“ZEN 2”

David Suggs | Dan Bouvier

Michael Clark, Kevin Lepak, Mahesh Subramony

August 19, 2019

THE PATH TO "ZEN 2"

CHALLENGES

- Meet multiple markets – server to mobile
- Ultimate performance
- Energy efficient performance

TECHNOLOGY

- Enable scale and power
- Balancing complexity of new technology node with time-to-market

PLATFORM

- Chiplets to deliver right technology to the needs
- Upgrade IO to meet system demands
- Compatibility

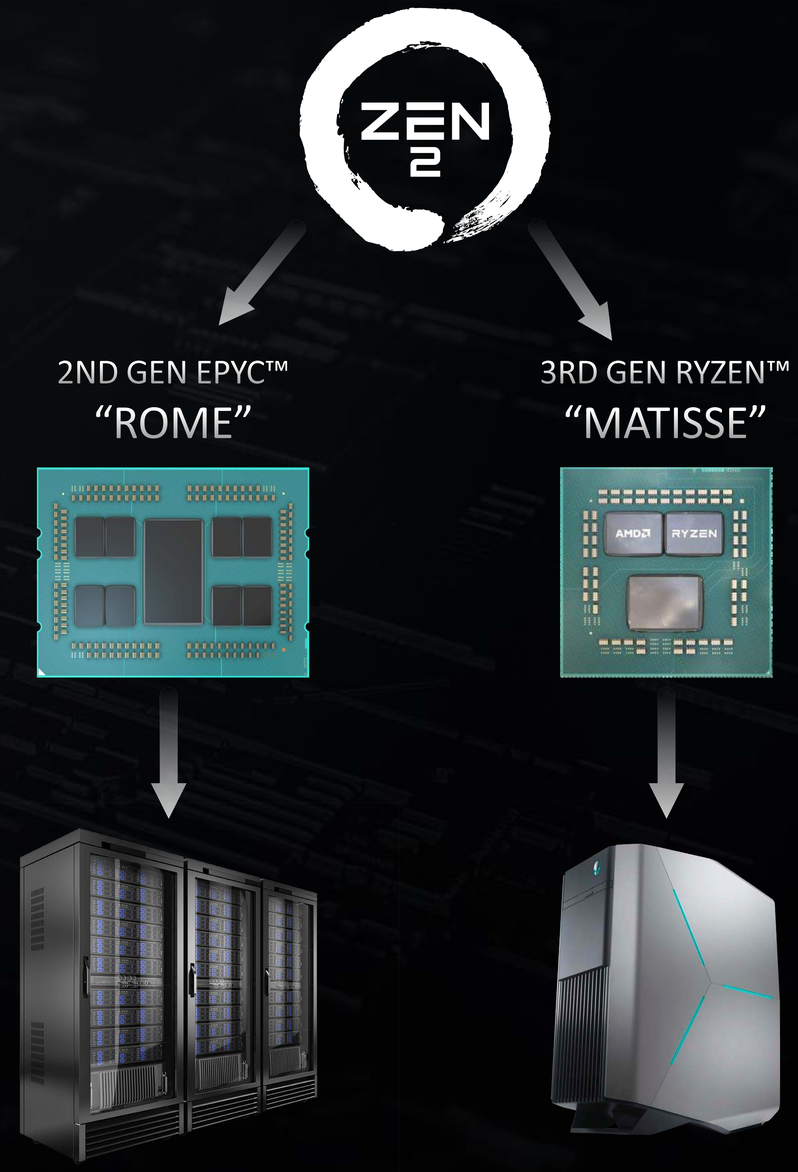

AMD

“ZEN 2”

MICROARCHITECTURAL HIGHLIGHTS

- New TAGE branch predictor
- 2x op cache capacity
- Reoptimized L1I cache
- 3rd address generation unit
- 2x FP data path width
- 3x L1 load+store bandwidth
- 2x L3 capacity
- Improved prefetch throttling
- 2 threads per core (SMT) carried forward

15% IPC IMPROVEMENT FROM “ZEN” TO “ZEN 2”

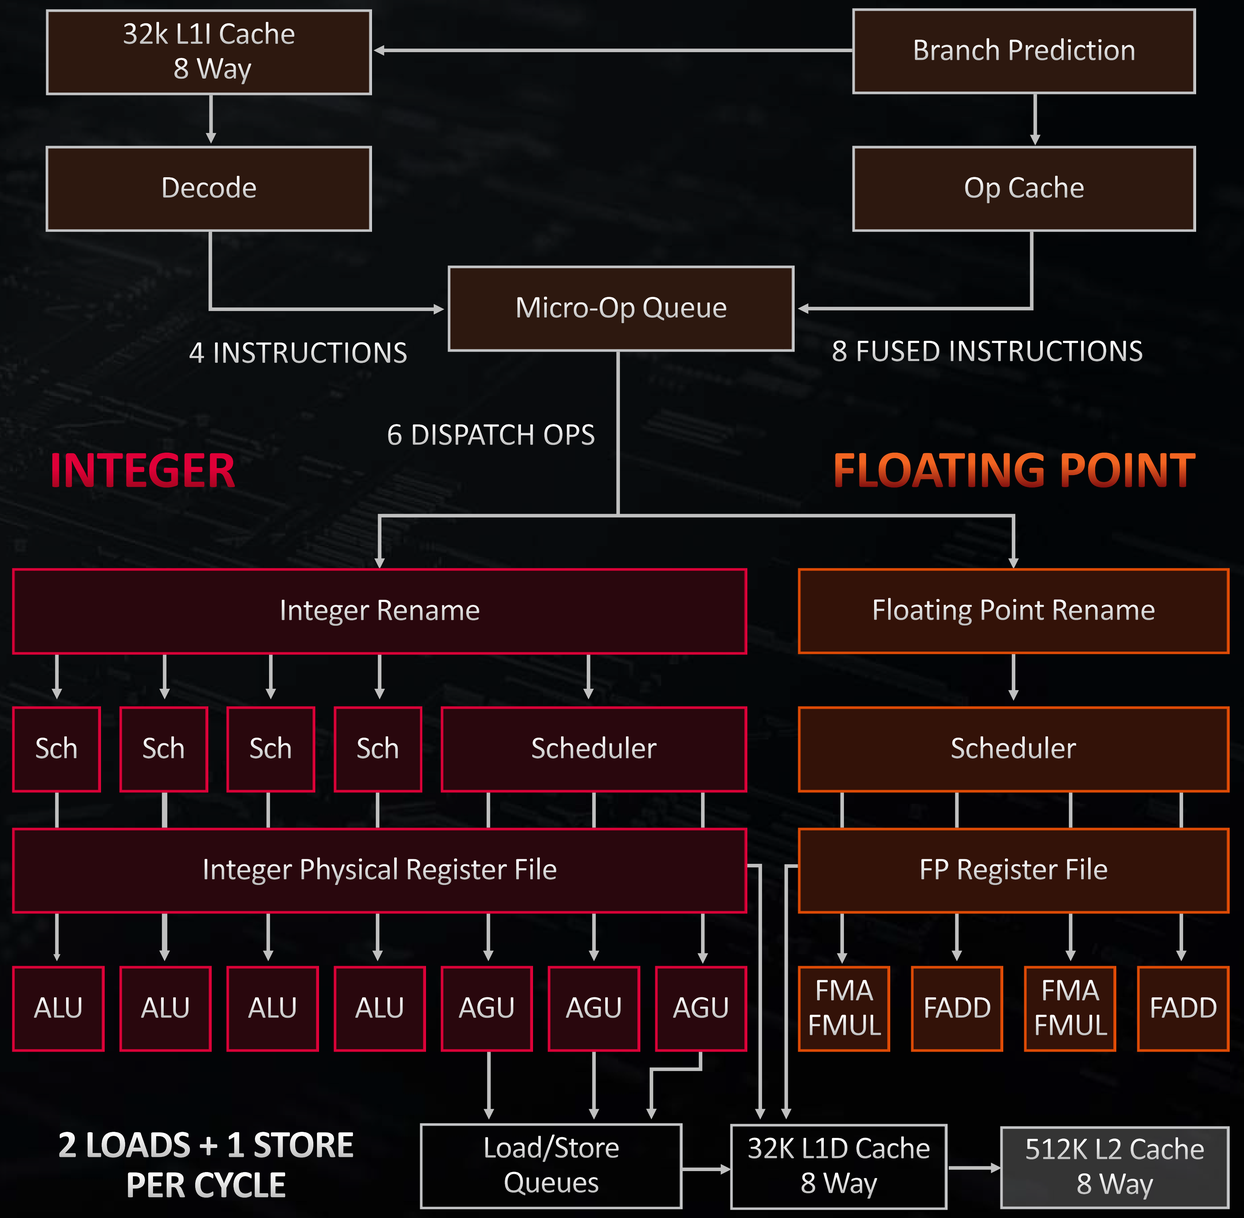

AMD

MAJOR EFFICIENCY IMPROVEMENTS

2x INSTRUCTIONS PER UNIT OF ENERGY

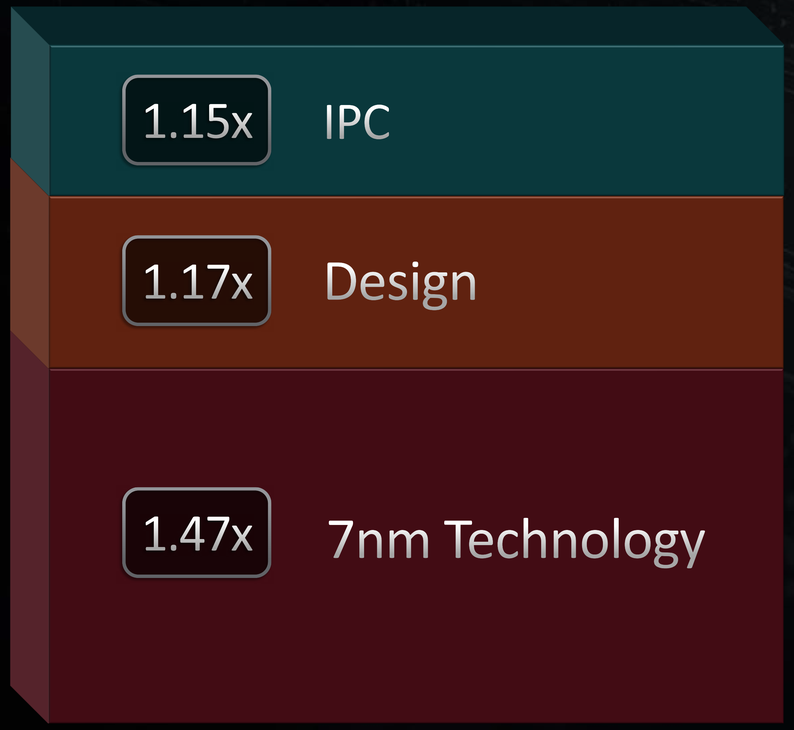

- Improved branch prediction accuracy
- Higher op cache hit rate
- New integer scheduler algorithms
- Clock and data gating improvements
- Low-power design methodology

DESIGNED FOR SECURITY

HARDWARE OPTIMIZED SECURITY MITIGATIONS

Spectre v2*	Indirect branch target injection
Spectre v4**	Speculative store bypass

SECURE MICROARCHITECTURE

Hardware permission checks
Permissions are checked prior to consuming data

SECURE VIRTUALIZATION

SEV-ES	SEV with encrypted state
More SEV Keys	SEV with up to 509 encrypted guests (was 15)
SEV-VTE	SEV with virtual transparent encryption
GMET	Guest mode execute trap. Hypervisor can trap guest supervisor-mode execution of pages for integrity checking

EXTENDING THE ARCHITECTURE

FEATURE	NOTES	"ZEN"	"ZEN 2"
x2APIC	APIC extension for high core count system support		✓
QOS	Quality-of-service monitoring/enforcement of L3 cache occupancy and memory bandwidth		✓
UMIP	User mode instruction prevention (CR4.UMIP controls access to SGDT, SIDT, SLDT, STR, and SMSW)		✓
CLWB	Non-volatile memory enhancement (cache-line writeback)		✓
WBNOINVD	Cleans caches, but does not invalidate		✓
RDPRU	Read processor register at a user level (currently MPERF and APERF)		✓

AMD Exclusive

PREDICTION, FETCH, AND DECODE

GREATER ACCURACY AND HIGHER CAPACITY

	"ZEN"	"ZEN 2"
IMPROVED BRANCH PREDICTION		
L2 Predictor	Perceptron	**TAGE**
2x L0 BTB	8	**16**
2x L1 BTB	256	**512**
1.75x L2 BTB	4K	**7K**
2X Indirect Target Array	0.5K	**1K**
30% Lower Mispredict Rate Target		✓
OPTIMIZED L1I CACHE	64K	**32K**
2x Associativity	4-way	**8-way**
Improved Utilization		✓
2X OP CACHE (FUSED INSTR)	2K	**4K**
Better Instruction Fusion		✓
Better Effective Throughput		✓

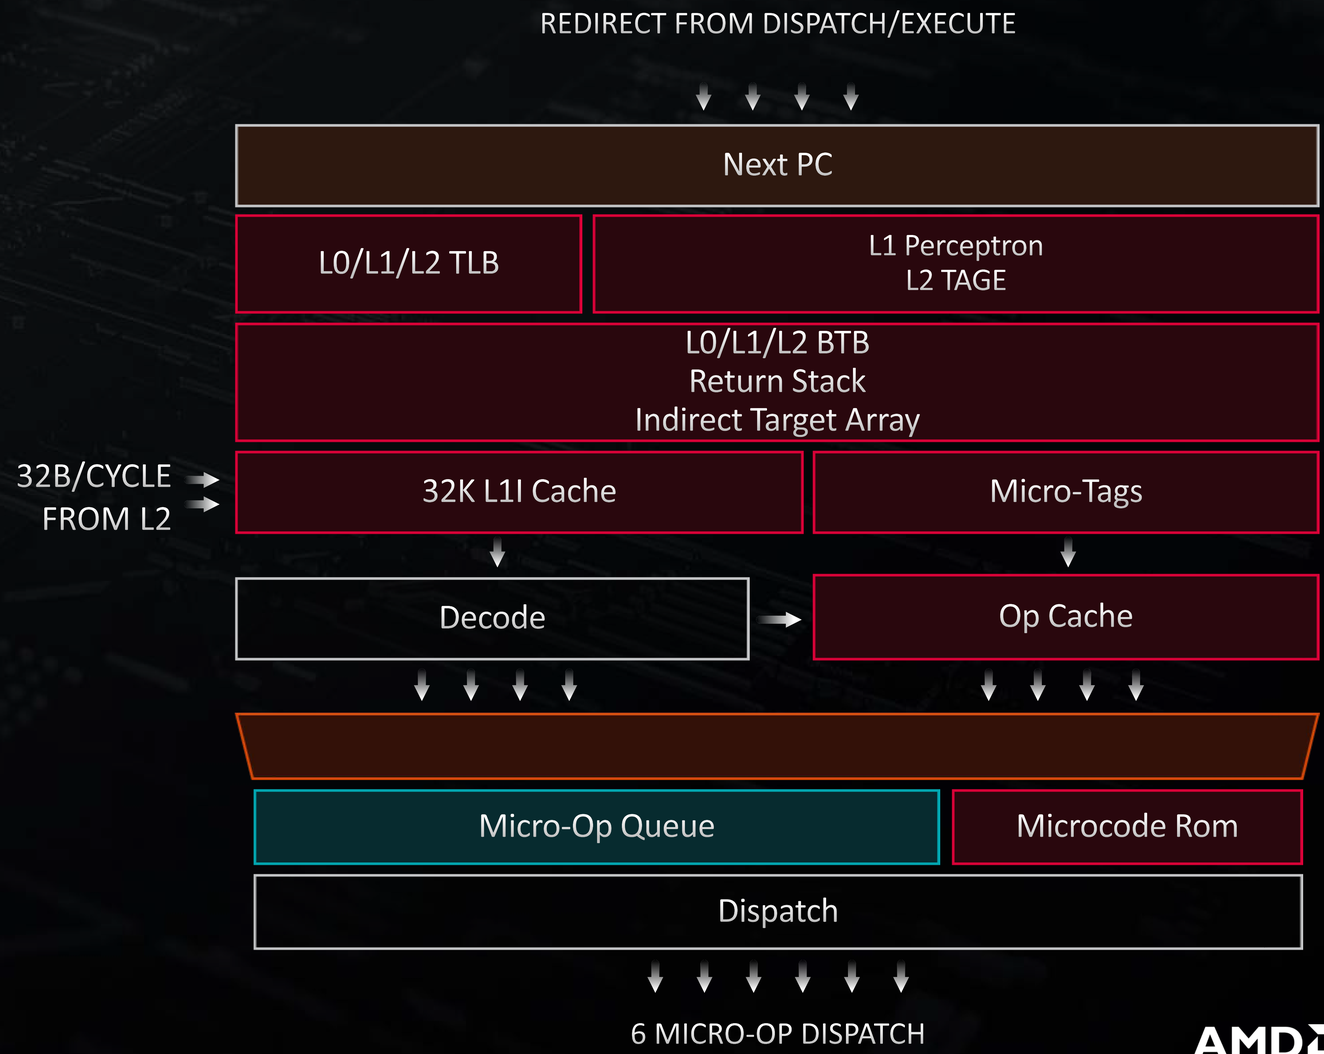

AMD

INTEGER EXECUTE

WIDER ISSUE AND DEEPER WINDOW

	"ZEN"	"ZEN 2"
WIDER ISSUE	6	**7**
ALUs	4	**4**
AGUs	2	**3**
DEEPER WINDOW		
Bigger ALU Scheduler	4x14	**4x16**
More Unified AGU Scheduler	2x14	**1x28**
Bigger Register File	168	**180**
Bigger Reorder Buffer	192	**224**
IMPROVED SMT FAIRNESS FOR ALU AND AGU SCHEDULERS		✓

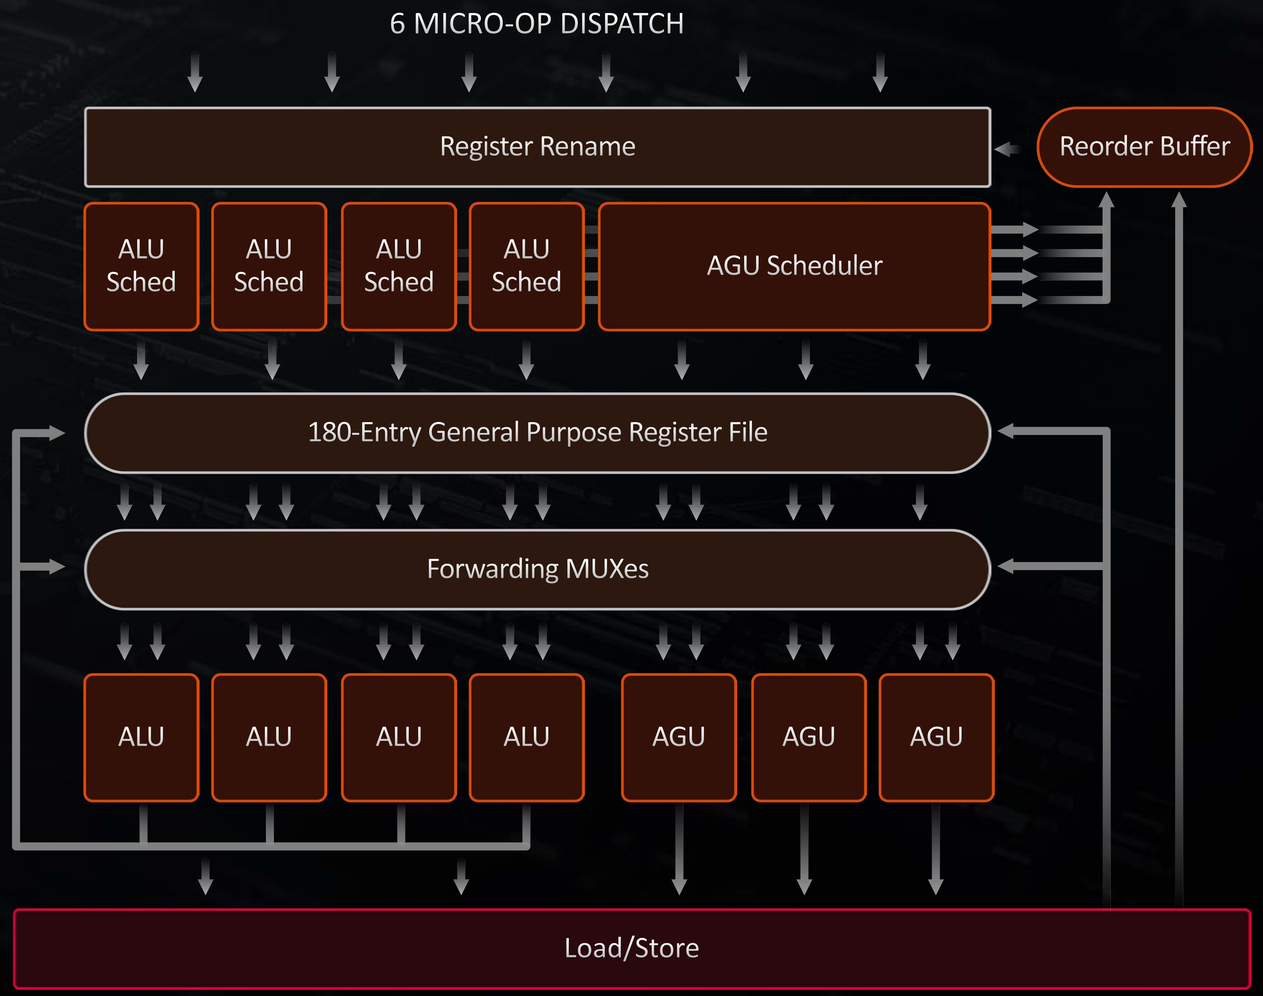

AMD

FLOATING-POINT/VECTOR EXECUTE

DOUBLE WIDE VECTORS

	"ZEN"	"ZEN 2"
AVX256 INSTRUCTION SUPPORT	✓	✓
2X WIDTH DATA PATH	128b	**256b**
EDC Management		✓
2X WIDTH VECTOR REGISTER FILE	128b	**256b**
2X WIDTH LOADS (2)	128b	**256b**
2X WIDTH STORES (1)	128b	**256b**
IMPROVED DOUBLE-PRECISION MULTIPLY LATENCY	4 cyc	**3 cyc**

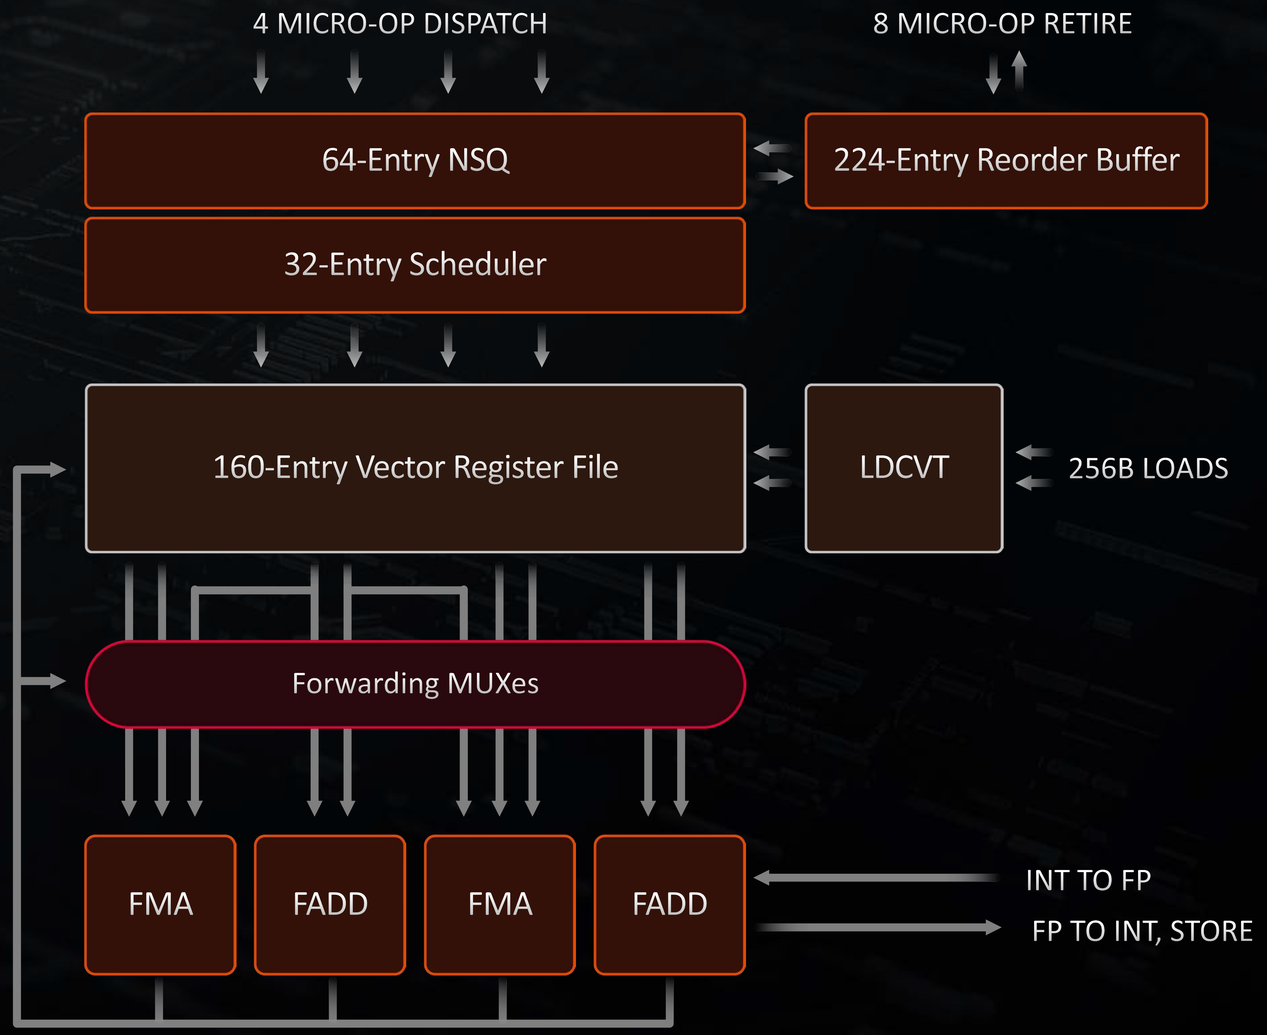

AMD

LOAD/STORE AND L1D CACHE

MORE THROUGHPUT AND BIGGER STRUCTURES

	"ZEN"	"ZEN 2"
BIGGER STORE QUEUE	44	**48**
BIGGER, BETTER L2 DTLB	1.5K	**2K**
1G Page Support (as 2M)		✓
Lower Latency	8 cyc	**7 cyc**
32KB, 8-WAY L1D CACHE		
2X Width Reads (2)	128b	**256b**
2X Width Writes (1)	128b	**256b**
3X Load + Store Bandwidth	32B/clk	**96B/clk**
IMPROVED WRITE-COMBINING BUFFER PERFORMANCE		✓
IMPROVED PREFETCH THROTTLING		✓

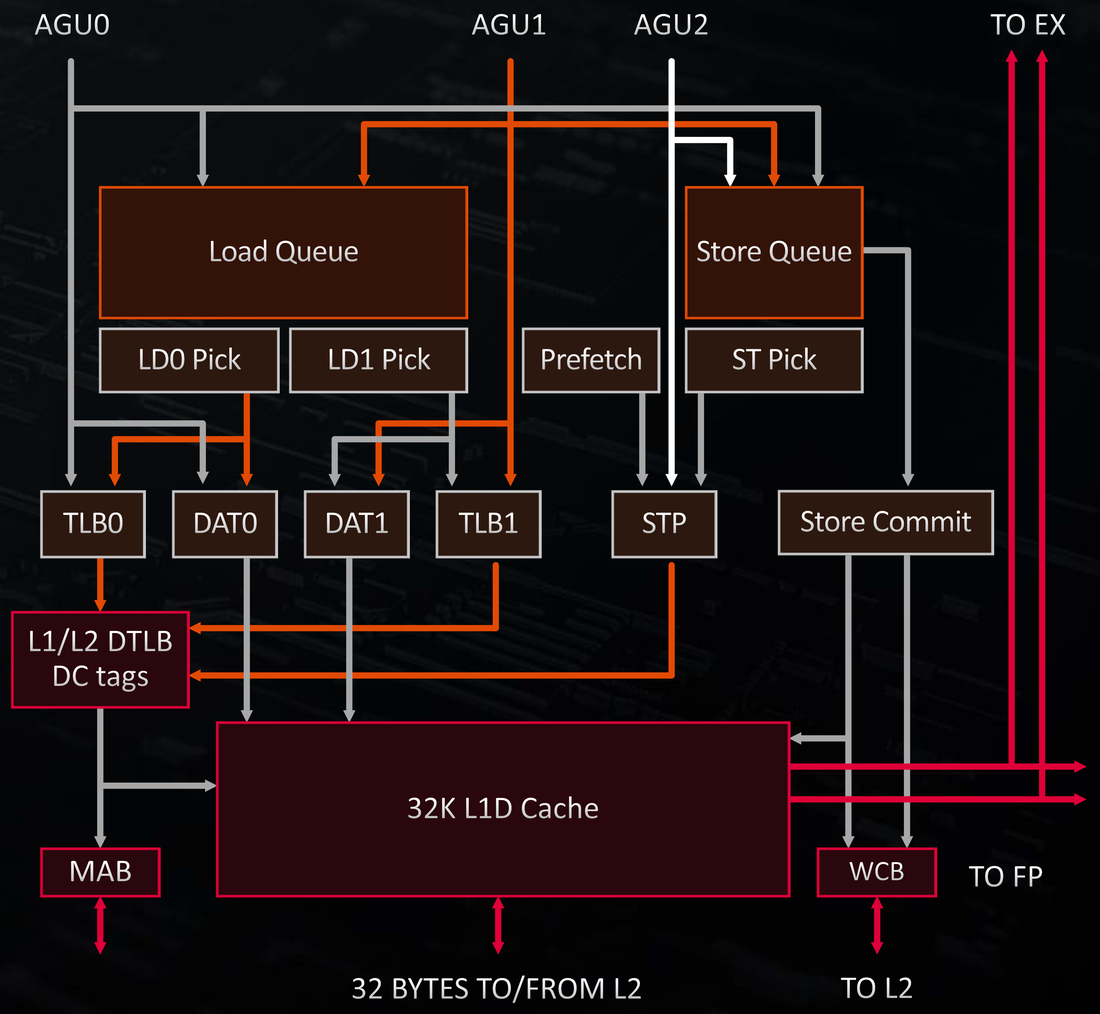

AMD

"ZEN 2" SOLUTIONS

3rd Generation Ryzen™ ("Matisse")

2nd Generation Epyc™ ("Rome")

CACHE HIERARCHY AND CPU COMPLEX

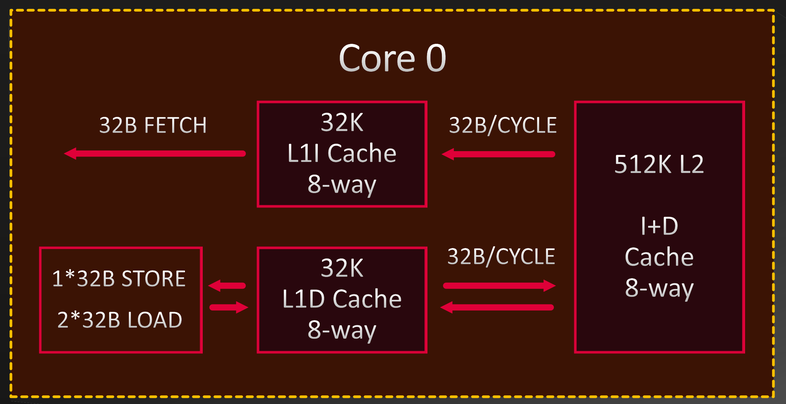

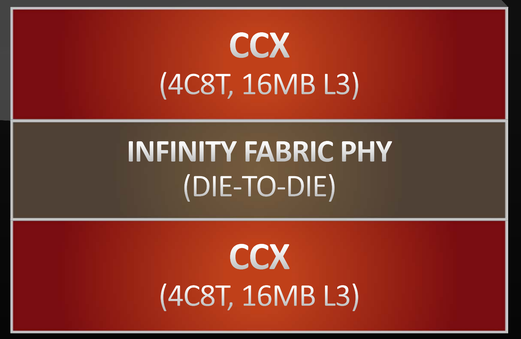

- 2X L1D CACHE LOAD/STORE BANDWIDTH *vs.* "ZEN"
  - *32B Per Cycle Everywhere*
- 2X L3 CACHE UP TO 16MB PER CCX, 2CCXs per CCD

AMD

AMD GOAL

DEPLOY LEADERSHIP PERFORMANCE SIMULTANEOUSLY INTO MULTIPLE MARKETS

CHALLENGES

- Higher cost per mm^2
- Complex analog/IO elements do not scale well and Can be challenging to port
- Wide range of diverse markets

SOLUTION | CHIPLET STRATEGY

- Small die improve yield
- Analog / IO IP in 12nm technology to reduce cost and complexity
- Enables product configurability

Next Generation Highly Optimized “Zen 2” Core

Leading Edge 7nm Technology for Density and Power Efficiency

Timely Delivery Of Products

REVOLUTIONARY CHIPLET DESIGN

Each IP in its Optimal Technology

Infinity Fabric™ Enables Modularity (MCM), Scaling (CCD Count)

Optimized I/O Die Enables Common Latency to All Cores/Caches

AMD

PUTTING IT TOGETHER
PERFORMANCE DESKTOP "MATISSE"

- A Desktop SoC and a Chipset
  - IOD combined with CCD(s) form the CPU
  - Standalone IOD re-purposed as Chipset
- Leading I/O
  - 48GB/s native PCIe™ BW
  - 4 USB 3.1 10Gb/s ports
- Memory BW
  - 51.2 GB/s Memory BW
  - Dual Channel DDR4 3200 MT/s
- Overclocking
  - Improved Memory overclocking (Phy, Package)
  - De-coupled various IO-die clocks for flexibility
- AM4 Platform Longevity
  - Compatible with AM4 platform

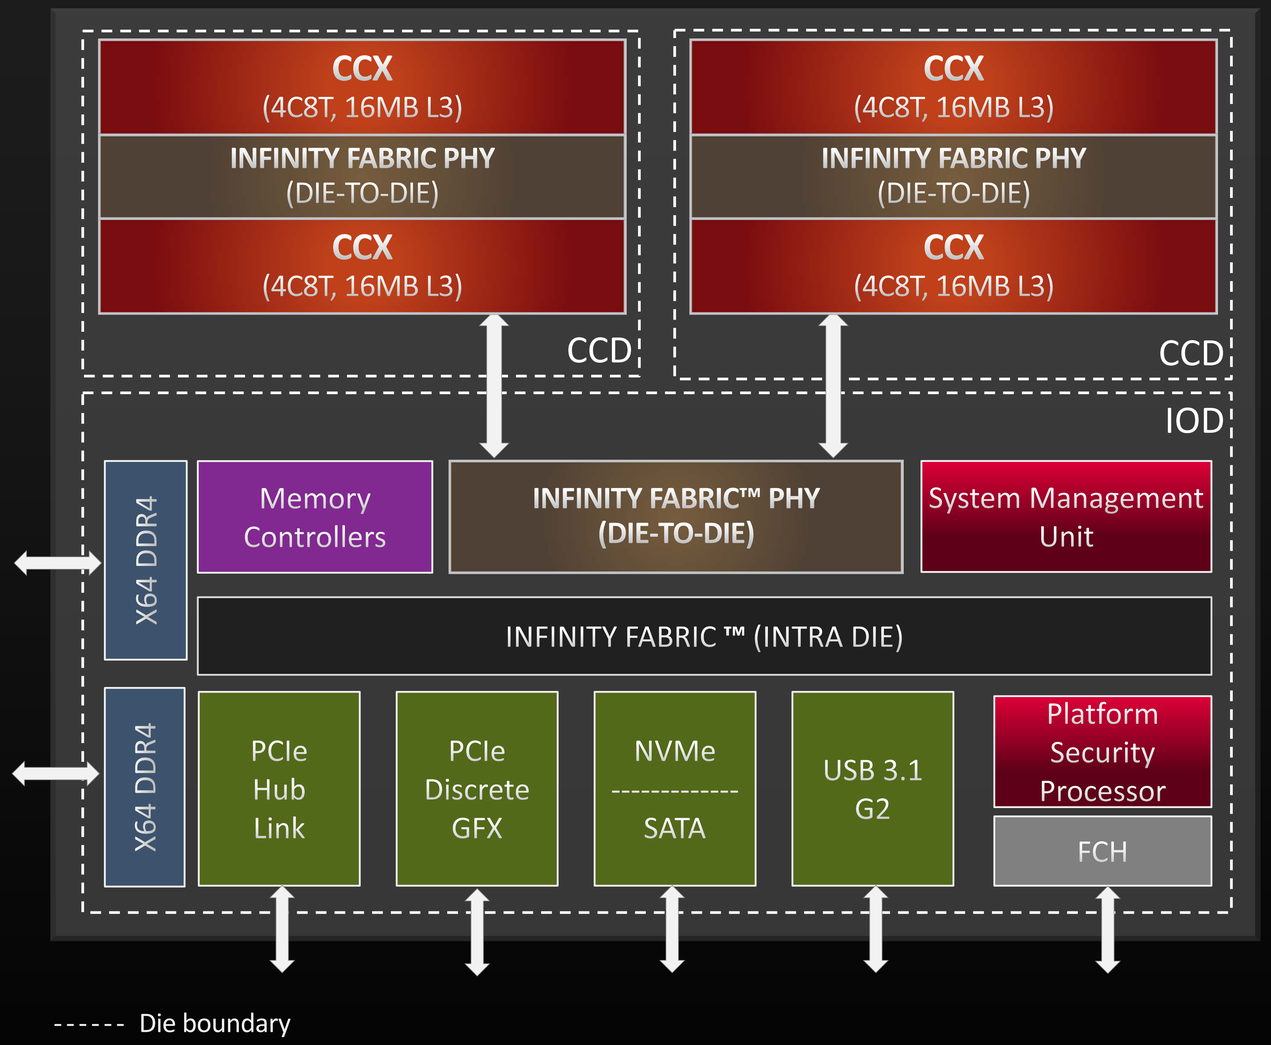

AMD

SIGNIFICANT DESKTOP PERFORMANCE IMPROVEMENT

Application Performance In Power Efficient TDP

- Compute heavy workloads benefit from power efficient design
- Content Creation, Rendering benchmarks see large gains
- Additional cores in 3900X deliver substantial compute power

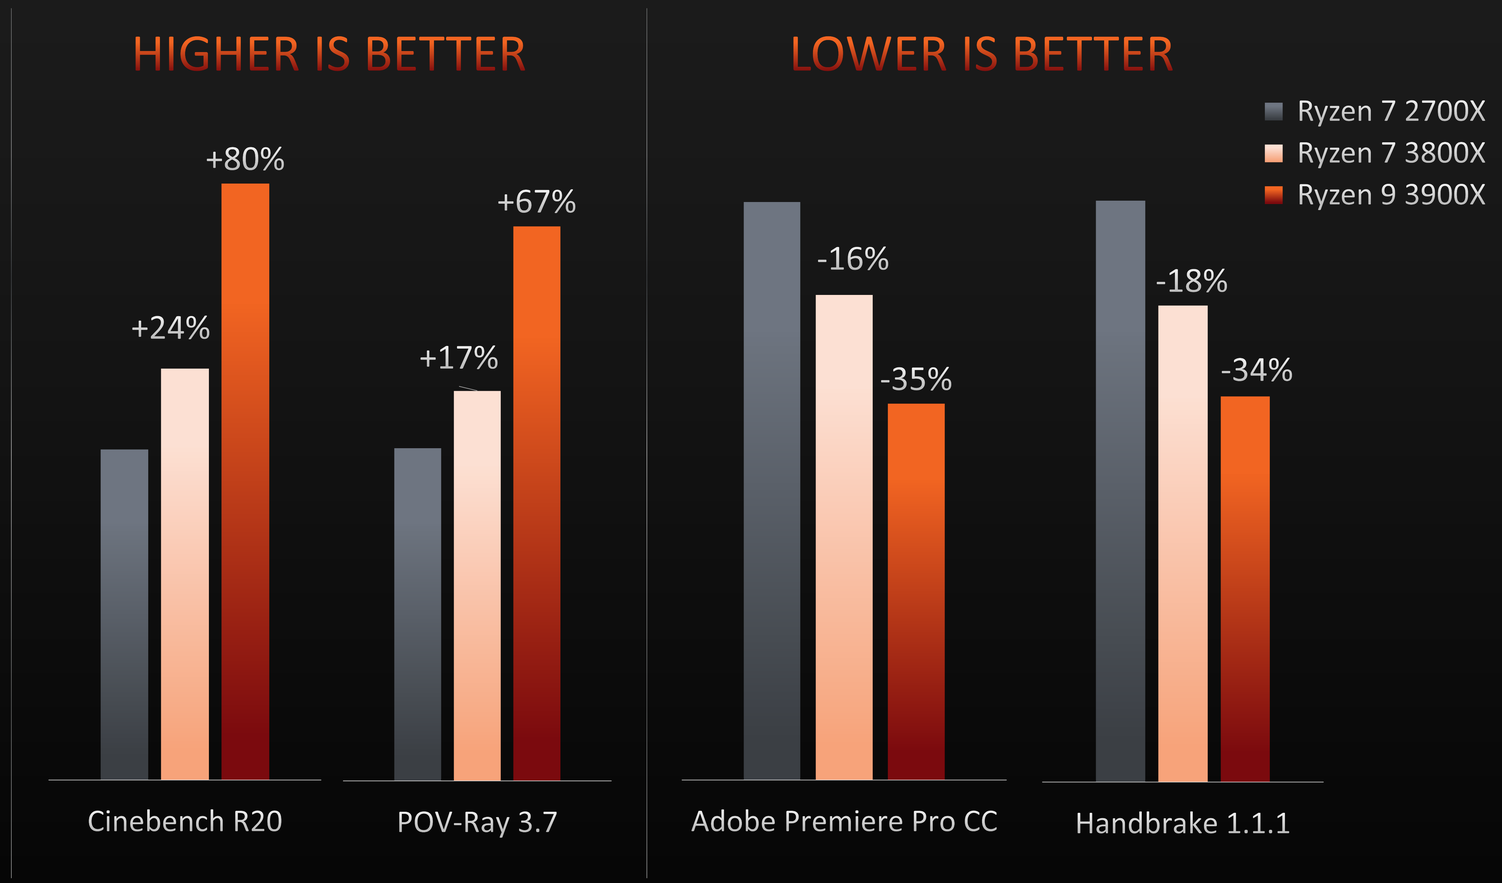

AMD

SIGNIFICANT DESKTOP PERFORMANCE IMPROVEMENT

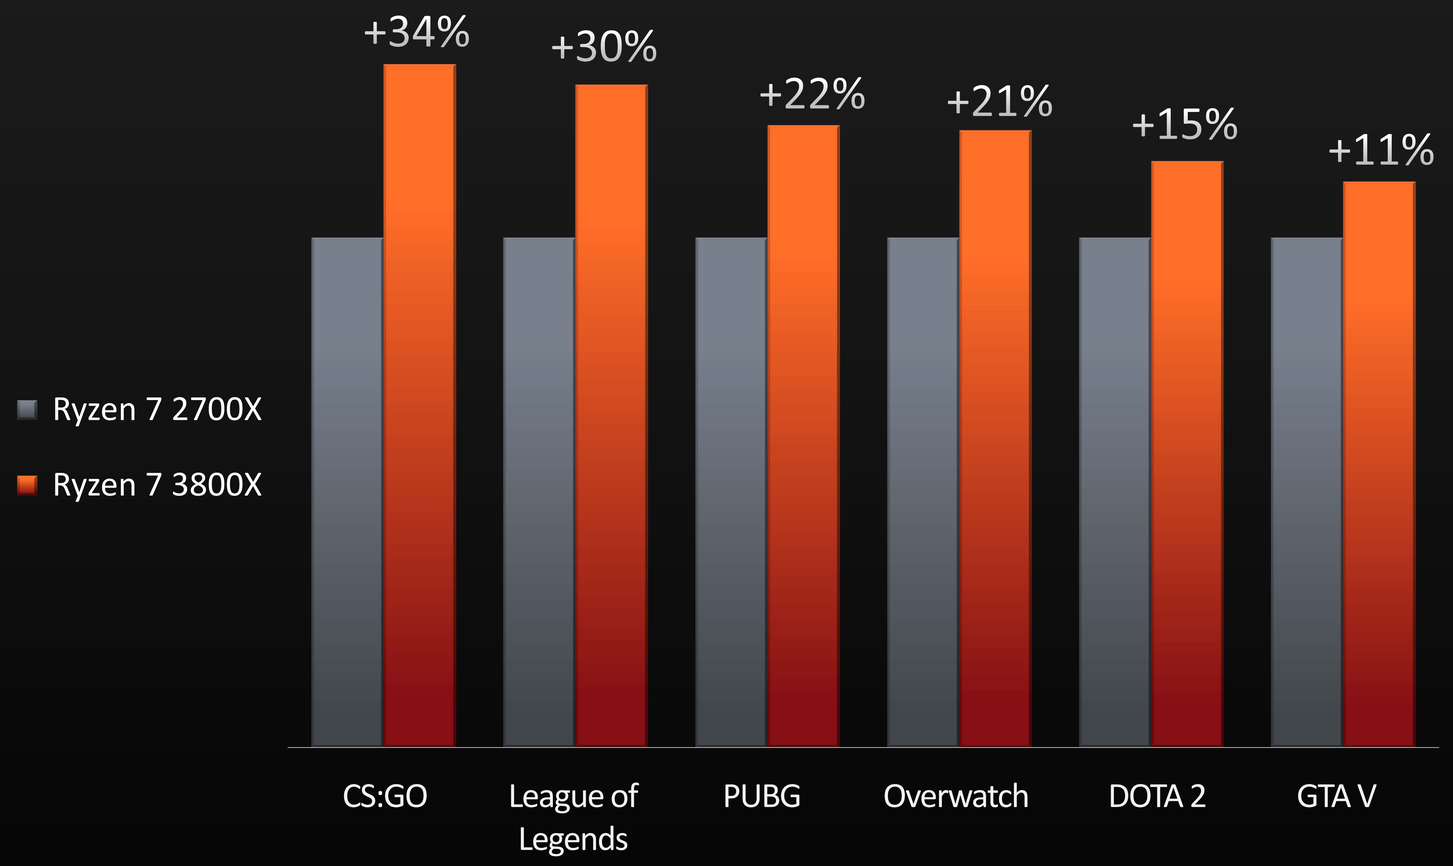

Gaming Performance

- Significant generational performance improvement
- Increased L3 (up to 64MB) reduces effective memory latency

AMD

APPLICATION BENEFITS FROM PCIE™ GEN4

IO Performance

- Up to 2x PCIe BW vs prior Ryzen generation

Storage Performance

- Large Block Sequential accesses are severely limited by link speeds

3DMark® PCIe Feature Test

- Vertex animation (game VFX) is sensitive to bus bandwidth, allowing significant upside

DaVinci Resolve

- Bus bandwidth is a significant limiting factor for non-linear editing (NLE) performance

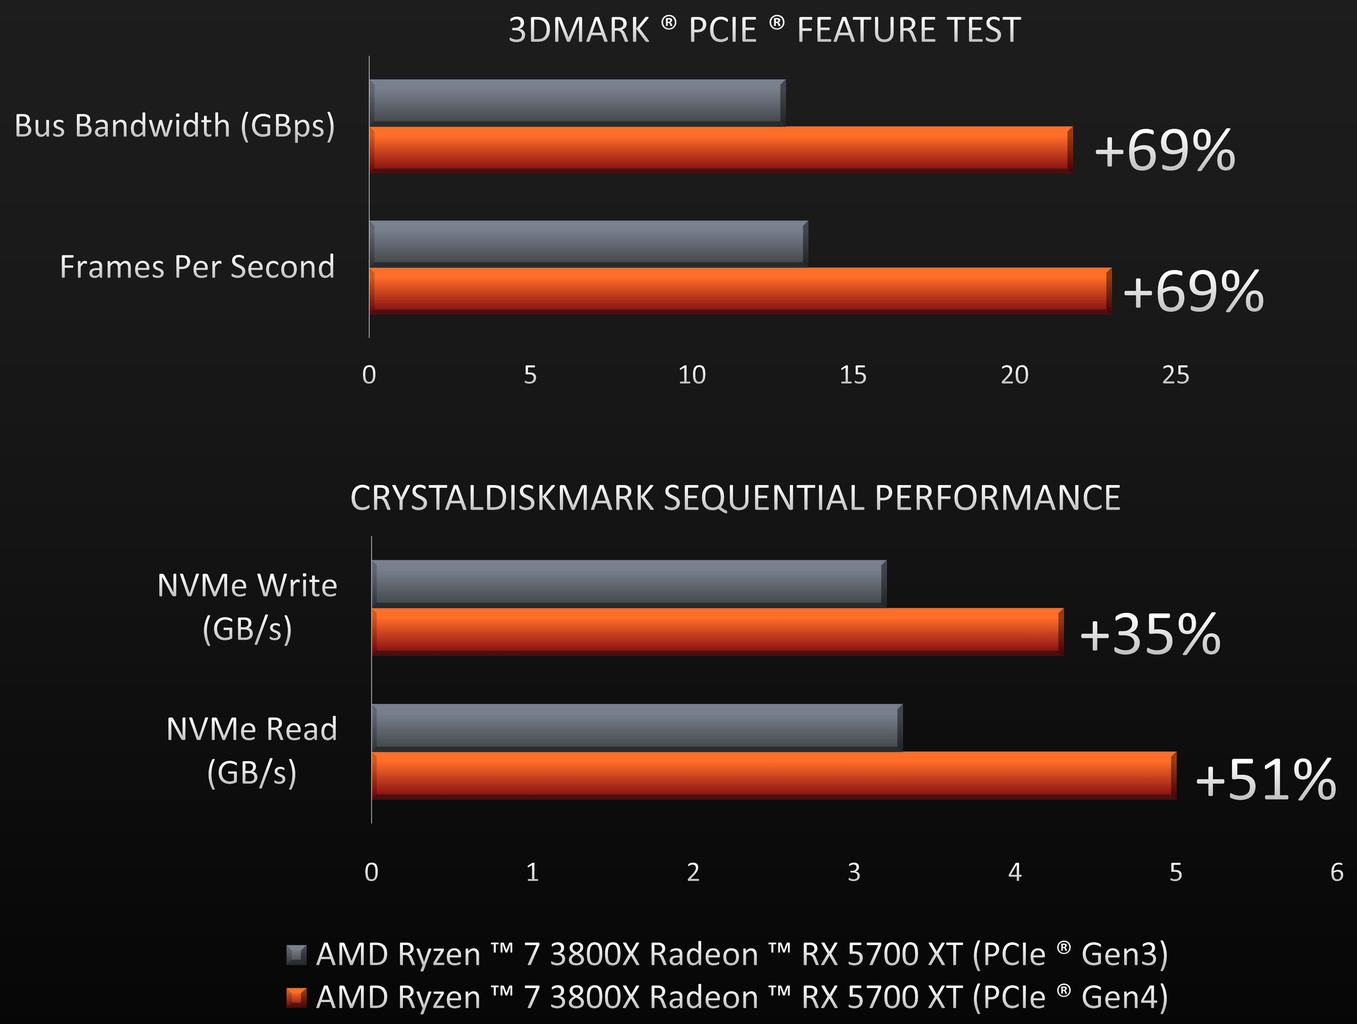

AMD

PUTTING IT TOGETHER
SERVER "ROME"

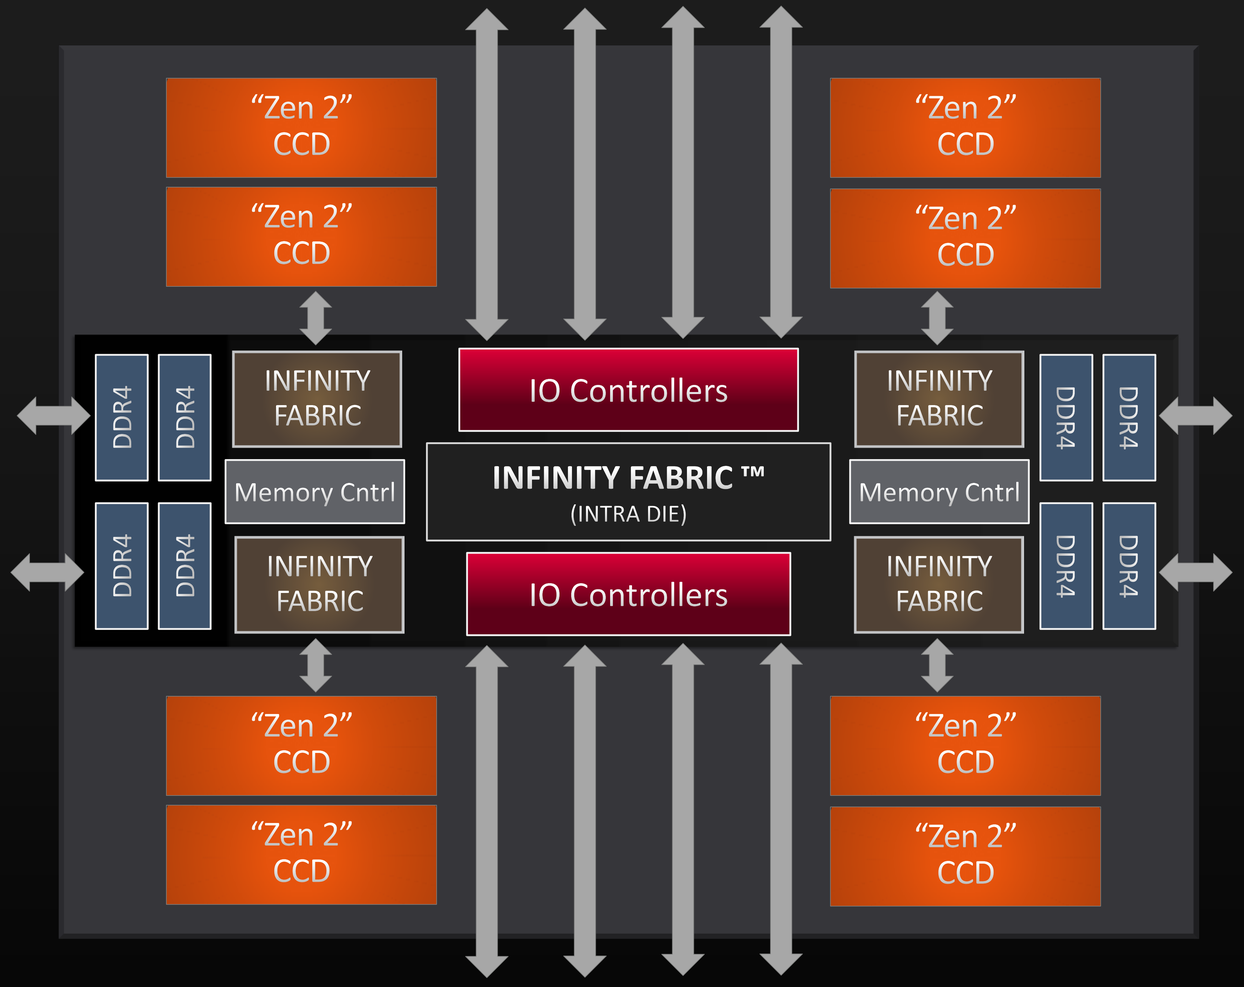

"Rome" CPU

- ~1000mm^2 cumulative area in leading-edge 7nm/12nm design
- Up-to-64 cores (8 CCDs) and 256MB L3-cache
- Improved NUMA
  - All memory and I/O hosted by single die
  - Single IOD results in effective latency improvements
- Class-Leading I/O and Memory Bandwidth
  - 8 Memory Channels 64b DDR4-3200 204.8 GB/s Bandwidth
  - PCIe™ Gen 4
- Platform Longevity; Compatible with Prior Generation EPYC™
- Advanced Platform Security Features: Expanded SEV keys, SEV-IO, SEV-ES

AMD

TWO SOCKET LEADERSHIP

2S INTEL® XEON® VS. 2S AMD EPYC™ SPEC CPU® 2017 PERFORMANCE

- Up to 128 cores (8 CCDs) and 256MB L3-cache per socket

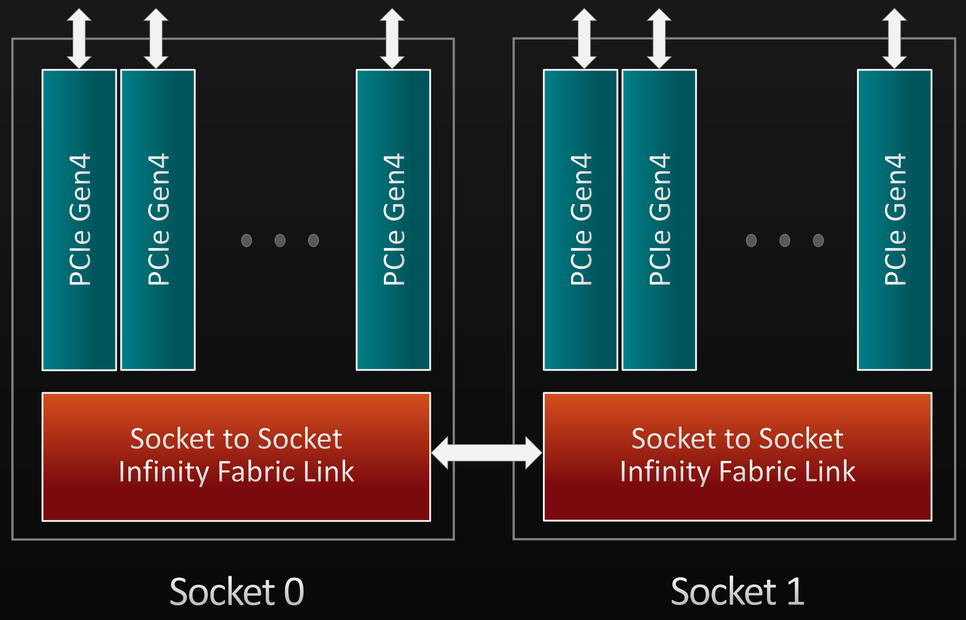

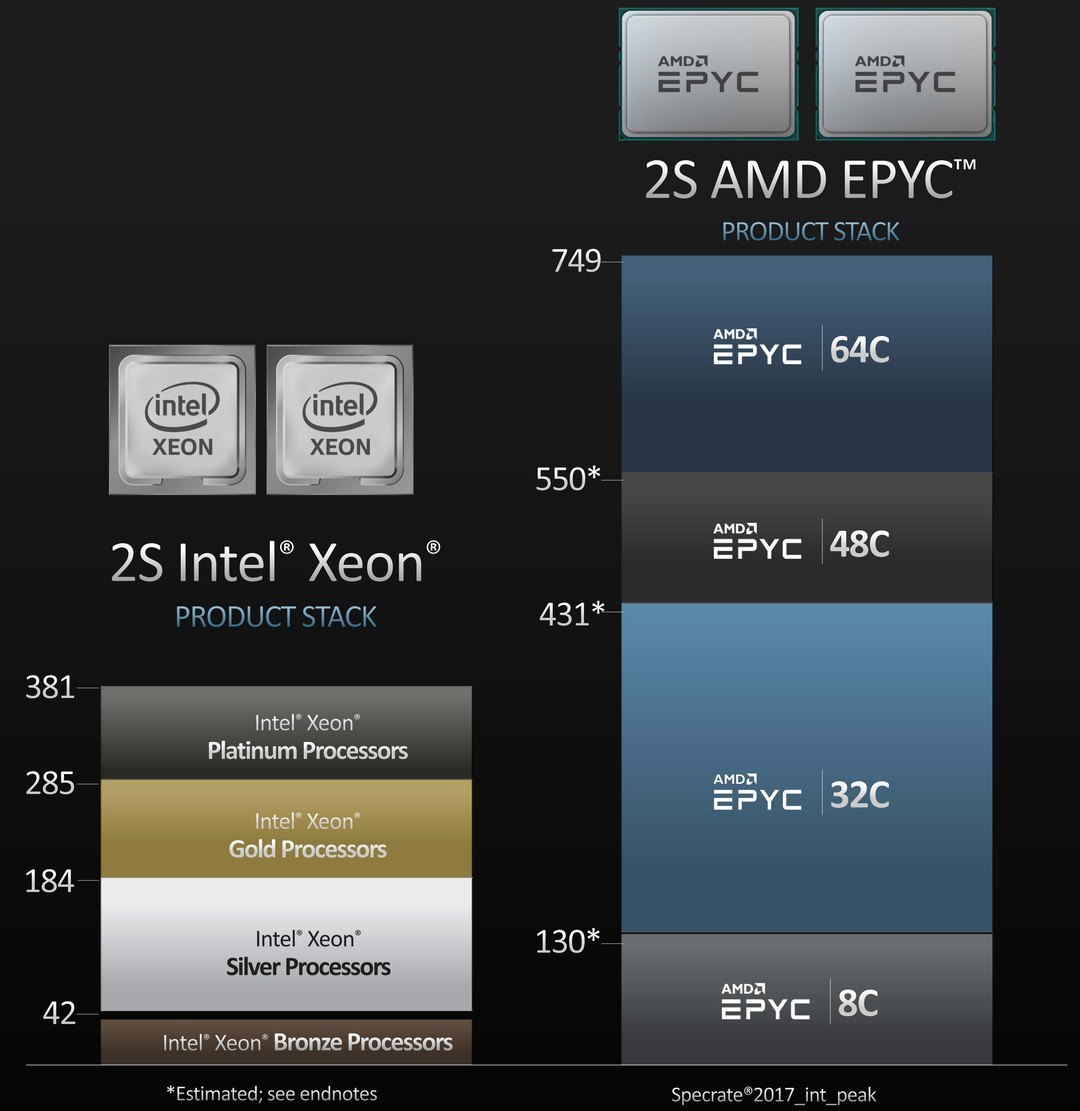

“ROME” LEADERSHIP PERFORMANCE

WITH REAL WORLD RESULTS

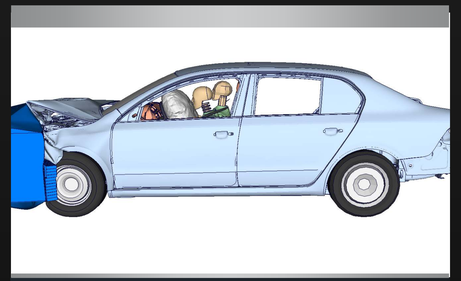

ENGINEERING SIMULATIONS

esi get it right®

UP TO 58%
HIGHER PERFORMANCE

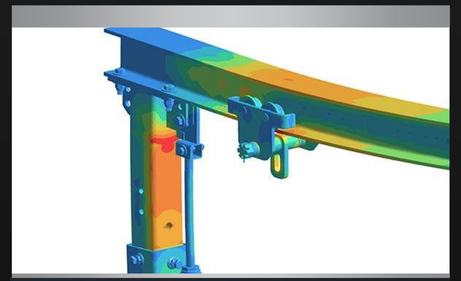

STRUCTURAL ANALYSIS

Altair

UP TO 72%
HIGHER PERFORMANCE

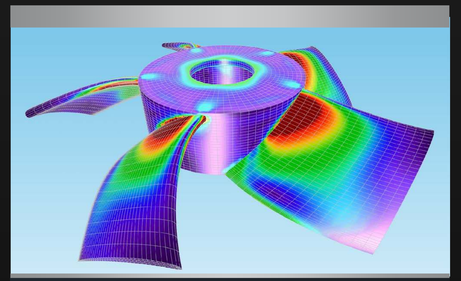

FINITE ELEMENT ANALYSIS

LSTC Livermore Software Technology Corp.

UP TO 79%
HIGHER PERFORMANCE

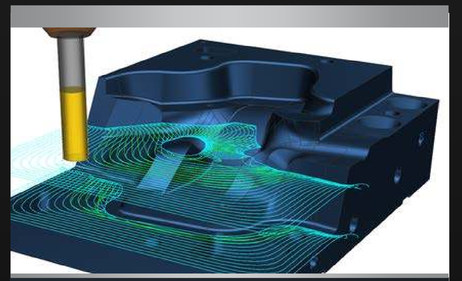

MANUFACTURING

SIEMENS

UP TO 95%
HIGHER PERFORMANCE

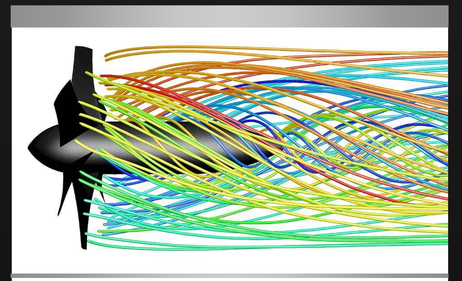

FLUID DYNAMICS

ANSYS®

UP TO 95%
HIGHER PERFORMANCE

SEE ENDNOTES ROM-63, ROM-56, ROM-49, ROM-42, ROM-70

AMD

CONCLUSION

“ZEN 2” Delivers Performance Uplift

- 15% IPC
- Up to 2x instructions per unit energy

Chiplet Deployment And Partitioning

- Enable efficient targeting of technology – performance/power/cost
- Faster deployment of product stack - Client CPU, Server CPU, Chipset

Resulting In Industry Leading Products

- 3rd generation Ryzen (“Matisse”)
- 2nd generation EPYC (“Rome”)

ENDNOTES

Slide 3: Claim: 15% IPC uplift.

- AMD "Zen 2" CPU-based system scored an estimated 15% higher than previous generation AMD “Zen” based system using estimated SPECint®_base2006 results. SPEC and SPECint are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org. GD-141.

Slide 4: Claim: Up to 2x instructions per unit energy.

- Testing conducted by AMD Performance Labs as of 7/12/2019 with 2nd Generation Ryzen and 3rd Generation Ryzen engineering samples using estimated SPECint®_base2006 results. PC manufacturers may vary configurations yielding different results. SPEC and SPECint are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org.

Slide 5

- * AMD recommended settings for these security features can be found at https://developer.amd.com/wp-content/resources/Architecture_Guidelines_Update_Indirect_Branch_Control.pdf.
- ** AMD recommended settings for these security features can be found at https://developer.amd.com/wp-content/resources/124441_AMD64_SpeculativeStoreBypassDisable_Whitepaper_final.pdf

Slide 15

- AMD’s product warranty does not cover damages caused by overclocking, even when overclocking is enabled via AMD hardware and/or software. GD-26

Slide 16, 17

See table at right for test system configurations.

- Compute: Testing by AMD performance labs using an AMD Ryzen™ 9 3900X, AMD Ryzen™ 7 3800X and AMD Ryzen™ 7 2700X in: DaVinci Resolve, Adobe Premiere, Cinebench R20, Handbrake 1.1.1, LAME MP3 Encoder, and POV-Ray 3.7. Results may vary.
- Gaming: Testing by AMD performance labs using an AMD Ryzen™ 9 3900X, AMD Ryzen™ 7 3800X and AMD Ryzen™ 7 2700X. All games tested at 1920x1080 with maximum in-game quality preset. Results may vary.
- Unless otherwise noted in the legend(s) of the chart, all performance analyses conducted and published throughout Computex, Next Horizon: Gaming, E3, and the Ryzen™ Processor Reviewer’s Guide were performed on the following system configurations in an air-conditioned climate of 70°F/21°C.

Slide 18

- 2x PCIe BW over prior Ryzen generation is based on theoretical maximum bandwidth of PCIe Gen3 vs PCIe Gen4 speeds.

Component	3rd Gen AMD Ryzen™ CPUs	2nd Gen AMD Ryzen™ CPUs
Processors Tested	AMD Ryzen™ 9 3900X AMD Ryzen™ 7 3800X AMD Ryzen™ 7 3700X AMD Ryzen™ 5 3600X AMD Ryzen™ 5 3600	AMD Ryzen™ 7 2700X
Cooler	Noctua NH-D15S PWM Auto	Noctua NH-D15S PWM Auto
Motherboard	AMD Reference Motherboard	AMD Reference Motherboard
BIOS	AGESA ComboPI-1003	AGESA PinnaclePI-1007
RAM	2x8GB DDR4-3600 CL16	2x8GB DDR4-3200 CL14
Storage	Phison PS5016	Phison PS5016
GPU	GeForce RTX 2080	GeForce RTX 2080
GPU Driver	430.39	430.39
Operating System	Windows® 10 v1903	Windows® 10 v1903
Security Mitigations	Windows® 10 1903 Default	Windows® 10 1903 Default
UEFI CPPC2	Activated	Activated
Topology Awareness	Activated	Activated
Multi-Core Enhancement	N/A	N/A
Precision Boost Overdrive	Disabled	N/A

AMD

ENDNOTES (CONT.)

Slide 19, 20

- Faster D2D interconnect and assembly of all memory controllers and PHYs on a single IOD, improves effective memory latency in NPS1 over prior generation EPYC™ .
- Each AMD EPYC processor has 8 memory channels. Each Intel Xeon Scalable processor has 6 memory channels. 8 – 6 = 2 ÷ 6 = 0.33 AMD EPYC has 33% more memory bandwidth. Class based on industry-standard pin-based (LGA) X86 processors. EPYC-06.

Slide 20:

- Slide represents both published and estimated SPECrate®2017_int_peak performance. Estimates as of July 3, 2019 for AMD EPYC 48C, 32C and 8C processors using computer modeling of preproduction parts and SPECrate®2017_int_peak internal testing results. Results may vary with production silicon testing. Published results for EPYC 64C processor as of August 7, 2019: https://spec.org/cpu2017/results/res2019q3/cpu2017-20190722-16242.html. Intel results as of June 2019: Xeon Platinum: http://spec.org/cpu2017/results/res2019q2/cpu2017-20190429-12779.pdf Xeon Gold: http://spec.org/cpu2017/results/res2019q2/cpu2017-20190404-11744.pdf Xeon Silver: http://spec.org/cpu2017/results/res2019q2/cpu2017-20190430-13444.pdf; Xeon Bronze: http://spec.org/cpu2017/results/res2019q3/cpu2017-20190624-15468.pdf. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. ROM-258

Slide: 21:

- Based on AMD internal testing of ANSYS FLUENT 19.1, lm6000_16m benchmark, as of July 17, 2019 of a 2P EPYC 7742 powered reference server versus a 2P Intel Xeon Platinum 8280 powered server. Results may vary. ROM-42

- Based on AMD internal testing of LSTC LS-DYNA R9.3.0, neon benchmark, as of July 17, 2019 of a 2P EPYC 7742 powered reference server versus a 2P Xeon Platinum 8280 powered server. Results may vary. ROM-49

- Based on AMD internal testing of Altair RADIOSS 2018, T10M benchmark, as of July 17, 2019 using a 2P EPYC 7742 powered reference server versus a 2P Xeon Platinum 8280 powered server. Results may vary. ROM-56

- Based on AMD internal testing of ESI VPS 2018.0, NEON4m benchmark, as of July 17, 2019 using a 2P EPYC 7742 powered reference server versus a 2P Xeon Platinum 8280 powered server. Results may vary. ROM-63

- Based on AMD internal testing of Siemens PLM STAR-CCM+ 14.02.009, kcs_with_physics benchmark, as of July 17, 2019 using a 2P EPYC 7742 powered reference server versus a 2P Xeon Platinum 8280 powered server. Results may vary. ROM-70