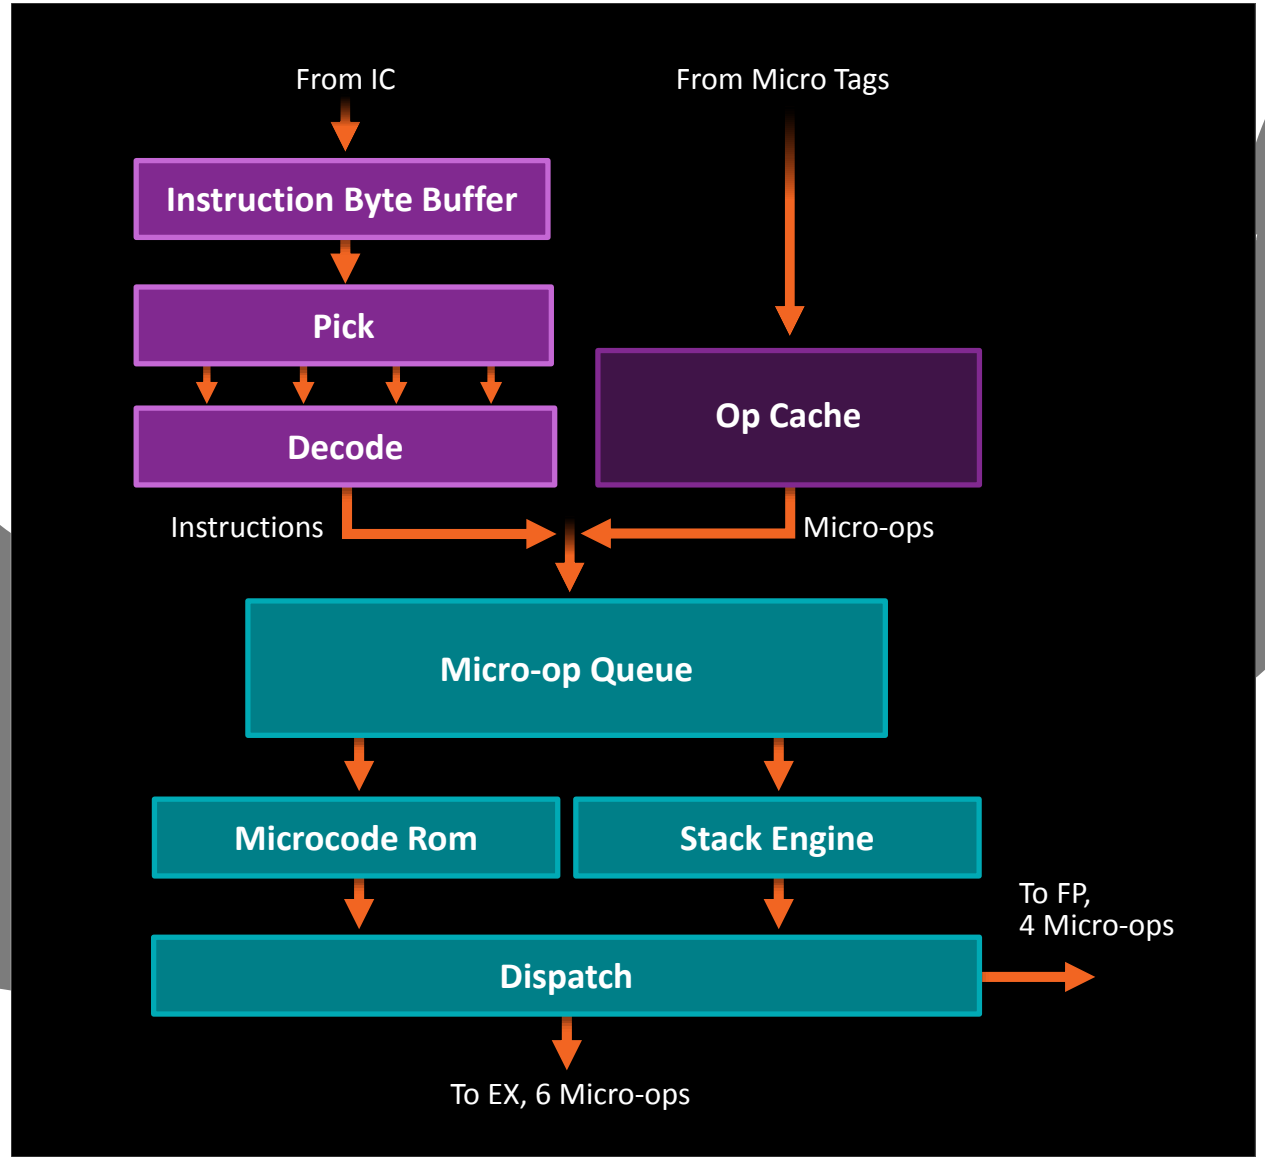

DECODE

- Inline Instruction-length Decoder
- Decode 4 x86 instructions
- Op cache
- Micro-op Queue
- Stack Engine
- Branch Fusion
- Memory File for Store to Load Forwarding

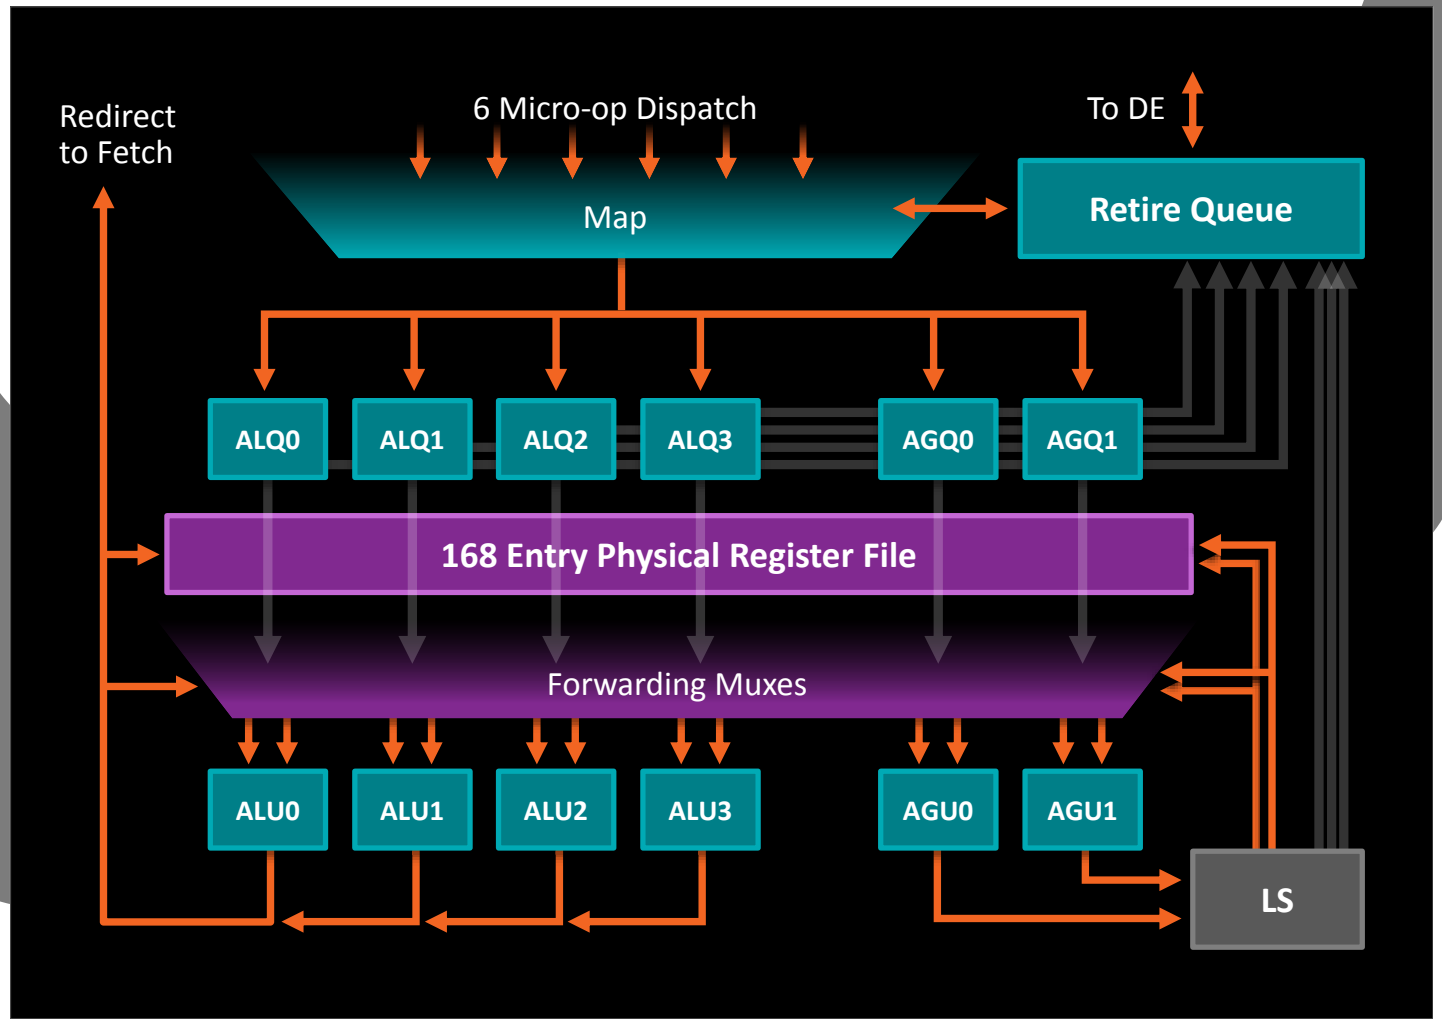

EXECUTE

- 6x14 entry Scheduling Queues
- 168 entry Physical Register File
- 6 issue per cycle
  - 4 ALU's, 2 AGU's
- 192 entry Retire Queue
- 2 Branches per cycle
- Move Elimination
- 8-Wide Retire

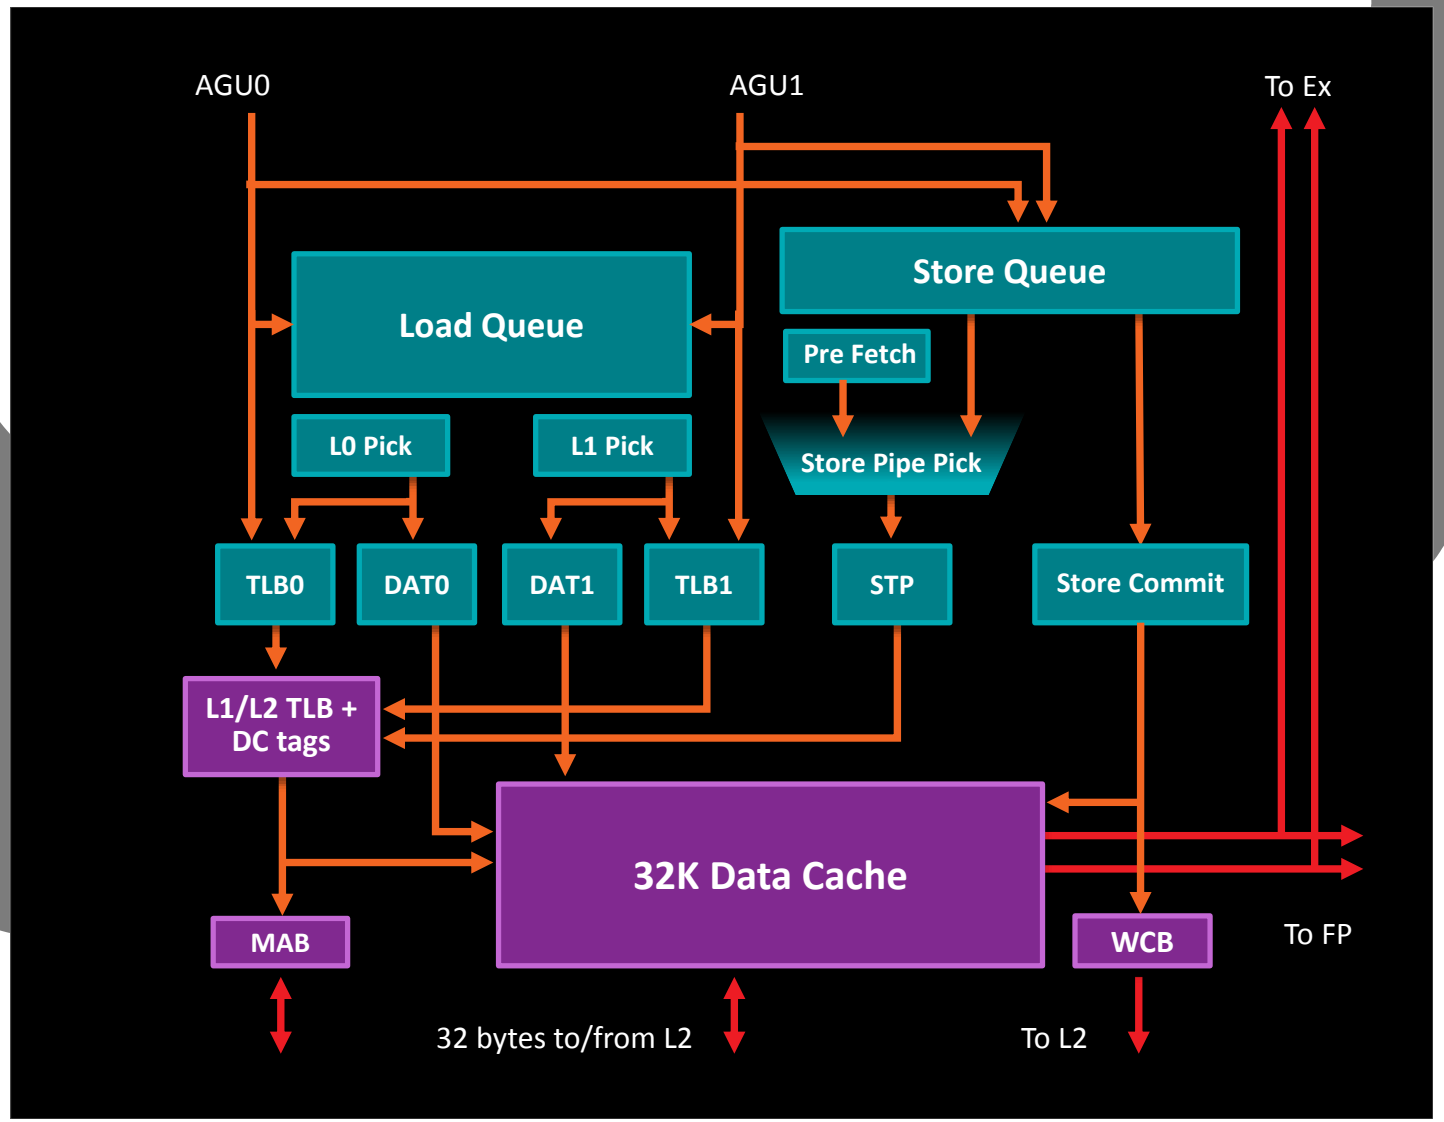

LOAD/STORE AND L2

- 72 Out of Order Loads
- 44 entry Store Queue
- Split TLB/Data Pipe, store pipe
- 64 entry L1 TLB, all page sizes
- 1.5K entry L2 TLB, no 1G pages
- 32K, 8 way Data Cache
  - Supports two 128-bit accesses
- Optimized L1 and L2 Prefetchers
- 512K, private (2 threads), inclusive

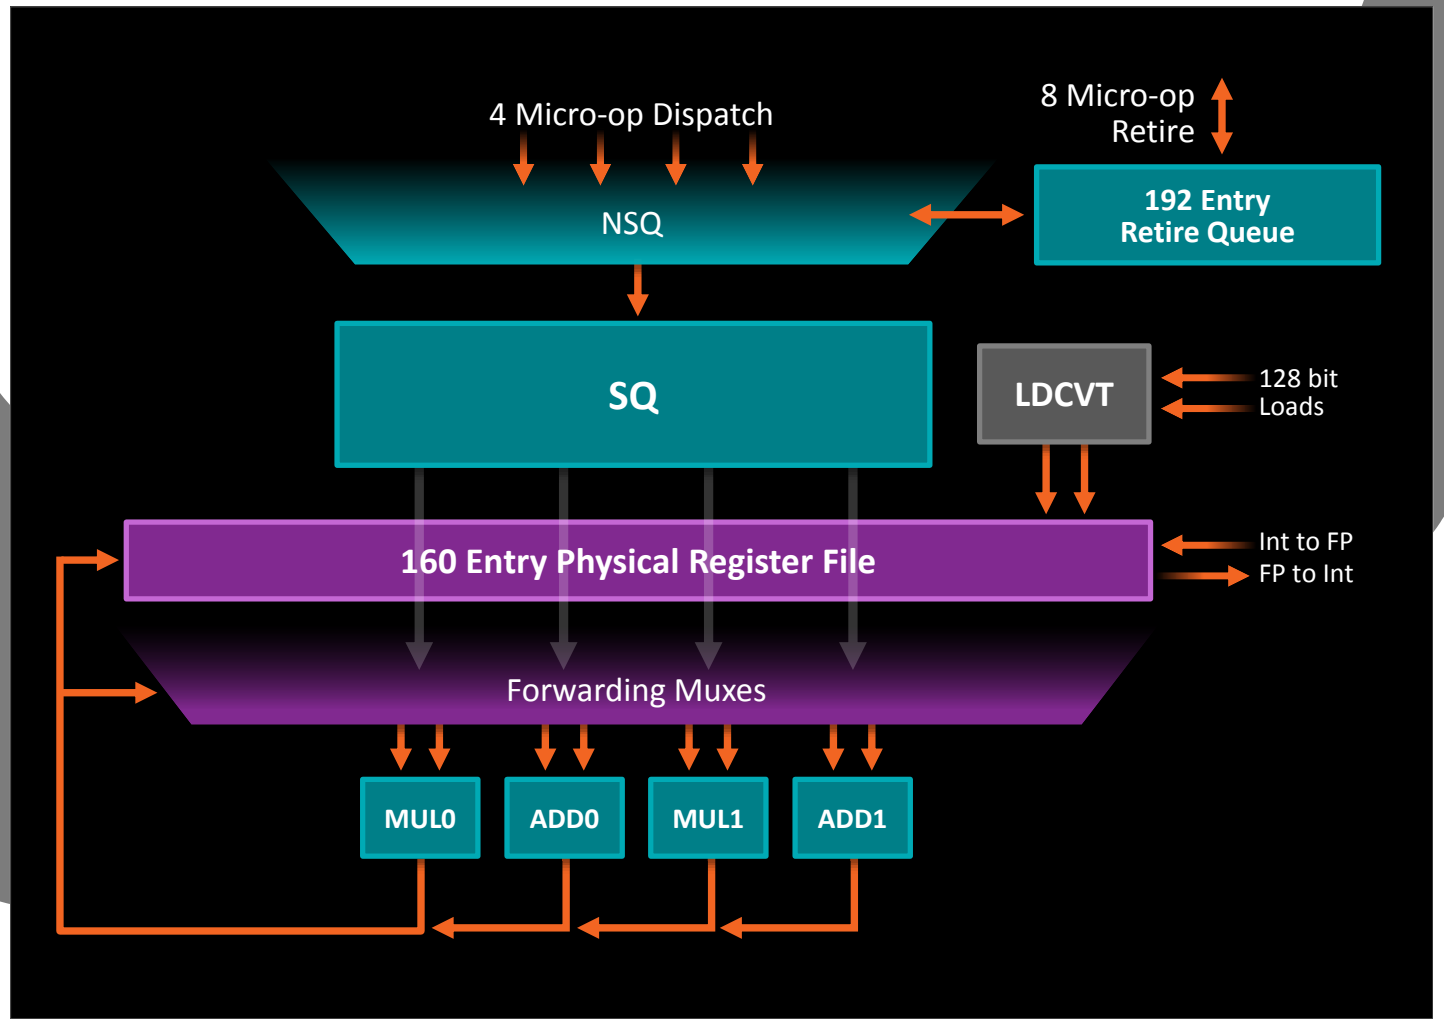

FLOATING POINT

- 2 Level Scheduling Queue
- 160 entry Physical Register File
- 8 Wide Retire
- 1 pipe for 1x128b store
- Accelerated Recovery on Flushes
- SSE, AVX1, AVX2, AES and legacy mmx/x87 compliant
- 2 AES units

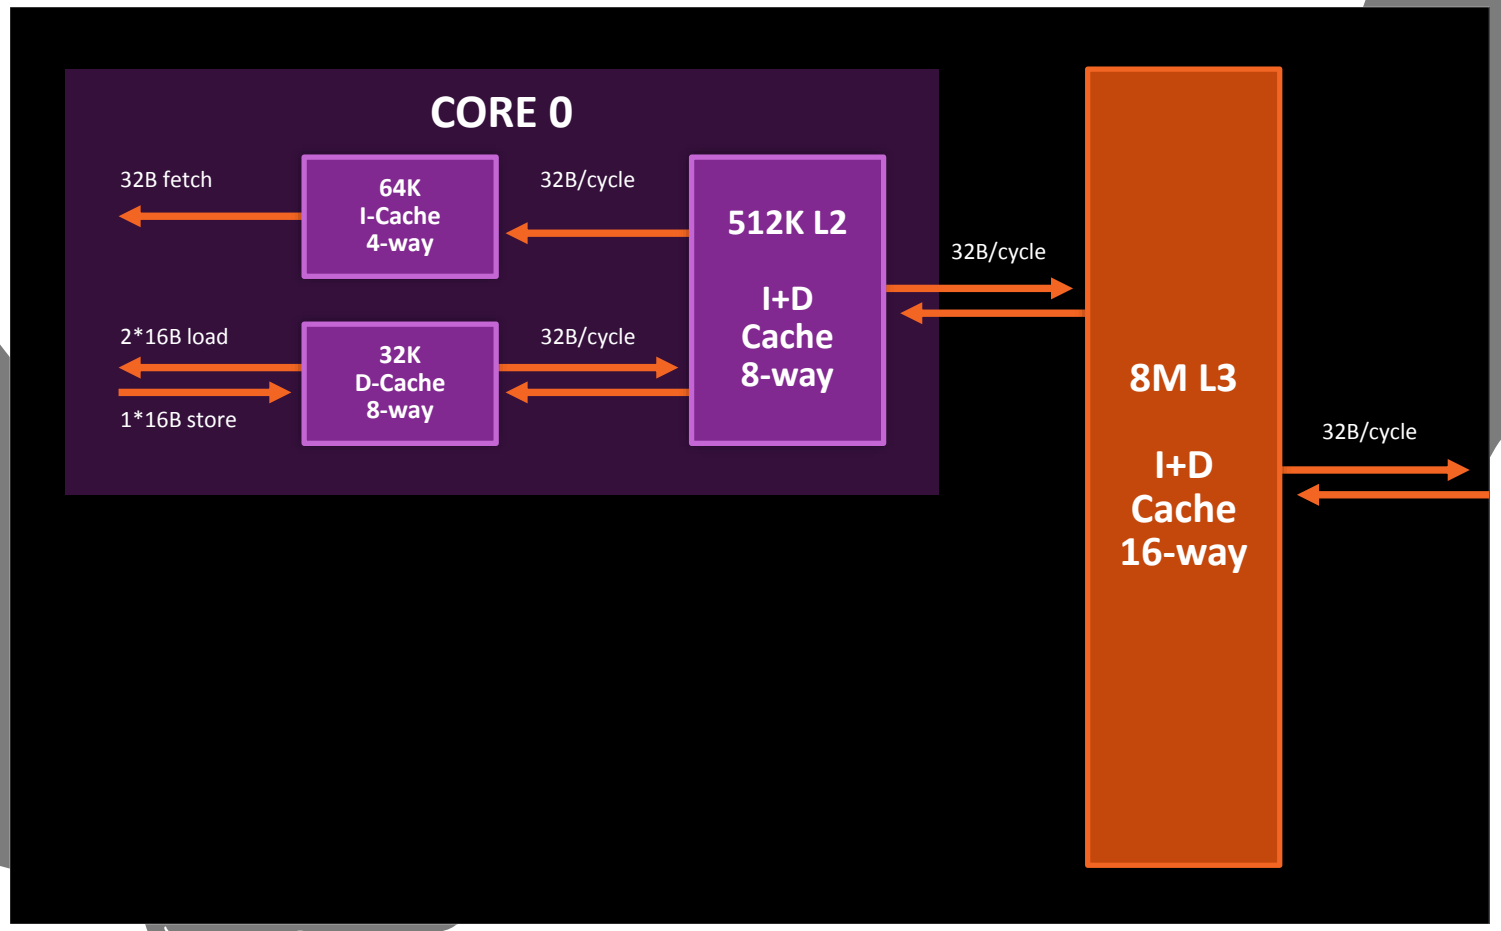

ZEN CACHE HIERARCHY

- Fast private 512K L2 cache
- Fast shared L3 cache
- High bandwidth enables prefetch improvements
- L3 is filled from L2 victims
- Fast cache-to-cache transfers
- Large Queues for Handling L1 and L2 misses

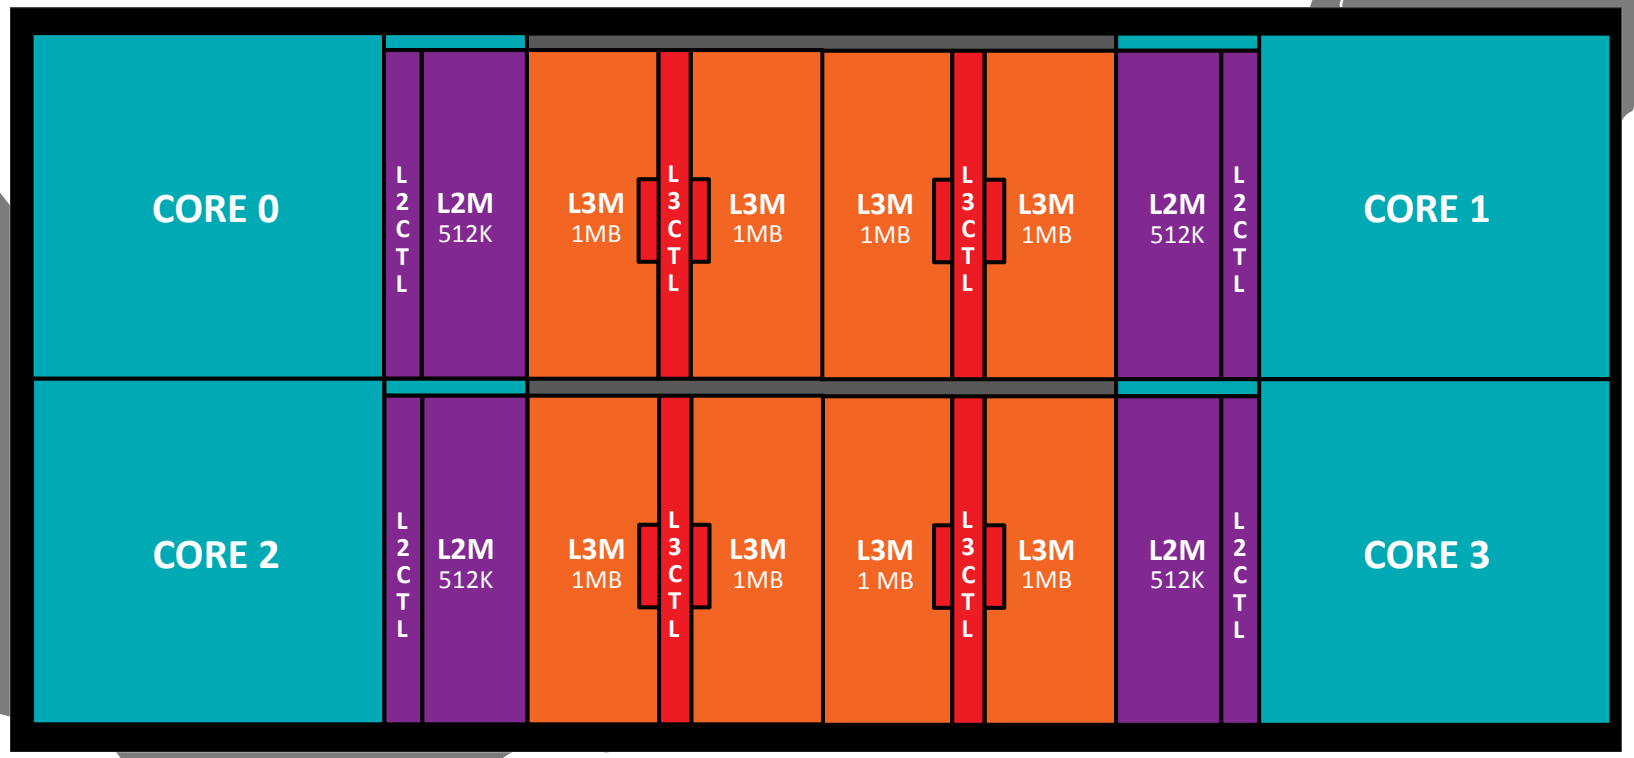

CPU COMPLEX

- A CPU complex (CCX) is four cores connected to an L3 Cache.
- The L3 Cache is 16-way associative, 8MB, mostly exclusive of L2.
- The L3 Cache is made of 4 slices, by low-order address interleave.
- Every core can access every cache with same average latency
- Ryzen has 2 CCXs
- Naples Server has 4 CCXs

SMT OVERVIEW

- All structures fully available in 1T mode
- Front End Queues are round robin with priority overrides
- Increased throughput from SMT

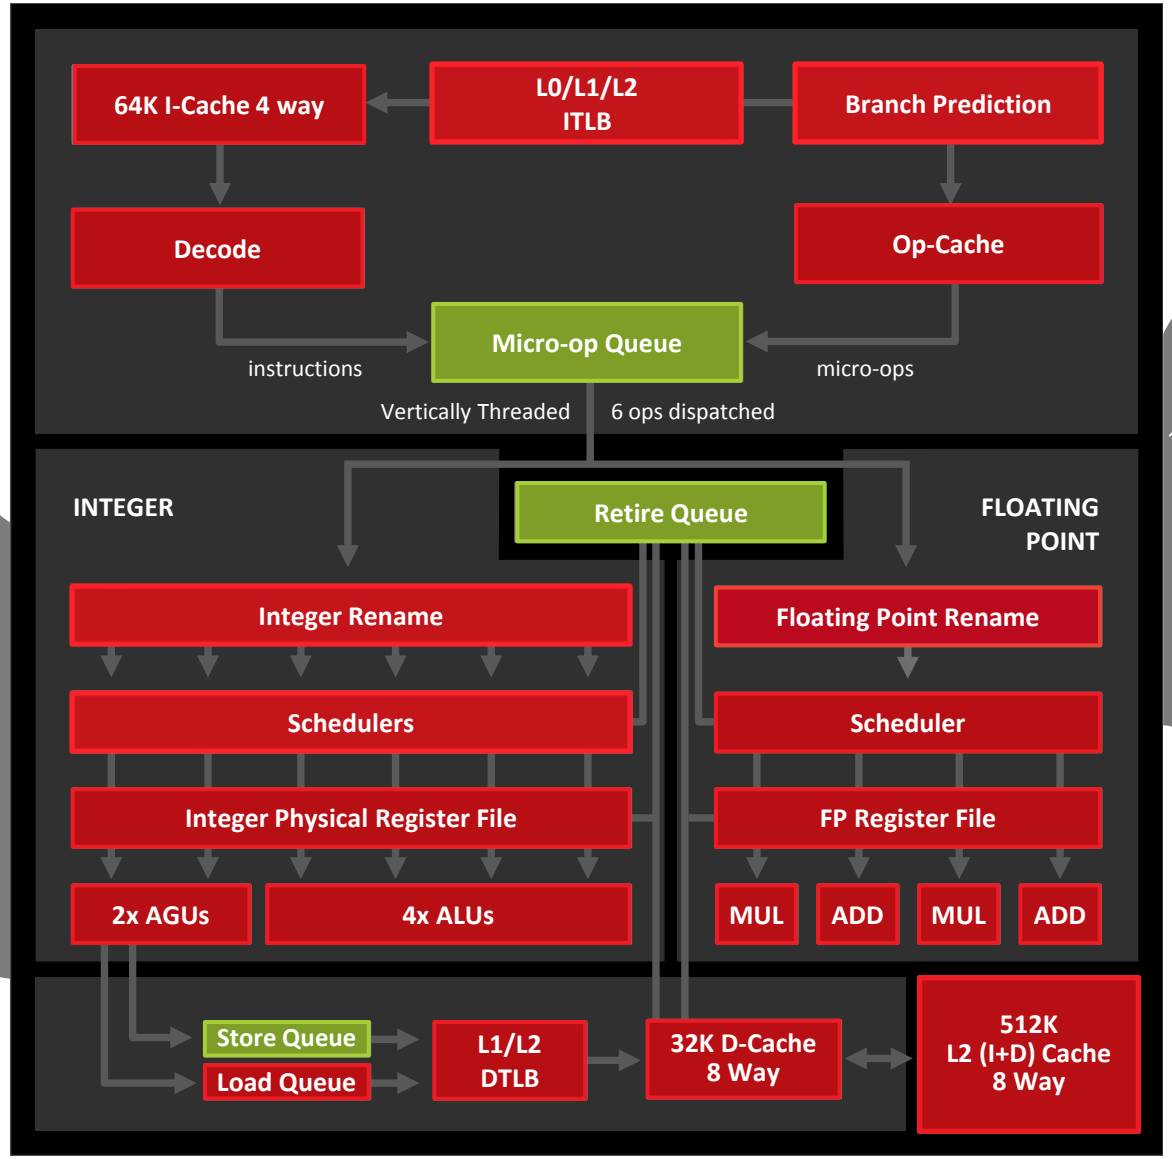

Competitively shared structures

Statically Partitioned

ZEN PERFORMANCE & POWER IMPROVEMENTS

AMD

BETTER CORE ENGINE

- SMT - two threads per core
- Better branch prediction
- Large Op Cache
- Wider micro-op dispatch 6 vs. 4
- Larger Instruction Schedulers
  Integer: 84 vs. 48 | FP: 96 vs. 60
- Larger retire 8 ops vs. 4 ops
- 4 issue FPU
- Larger Retire Queue 192 vs. 128
- Larger Load Queue 72 vs. 44
- Larger Store Queue

BETTER CACHE

- Write back L1 cache
- Faster L2 cache
- Faster L3 cache
- Faster Load to FPU: 7 vs. 9 cycles
- Better L1 and L2 data prefetcher
- Close to 2x the L1 and L2 bandwidth
- Total L3 bandwidth up 5x

LOWER POWER

- Aggressive clock gating with multi-level regions
- Write back L1 cache
- Large Op Cache
- Stack Engine
- Move elimination
- Power focus from project inception

52% IPC PERFORMANCE UPLIFT

RYZEN BASED SOCS

- **EPYC SERVER**
  - 8 CCX, 4 Dies in each socket
  - Connected by fast Infinity fabric
  - Support for virtualization, encryption, ECC
  - Targeted for Data center
- **THREAD RIPPER DESKTOP**
  - High Frequency
  - Gaming performance
- **RAVEN RIDGE LOW POWER APU (CPU + GPU)**
  - 4 M L3
  - Power optimization

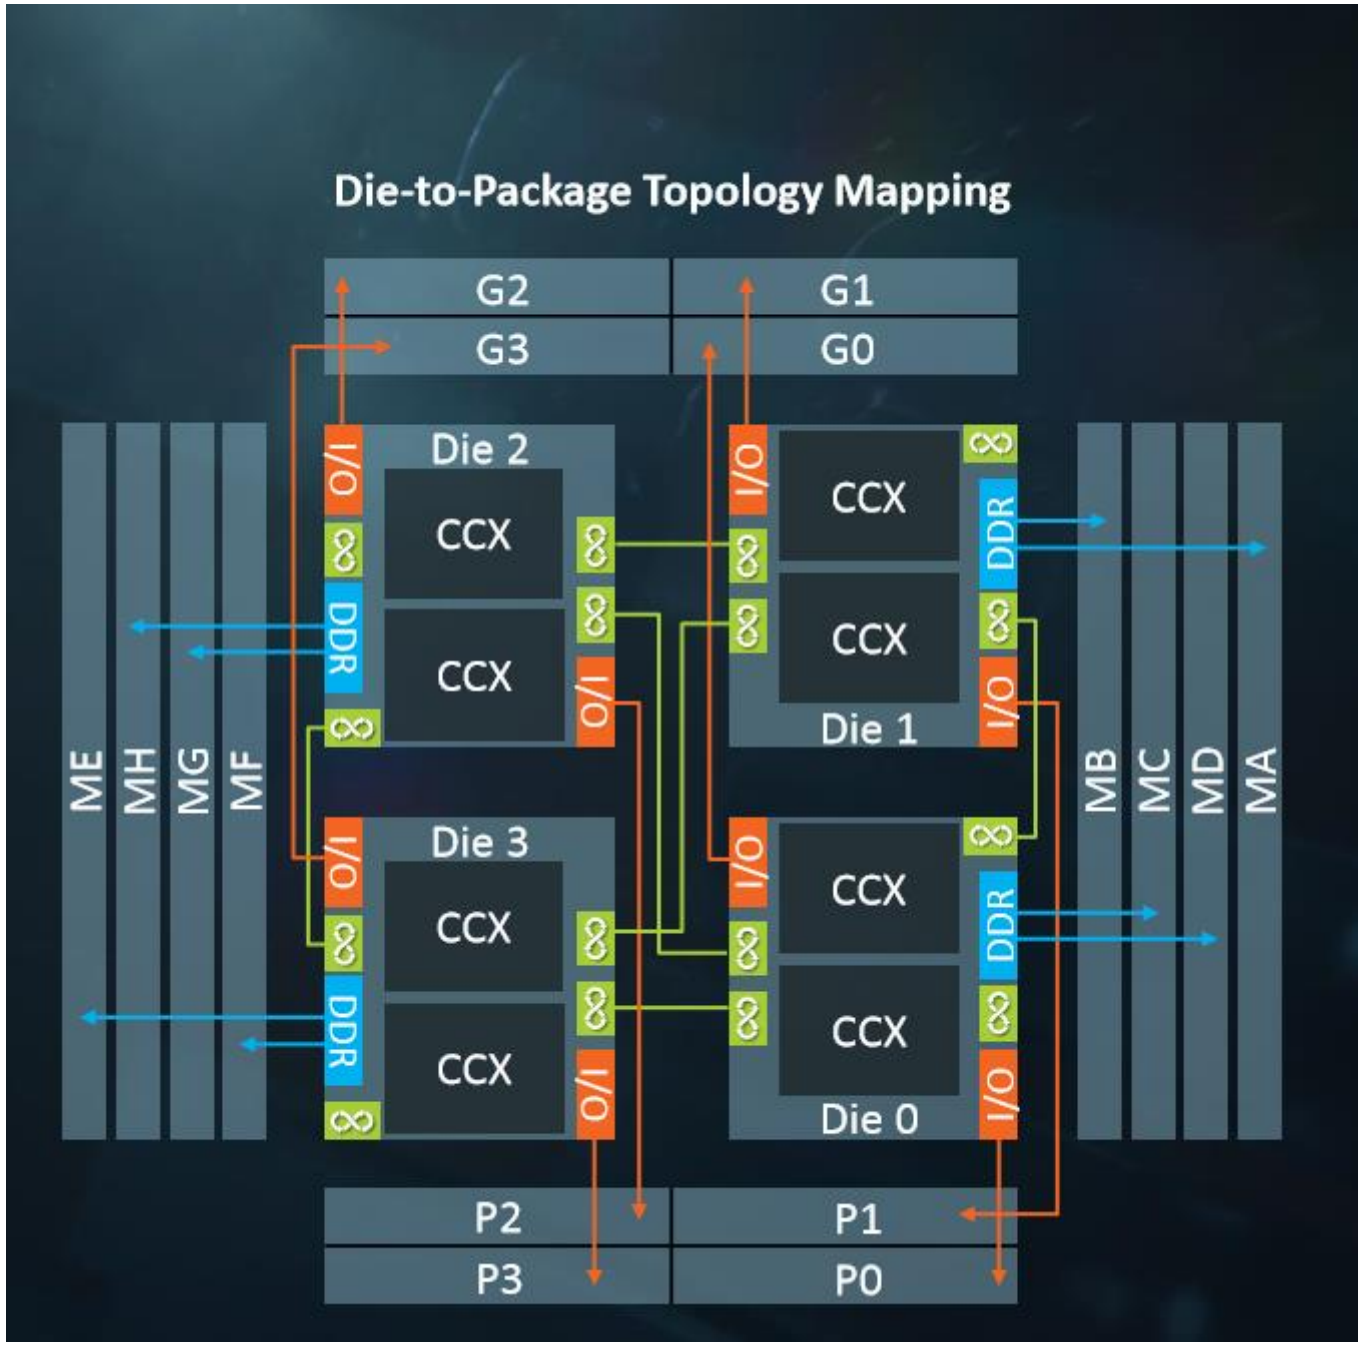

DESIGN CHALLENGES

AMD

- Long and expensive processor development project
  - More than 2 million engineering hours for Zen over four years
- Need to make key design decisions early in the project which has long lasting impact
  - Dependence on unknowns like process readiness in the fab, yield
  - Area, power
  - Availability of technology
- Conflicting needs of various SoCs
  - Server : Bigger LLC, MT performance more important than 1T performance
  - Desktop: High frequency, 1T is more important (aggressive prefetching)
  - Client: Area and power sensitive, smaller LLC
  - Semicustom: Specific ISA supports
  - Ultra low power
- Workloads to consider
- Competitiveness of the parts

CONTRIBUTIONS FROM INDIA DESIGN CENTER

AMD

- Performance modeling, micro-architectural feature exploration, performance verification
  - Development of cache hierarchy performance model (L2, L3, prefetcher) and architecture
- Traffic optimization between cores and memory
- All (Server, Client and Cores) workloads analysis, characterization and collection
- Core, Data Fabric and SoC level verification
- L3 RTL development for Zen based client part
- Zen based client SoC architecture development
- Compiler optimizations for Zen based server (Epyc)

WORK IN PROGRESS

AMD

- Zen2 hardware is in the lab
  - 7 nm core with significant performance uplift over Zen
- Microarchitecture development, performance modeling and correlation for core following Zen2
- Exploring micro-architectural ideas for future processors
- Improving performance analysis and modeling methodologies
  - Performance bottleneck analysis
  - Very high thread count simulation, impact of sharing and locks
  - Warming-up large caches
  - Execution driven vs trace driven simulation
  - SoC level modeling – CPU/GPU/accelerators interaction
- Workloads analysis
  - Cloud, virtualized workloads, machine learning, HPC workloads

Lots of challenging problem to work on!!

PROJECT TIMELINE

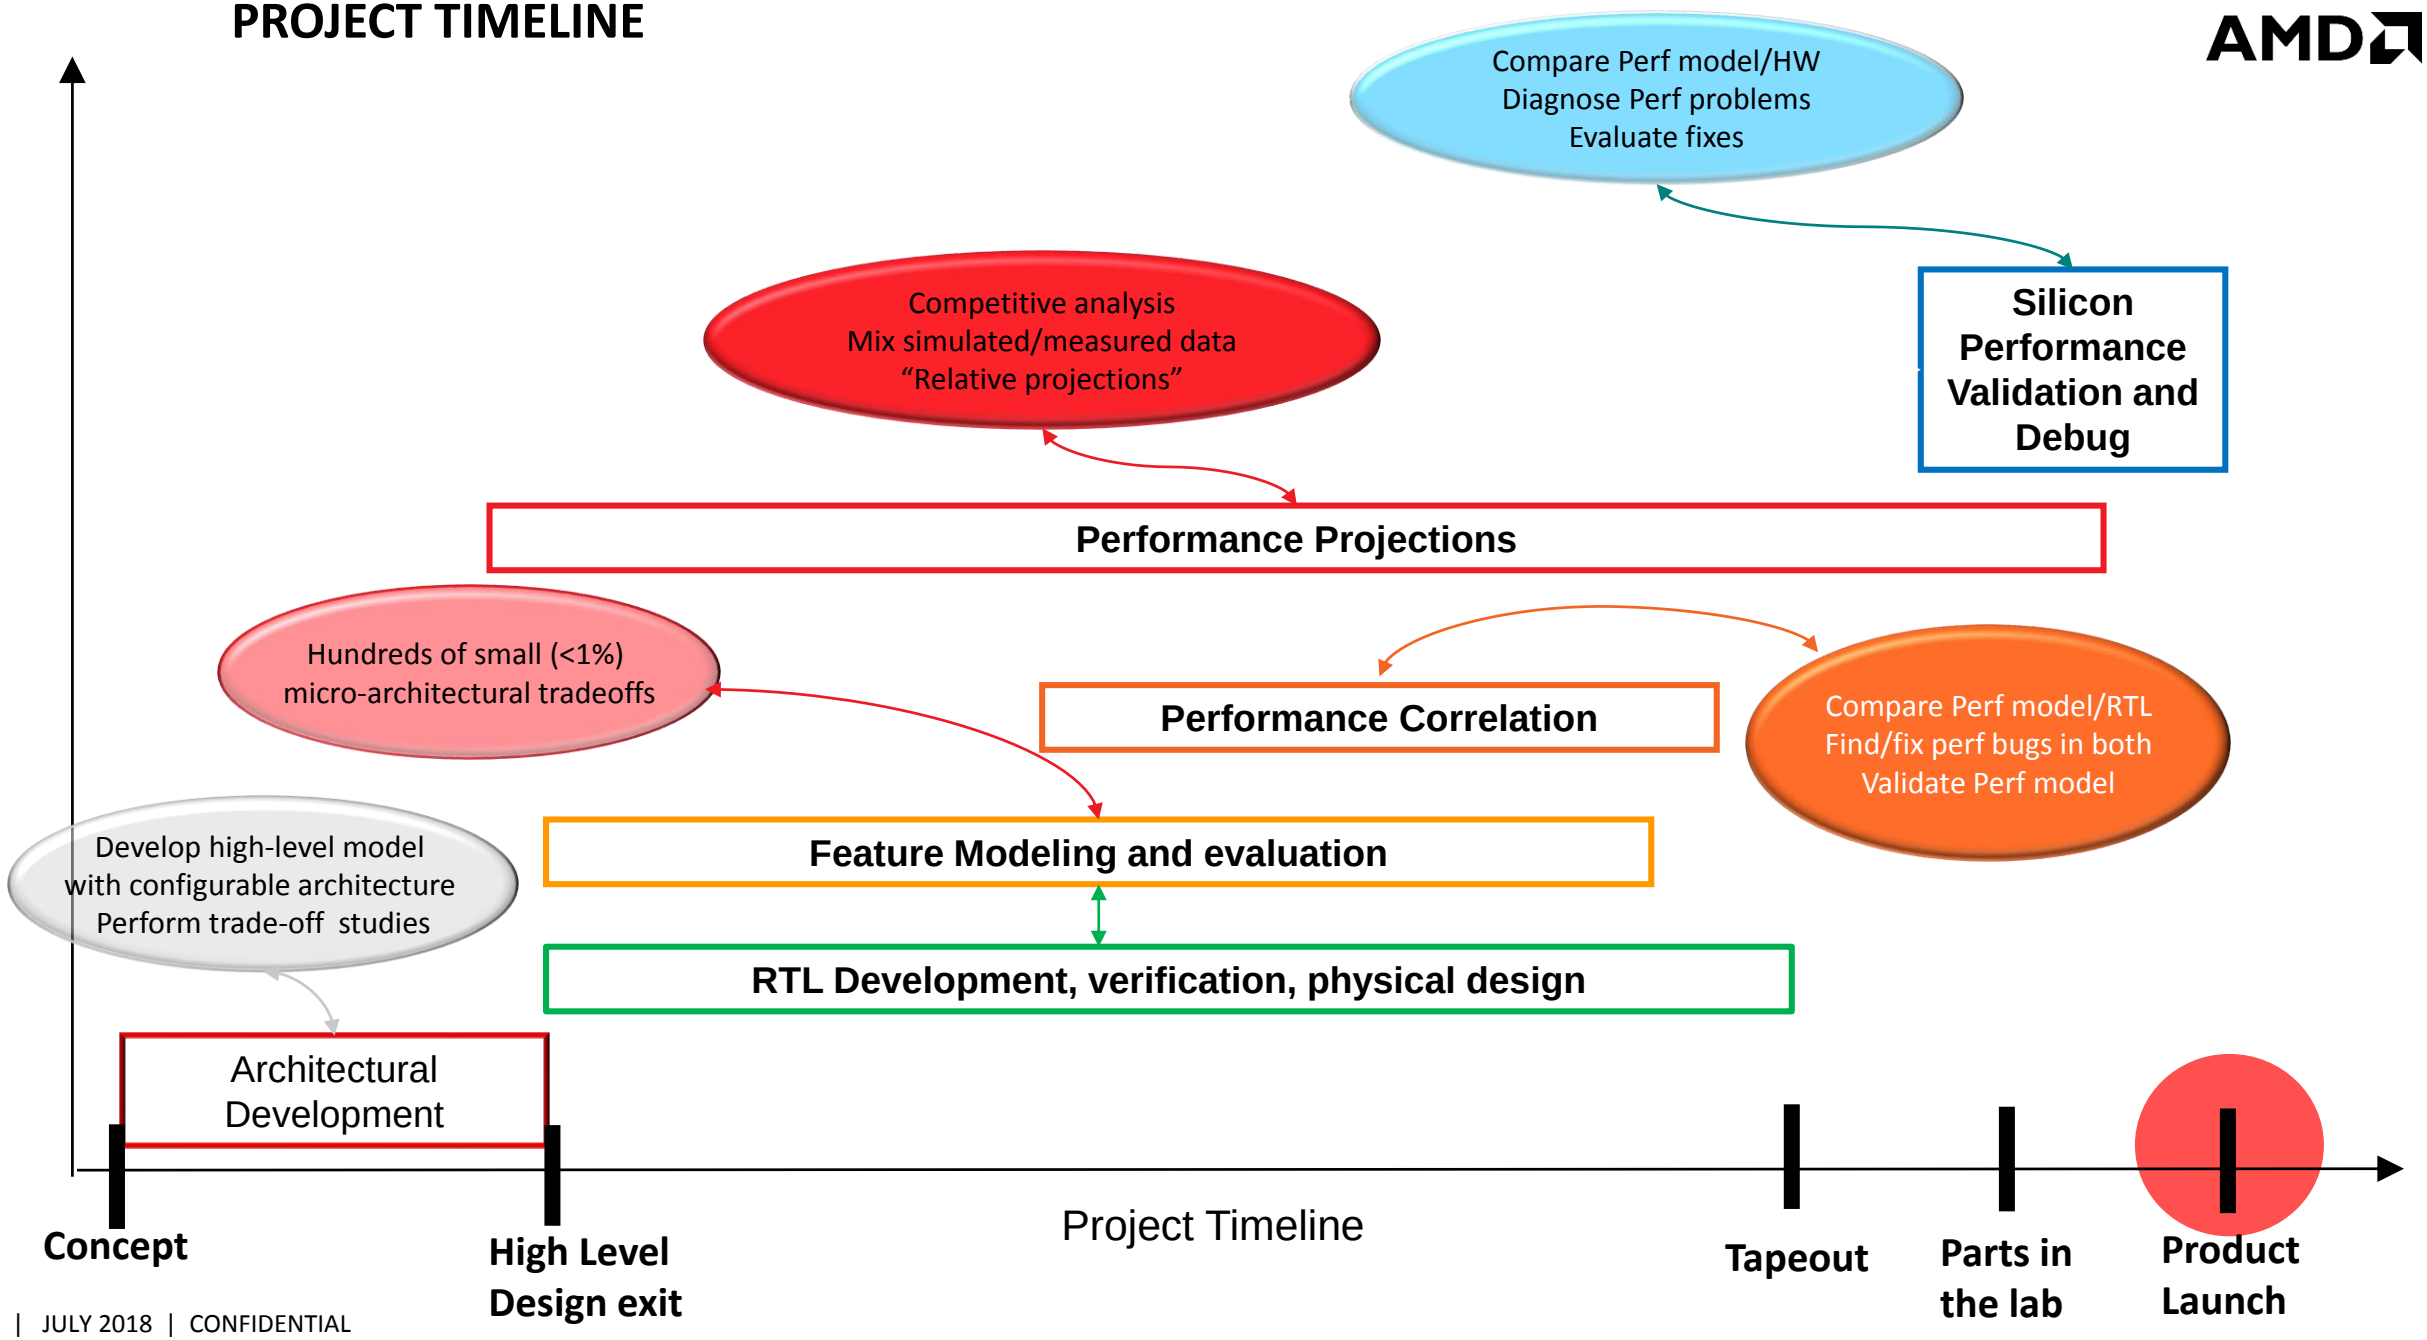

PERFORMANCE ANALYSIS AND MODELING

AMD

- **Develop and explore micro-architectural ideas for performance improvement**

- Variety of tools and technologies deployed for performance analysis & architectural innovation
  - **Cycle accurate simulator**
- Constant state of evolution, due to changing workloads, and architecture
  - More IPC, MP, SMT, large caches
  - SW and HW more tightly linked than ever
- Correct performance projection is critical for the success of the project
- Workloads analysis, characterization and tracing
- Server workloads – SPEC CPU 2006, SPEC CPU2017, Legacy Enterprise workloads, Spec JBB, Cloud, NoSql database
- Client workloads – Games, Cinebench, Productivity suite, various use case scenarios

DISCLAIMER & ATTRIBUTION

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION

AMD X86 CORES: DRIVING COMPETITIVE PERFORMANCE

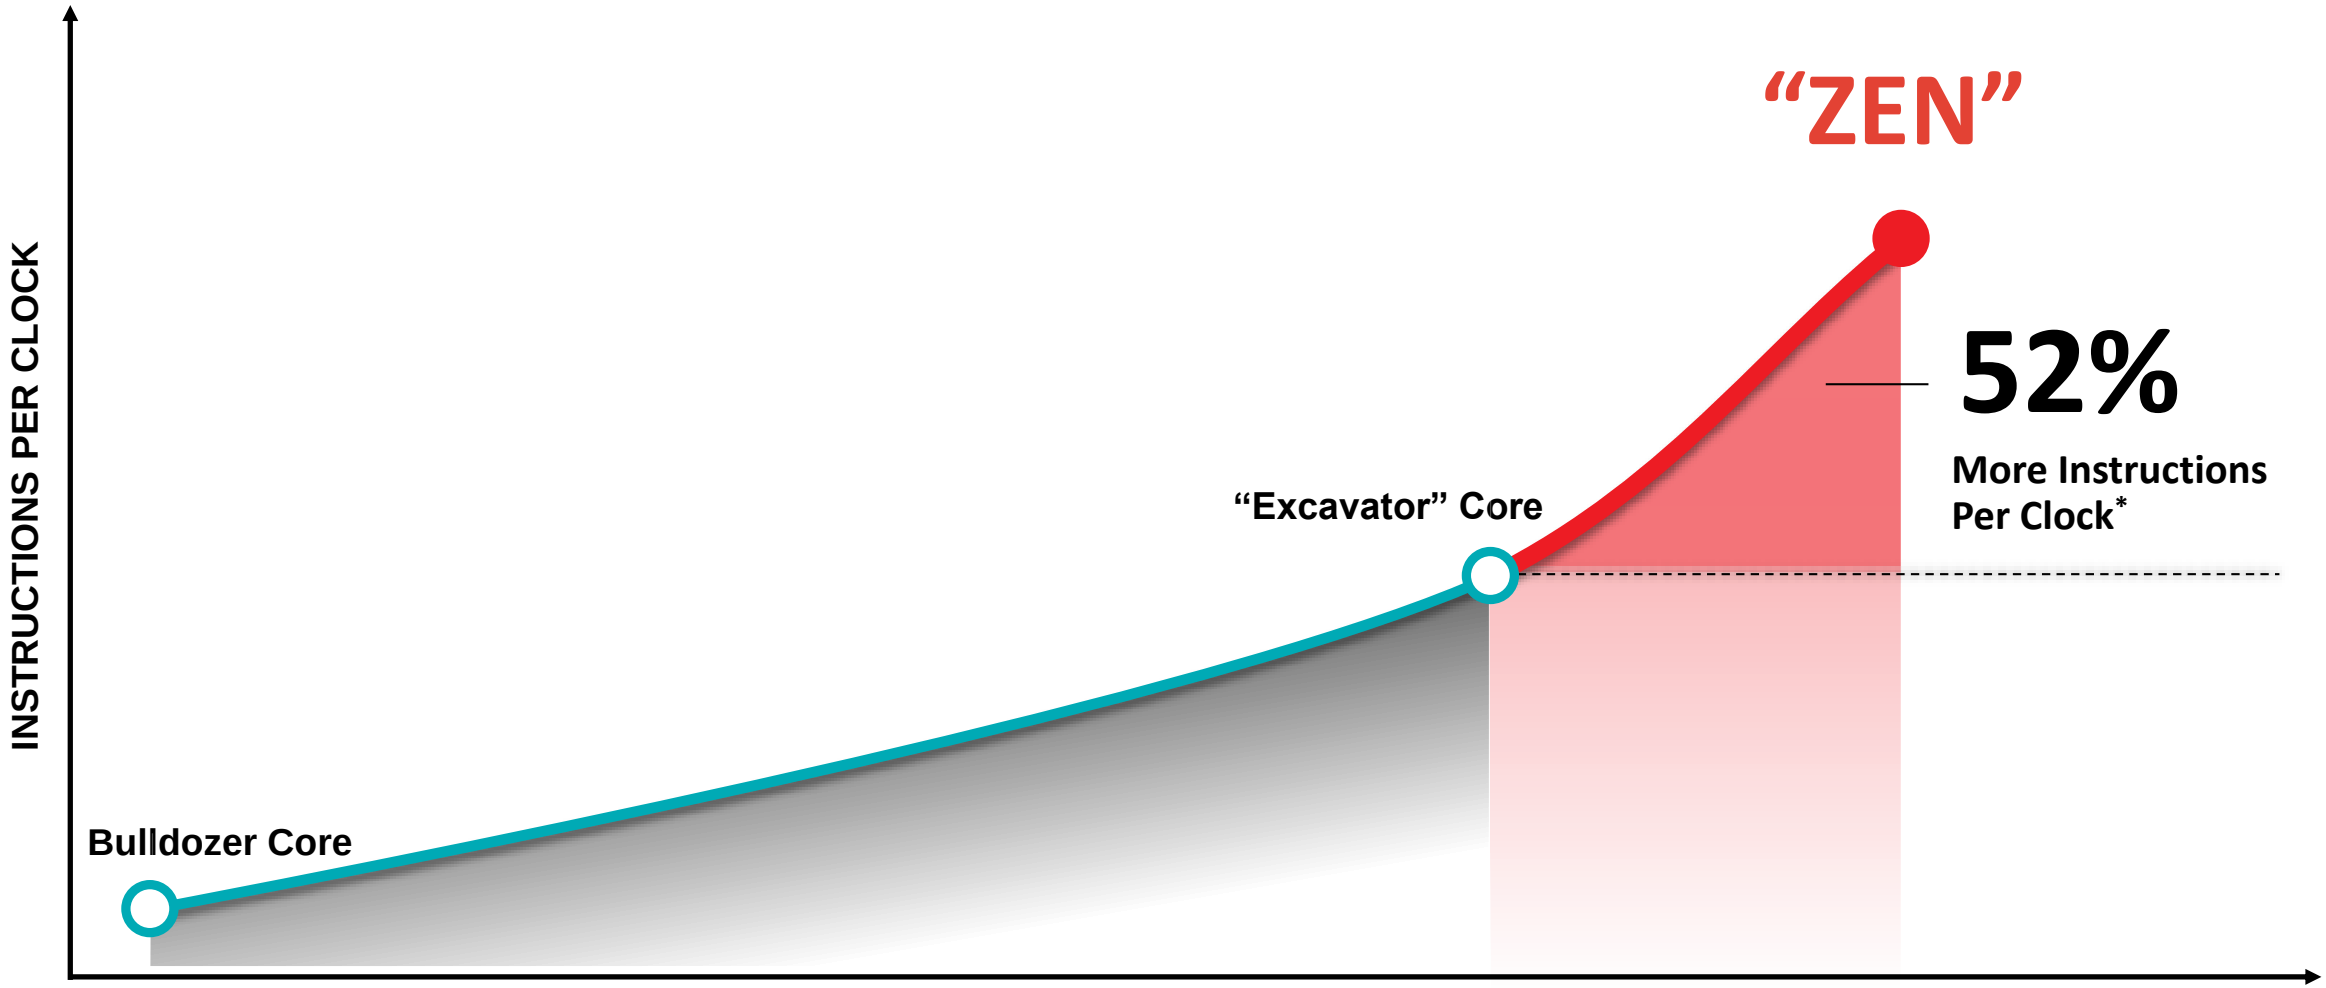

*Based on internal AMD estimates for "Zen" x86 CPU core compared to "Excavator" x86 CPU core.

PERFORMANCE ANALYSIS TOOLS