FROM RESEARCH TO PRODUCT:
RAS FEATURES IN EPYC AND RADEON INSTINCT

VILAS SRIDHARAN

CLOUD

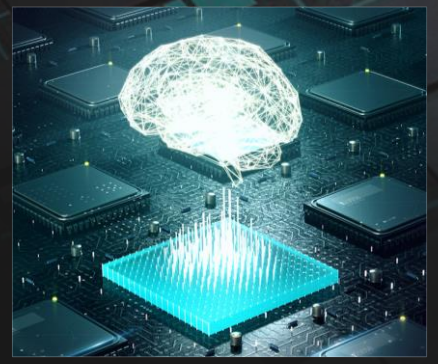

MACHINE INTELLIGENCE

MEDICINE

PERSONAL COMPUTING

GAMING

HIGH PERFORMANCE COMPUTING

SECURITY

AUTOMOTIVE

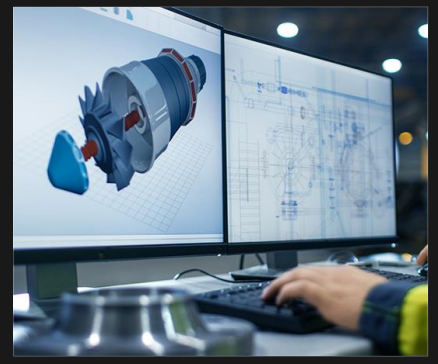

INDUSTRIAL

FINANCIAL SERVICES

VR & AR

DEMAND FOR BETTER EXPERIENCES

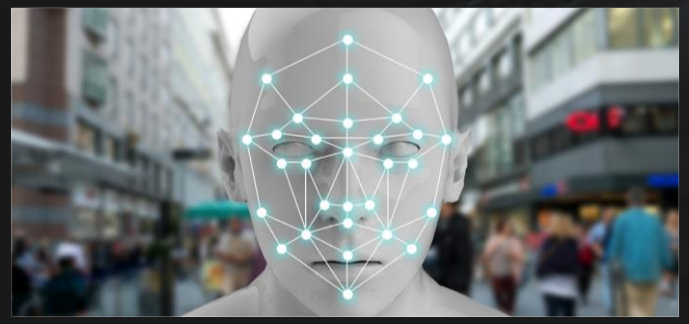

VOICE, GESTURE, FACE RECOGNITION

SUPER HIGH RESOLUTION DISPLAYS

VR, AR

HUGE DEMAND FOR MORE COMPUTE

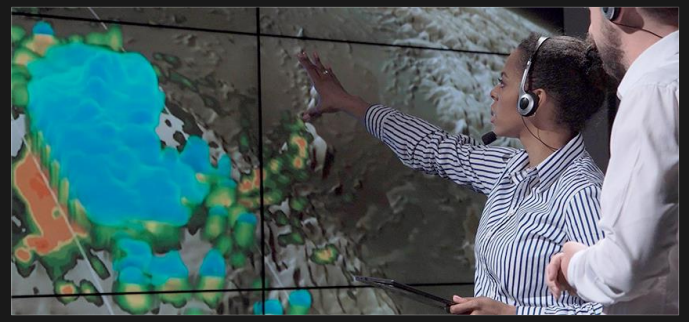

BIG DATA ANALYTICS

HIGH-PERFORMANCE COMPUTING

MACHINE LEARNING

AMD EPYC™ LEADERSHIP

Cloud Service Providers

IaaS/PaaS

Media

Social

SaaS

Enterprise IT

Virtualization

SDS/HCI

Hadoop

NoSQL

High Performance Compute

Design & Simulation

Research & Academia

Machine Learning

Supercomputing

DESIGNED FOR THE CLOUD

AMD RADEON INSTINCT™ MI50

World's First 7nm GPU

Machine Learning Operations for Training and Inference

Flexible Architecture for Different Workloads

End-to-End ECC Protection

DATA CENTER TRENDS

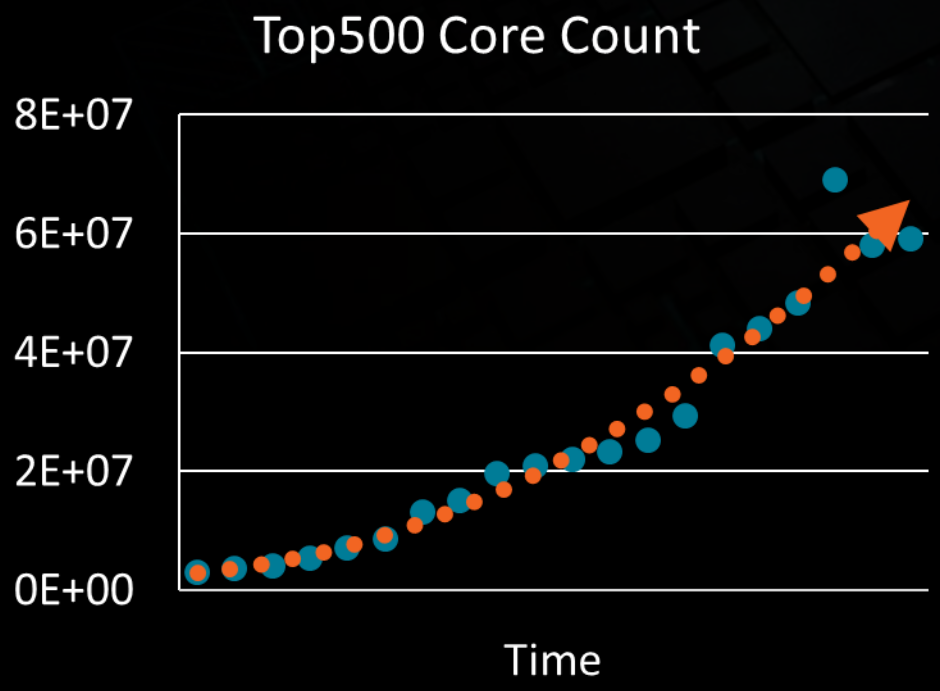

- High reliability to help enable data center growth
- Advanced availability to help improve customer experience
- Robust serviceability to help reduce data center costs

- Justify RAS features with data

FROM RESEARCH TO PRODUCT:
RAS FEATURES IN EPYC AND RADEON INSTINCT

ENDNOTES: 1


MEMORY TRENDS

DRAM BEHAVIORS

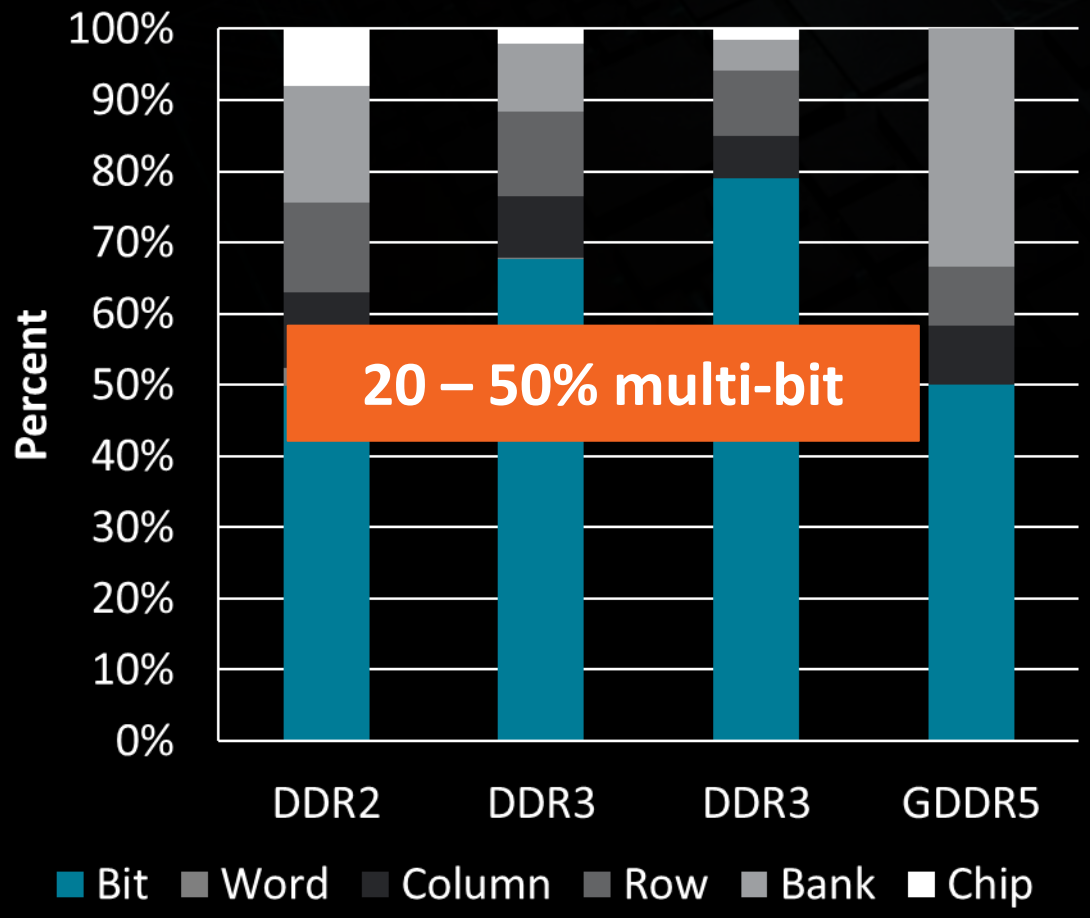

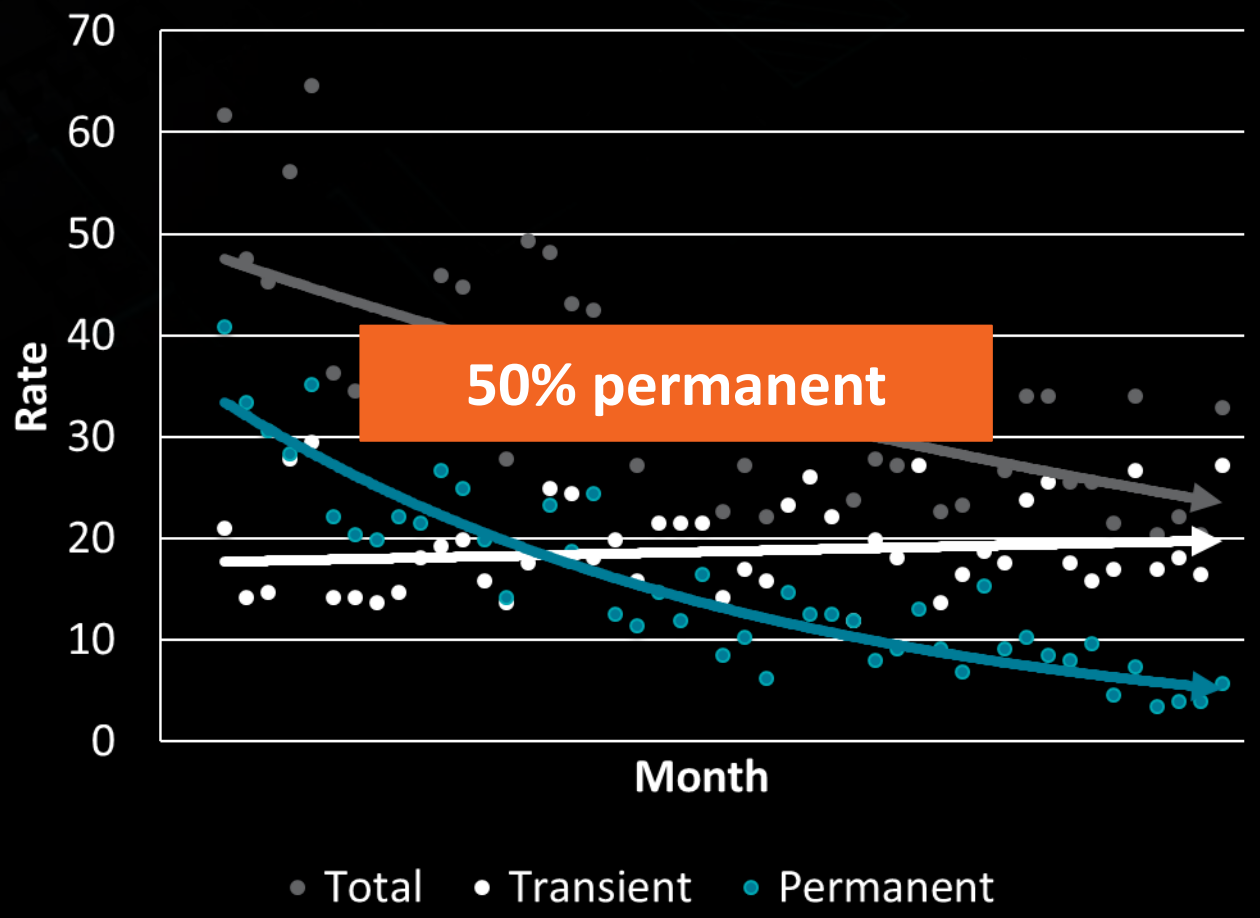

ENDNOTES: 2, 3, 4, 5

BUS SPEED

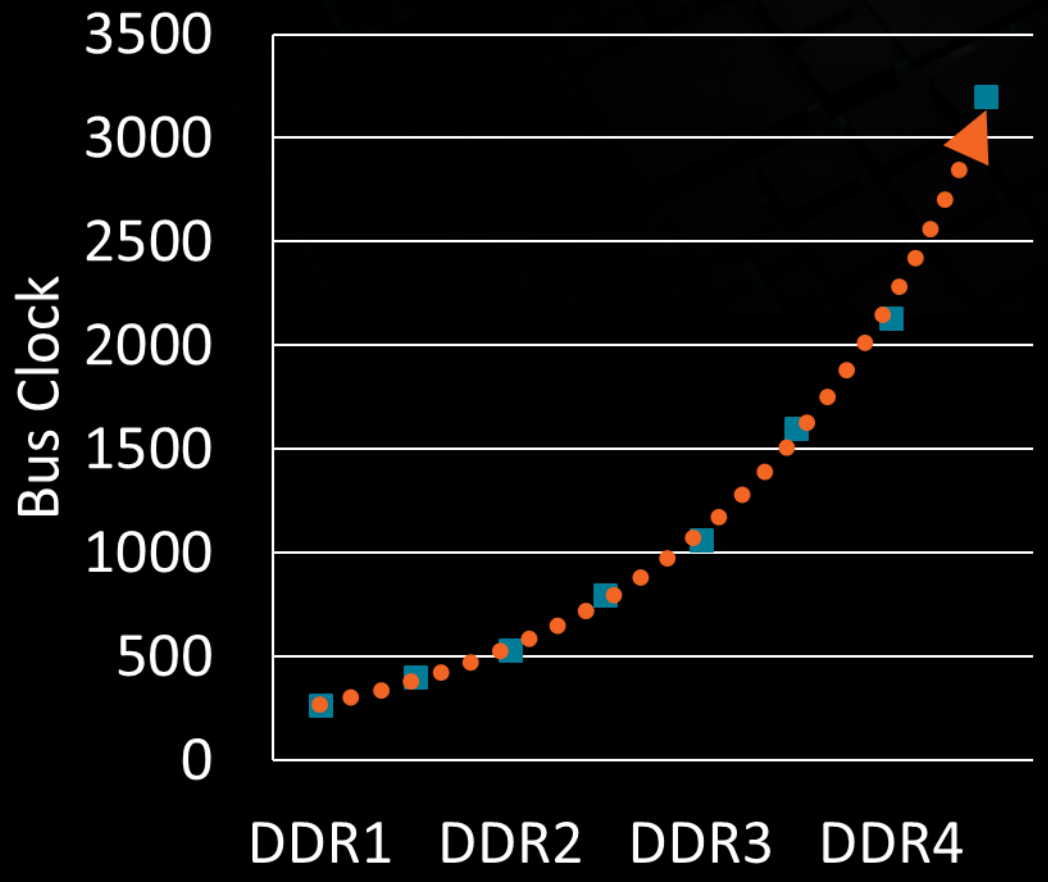

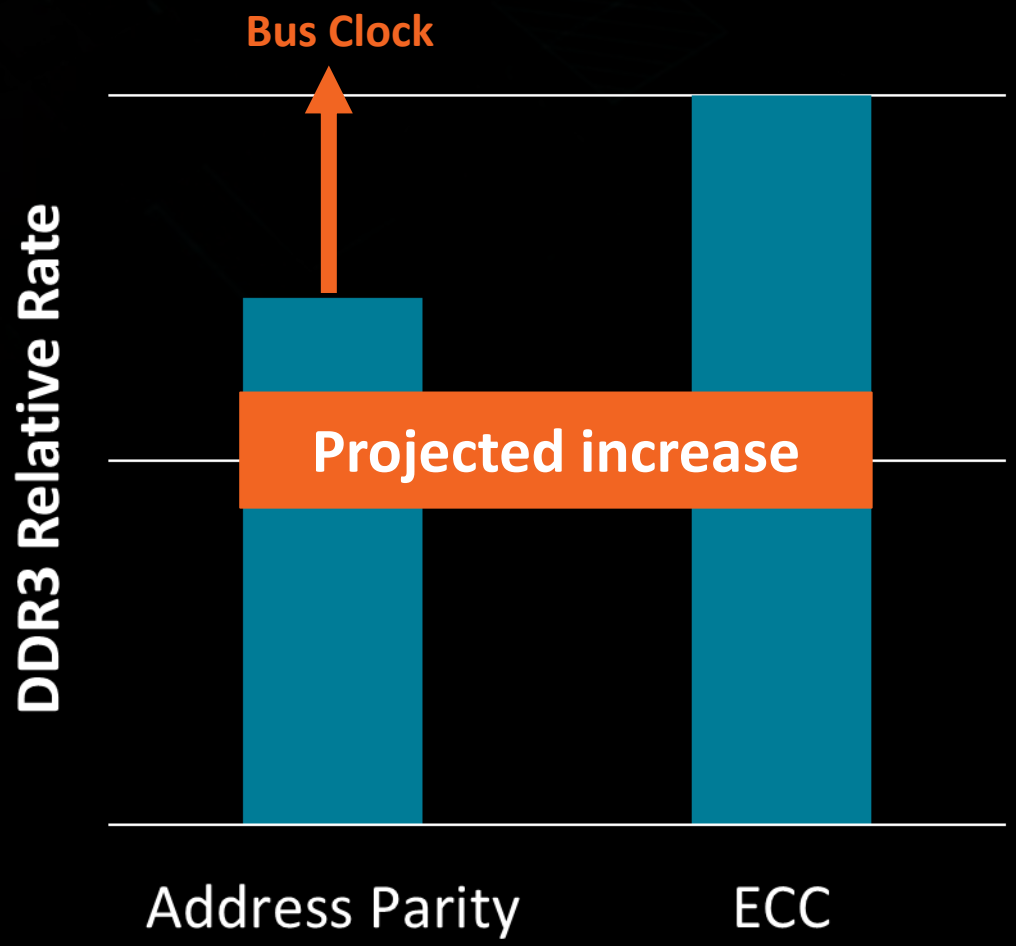

ENDNOTES: 4, 6

EFFECTIVE REMEDIATION

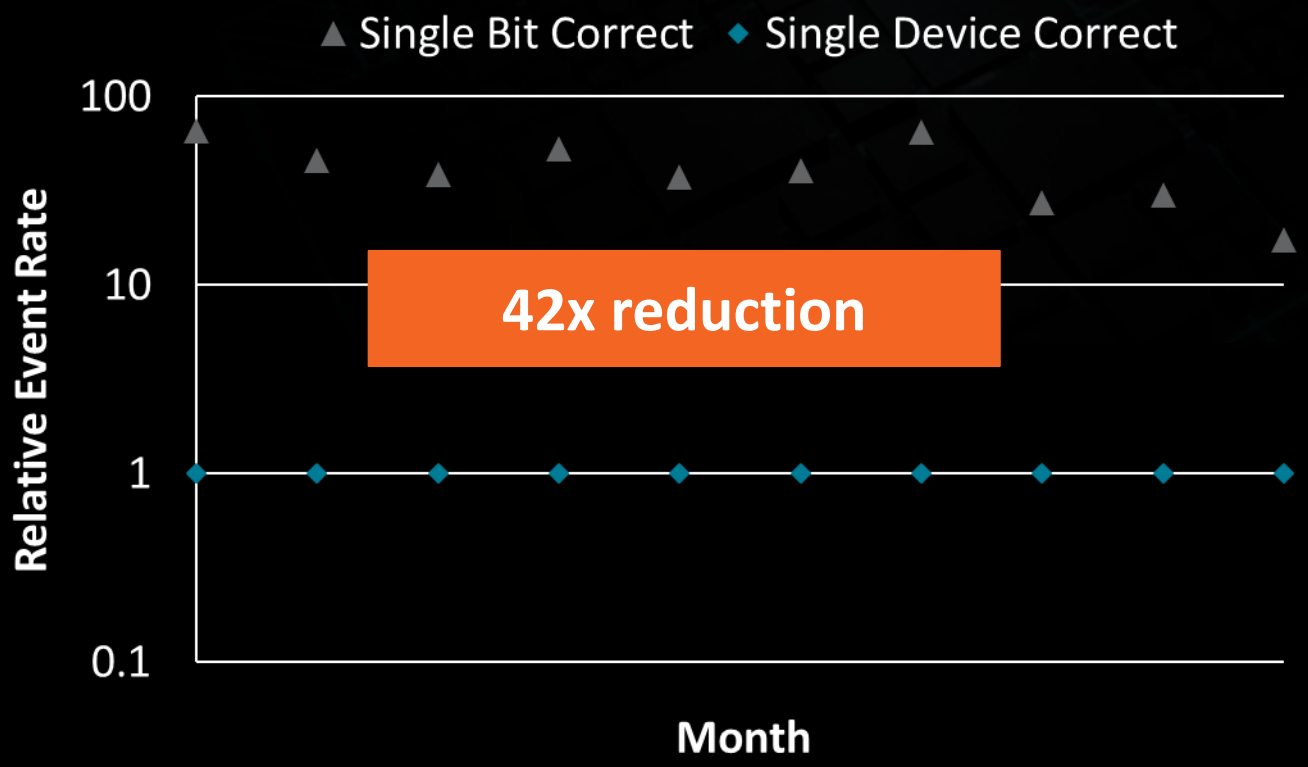

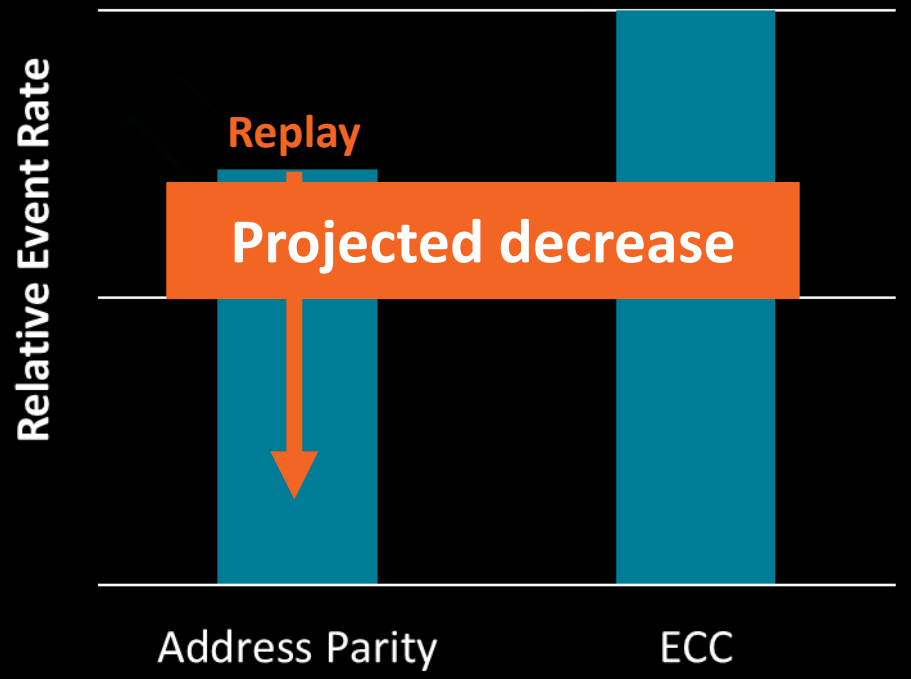

ENDNOTES: 2, 4

PRODUCT FEATURES

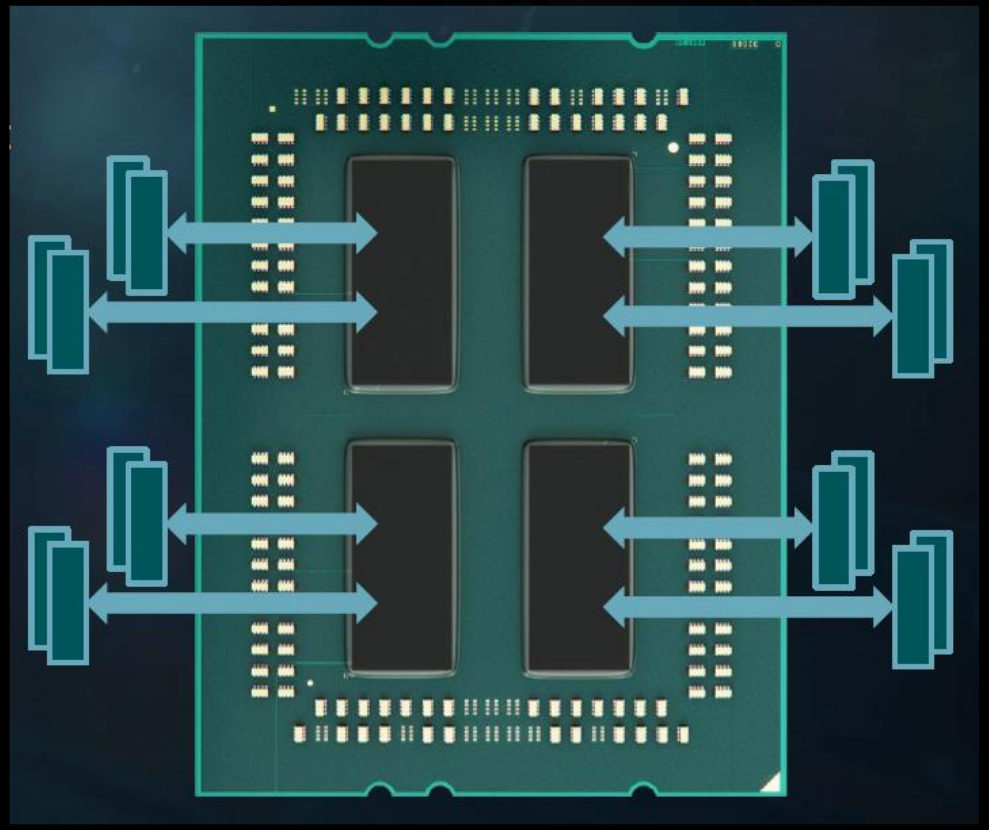

DDR4 SUBSYSTEM

- DRAM ECC with x4 DRAM device correction
- DRAM address/command parity, write CRC—with replay
- Patrol and demand scrubbing
- Data poisoning and Machine Check recovery

AMD EPYC

SOURCE: AMD

SERVICE COSTS

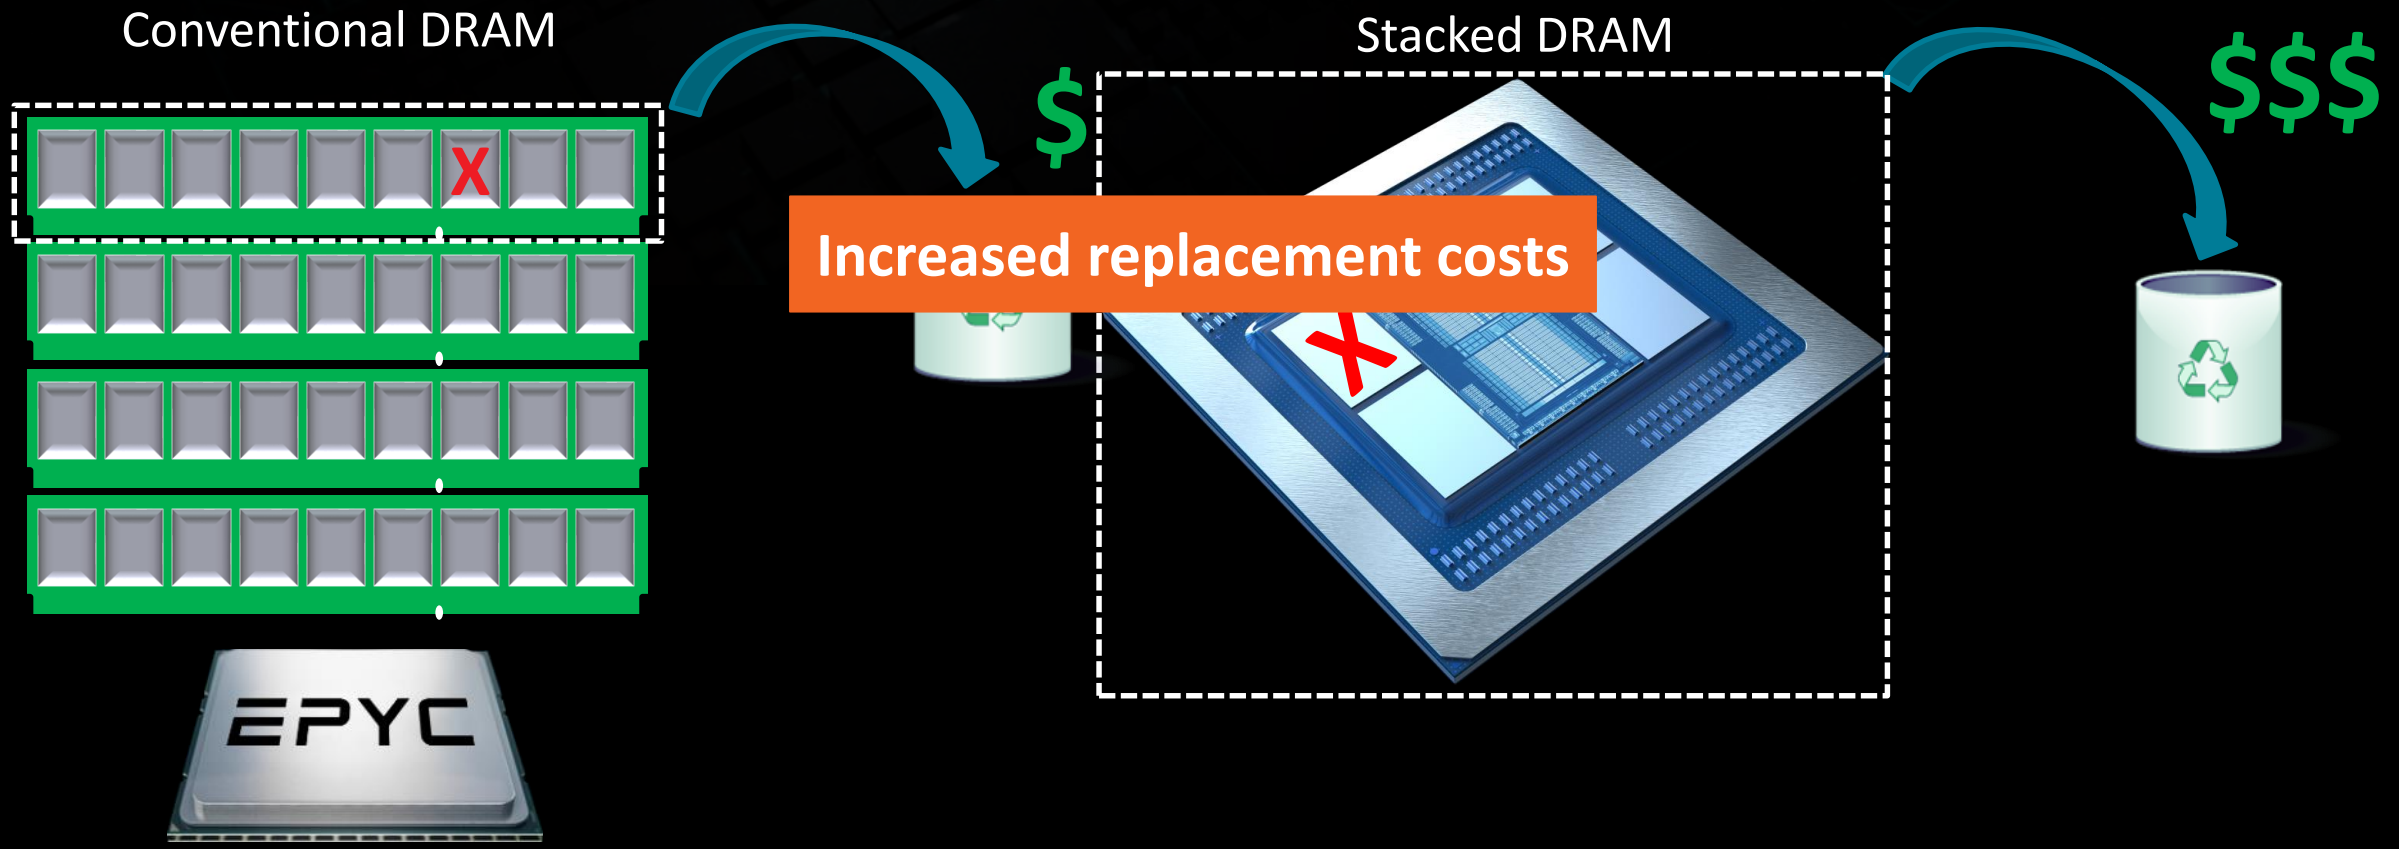

ENDNOTES: 7

MEMORY BANDWIDTH

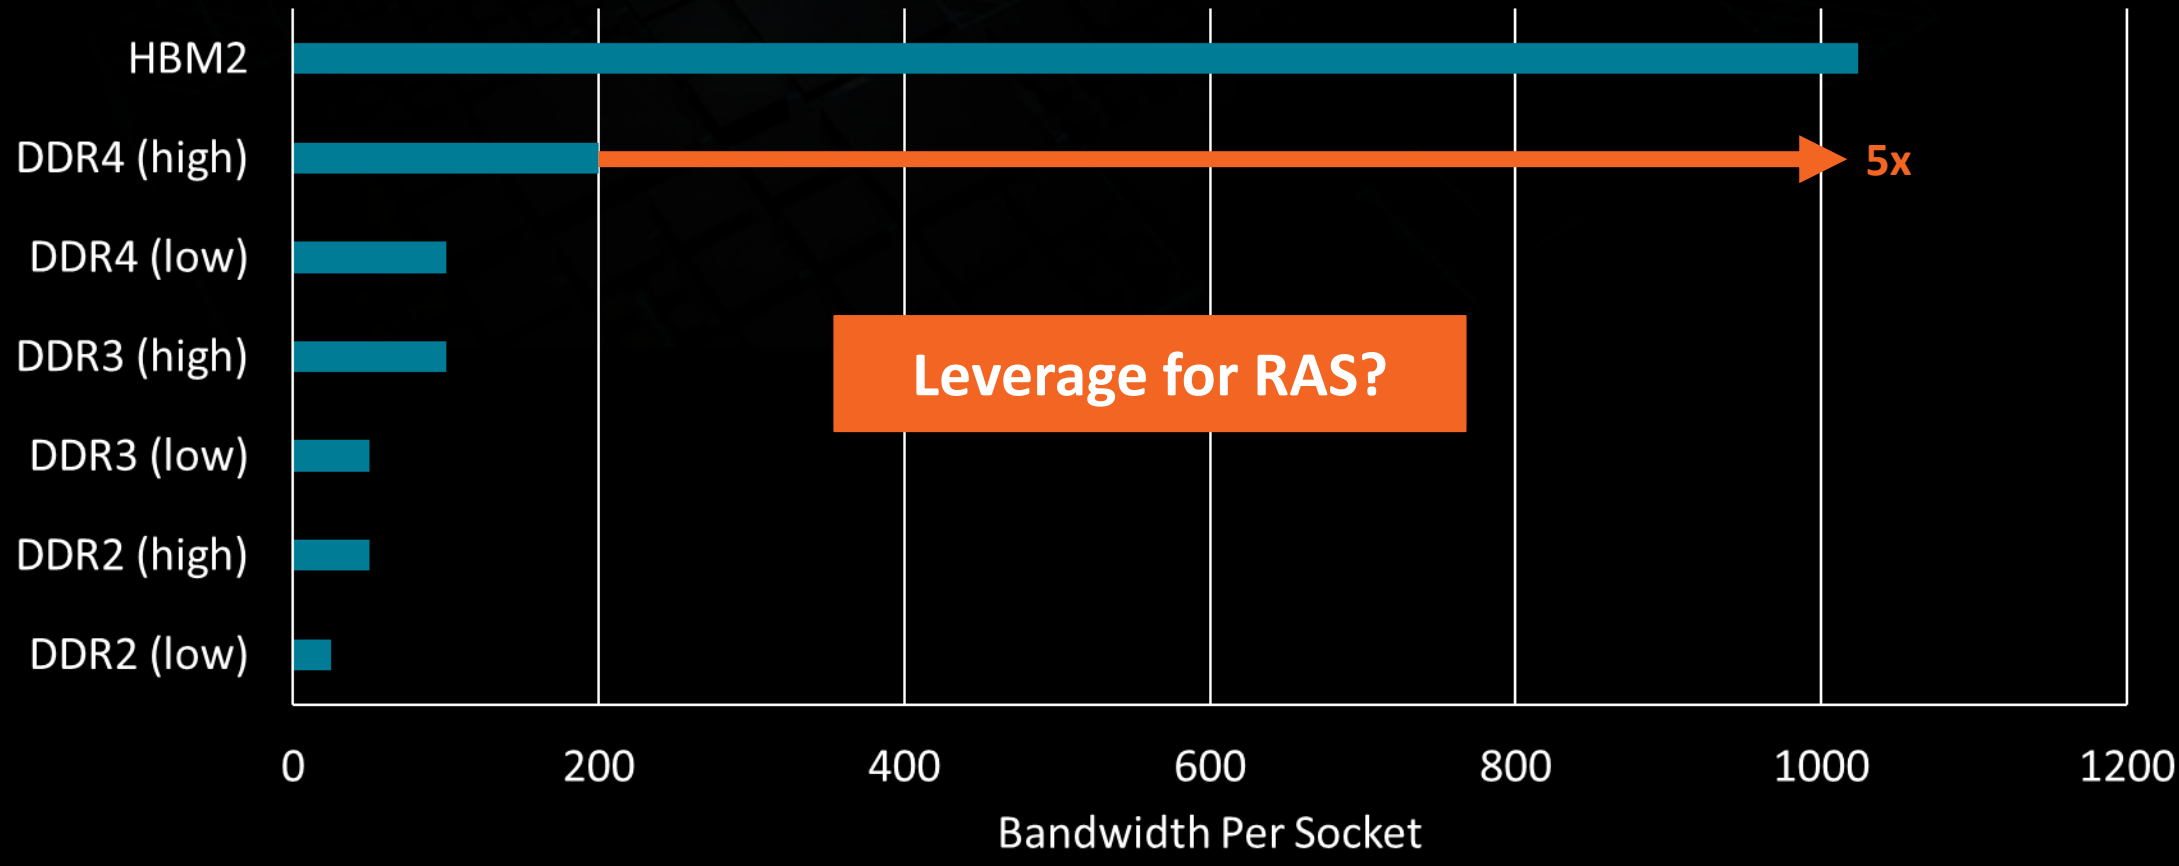

ENDNOTES: 6

REDUNDANT MEMORY

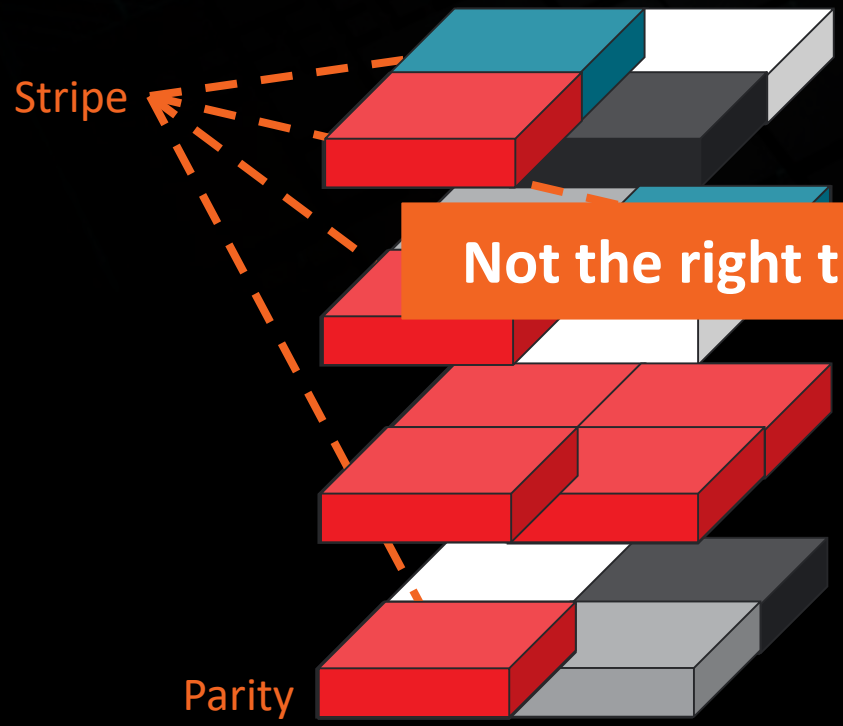

- Can provide
  - Improved reliability

Not the right tradeoff for many markets

  - Lower capacity
  - Reduced bandwidth
  - Unpredictable performance

ENDNOTES: 8

PRODUCT FEATURES

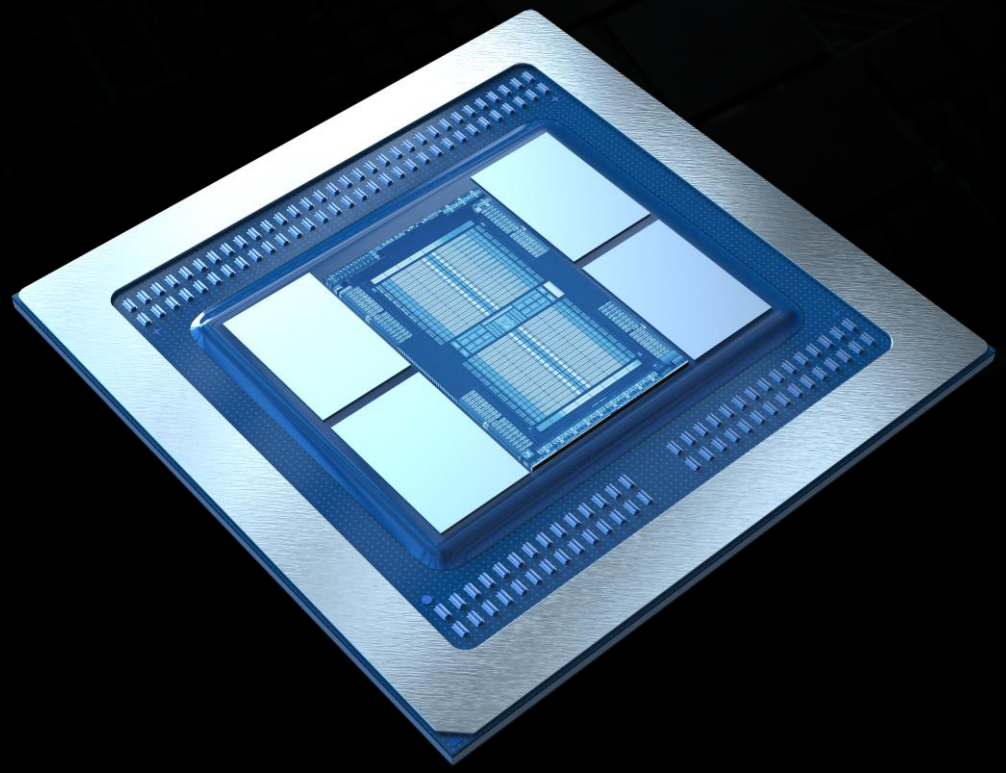

HBM2 SUBSYSTEM

- Single bit correction ECC
- Multi-bit detection CRC
- Stores data XOR address

AMD
RADEON
INSTINCT

SOURCE: AMD


PROCESSOR TRENDS

TRANSIENT UPSETS

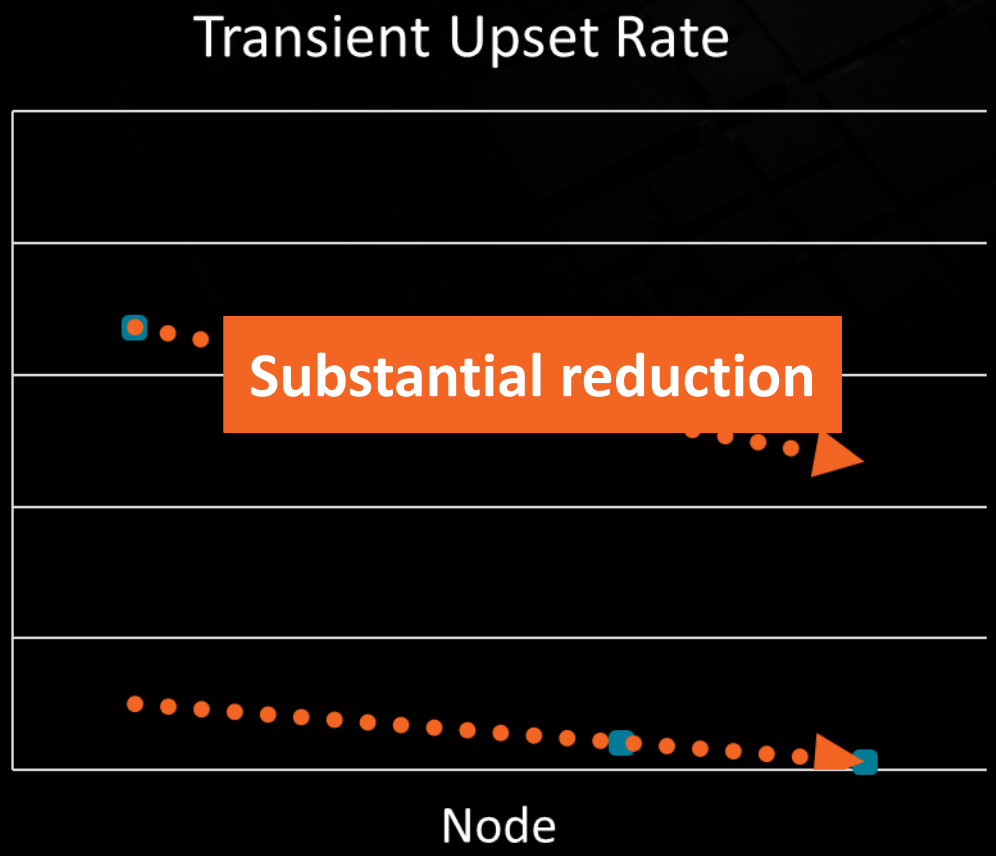

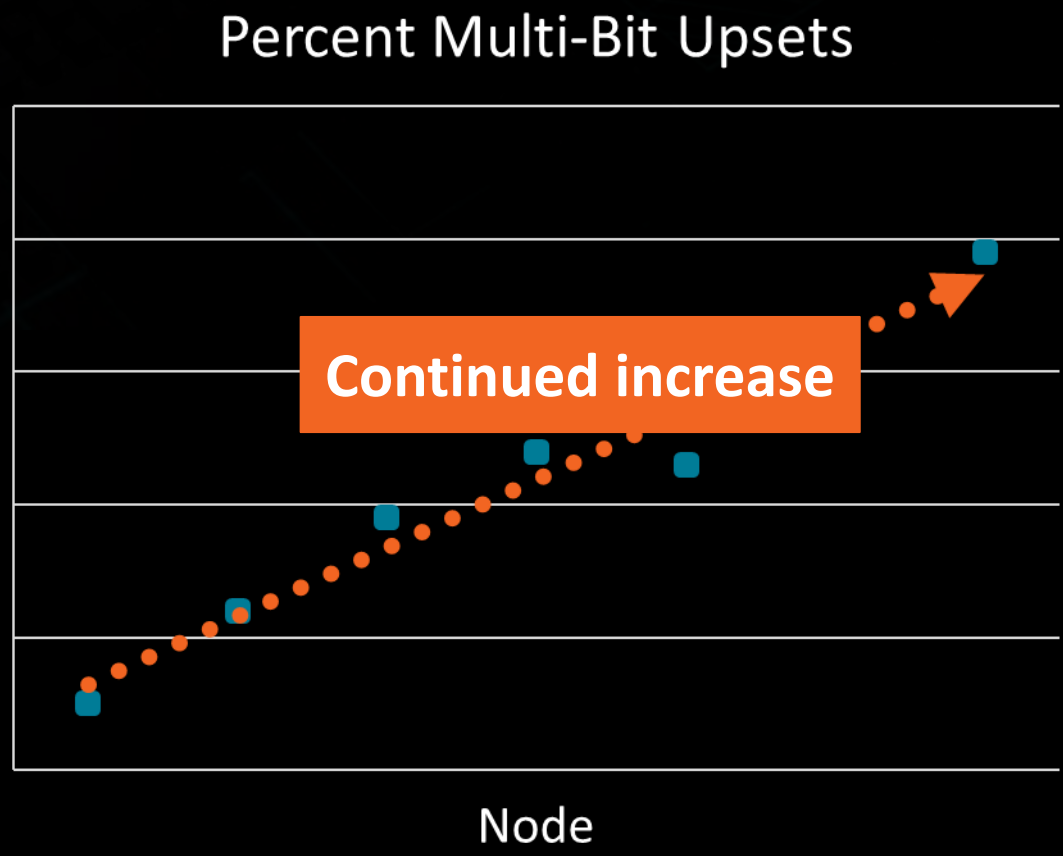

SOURCE: AMD ENDNOTES: 9, 10

REDUCED VOLTAGE

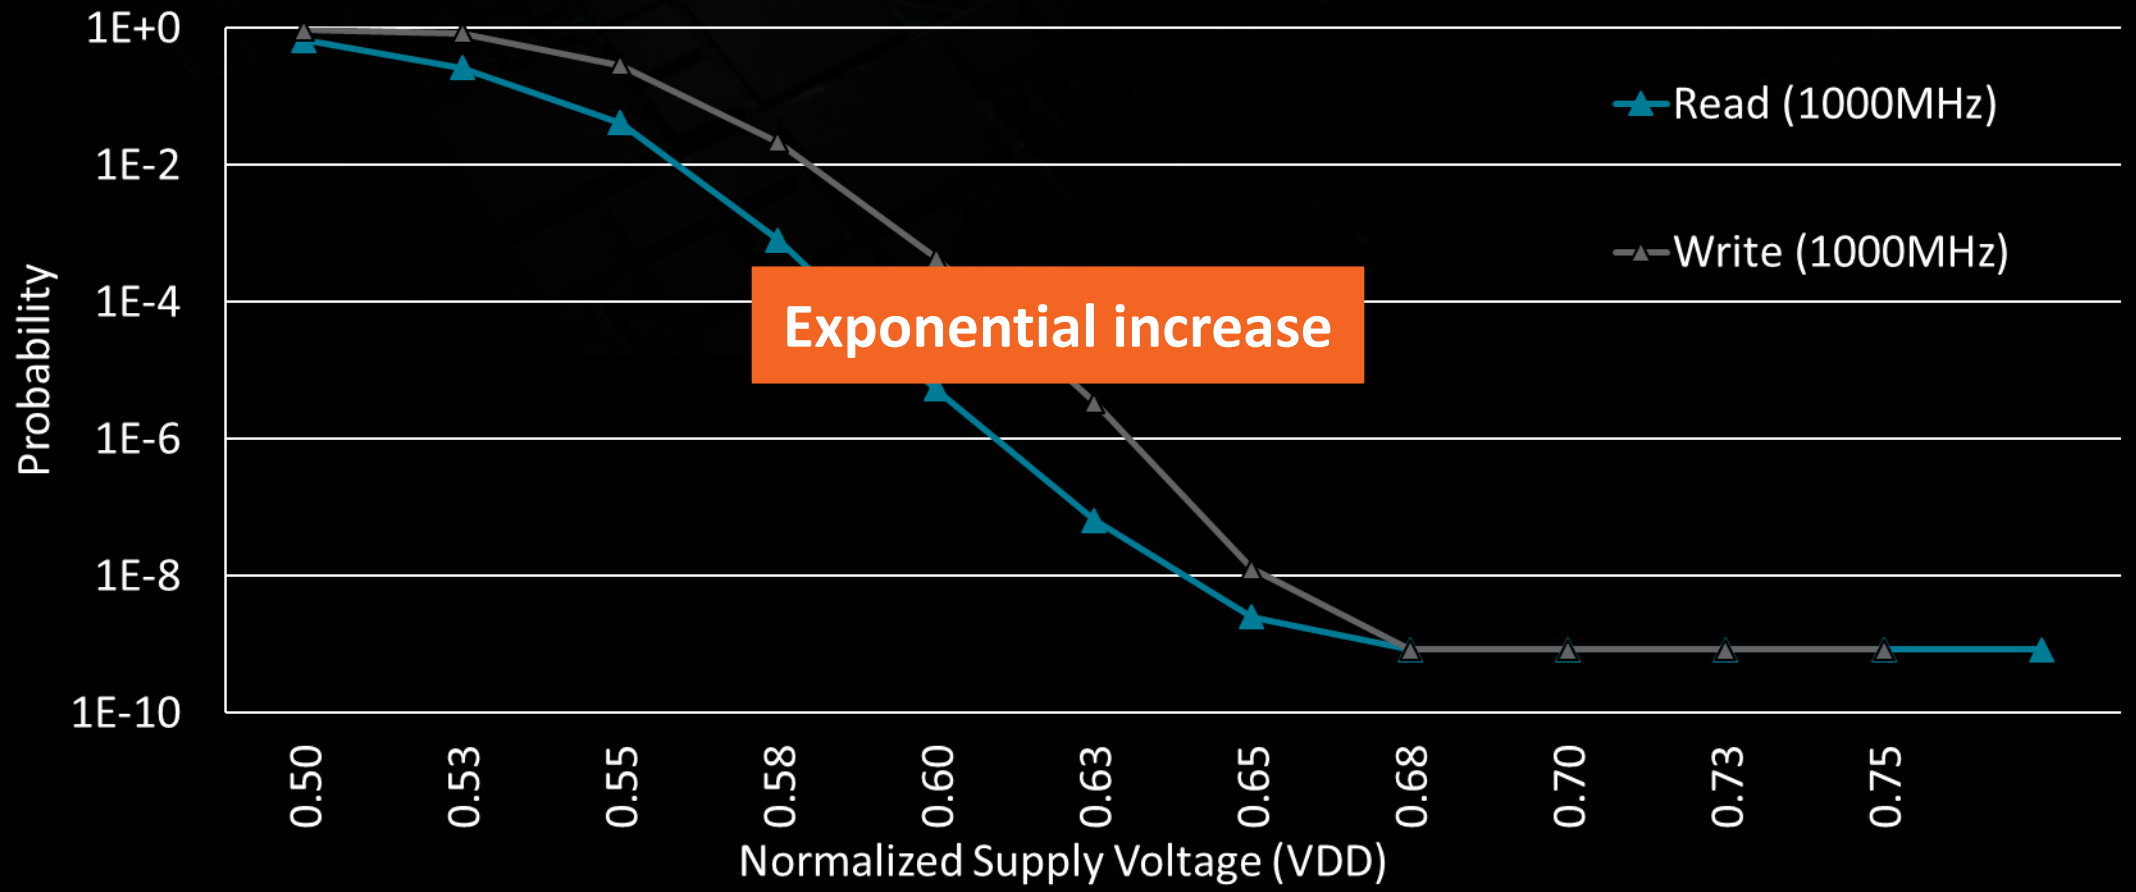

ENDNOTES: 11

AVF ANALYSIS

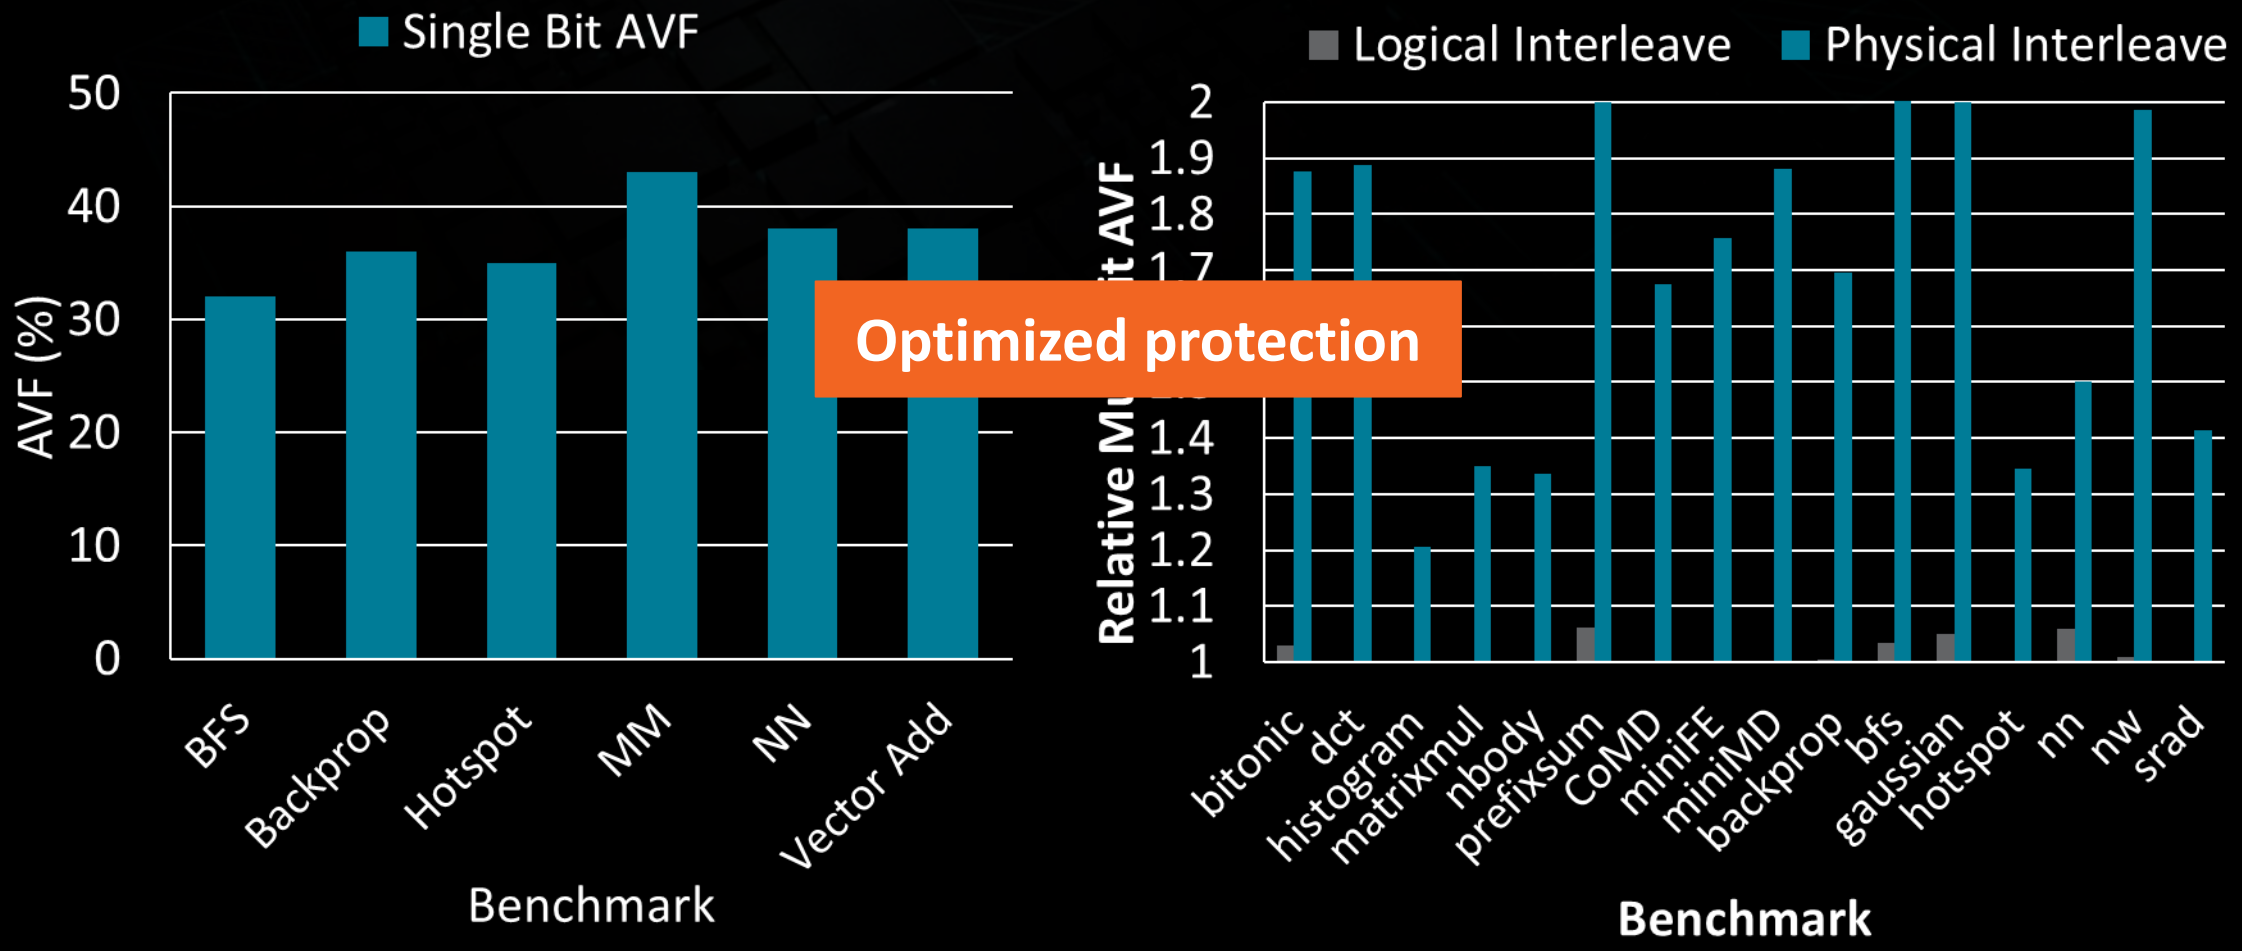

Optimized protection

ENDNOTES: 9, 14

PRODUCT FEATURES

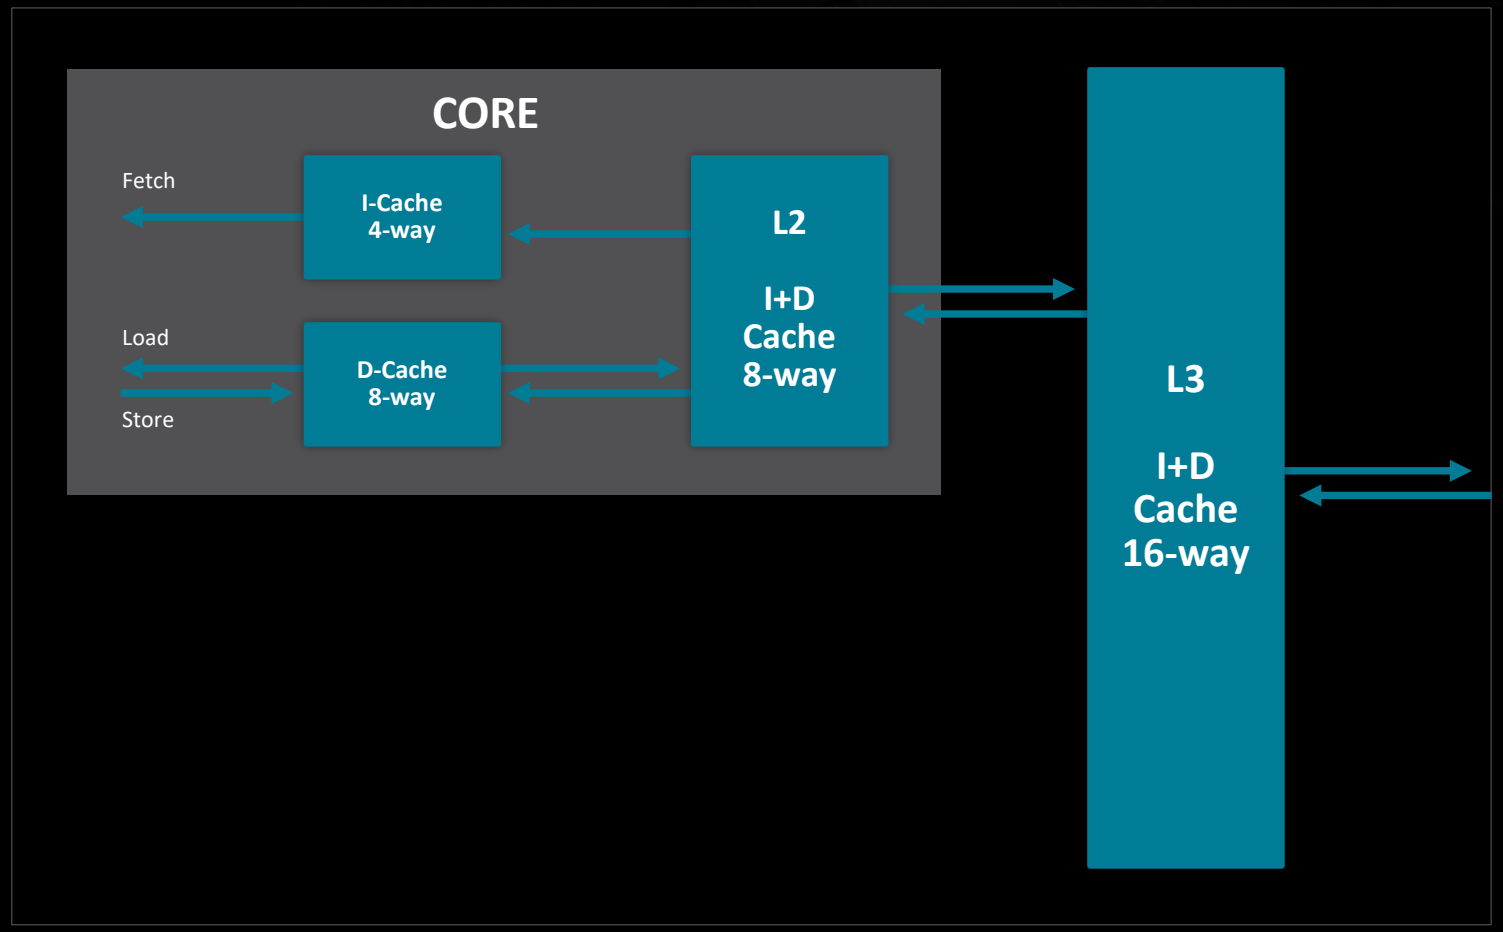

CACHE HIERARCHY

- Fast private L2 cache
- Fast shared L3 cache
- Double bit correct, triple bit detect ECC on L2, L3, and queues
- Interleaving in L2 and L3
- Separate L2/L3 voltage rail (Vddm)

AMD EPYC

ENDNOTES: 13


GPU TRENDS

COMPUTE THROUGHPUT

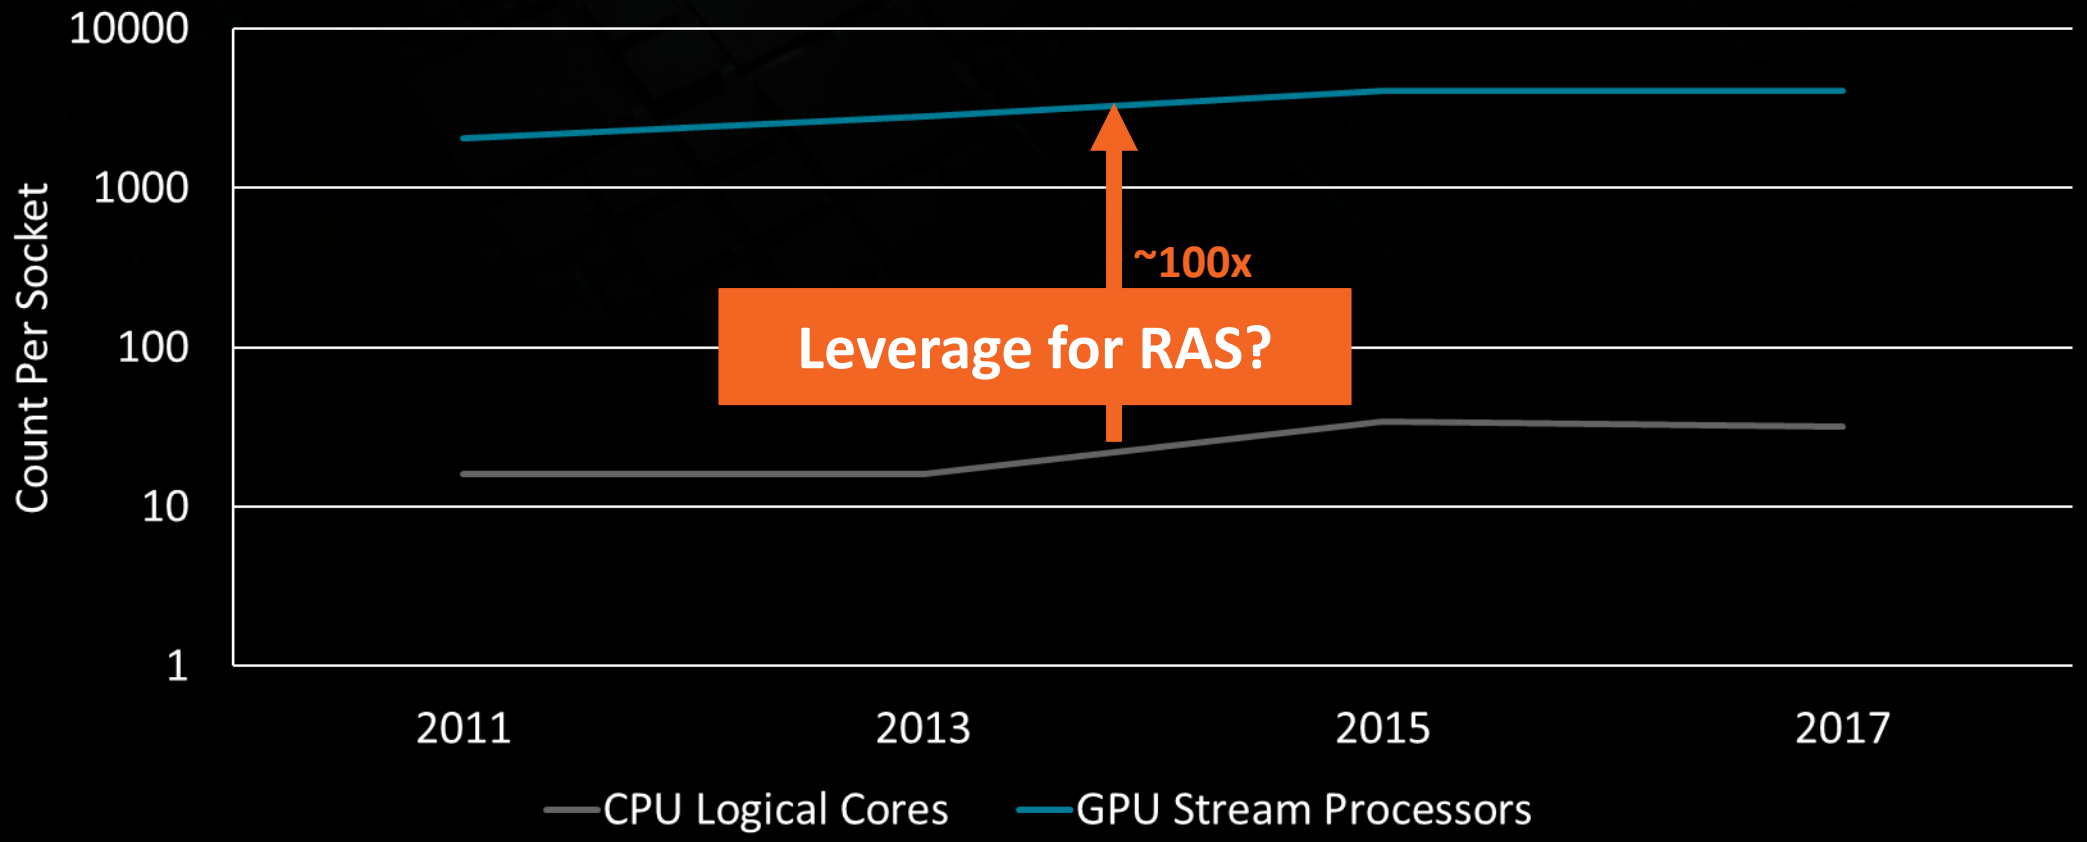

SOURCE: AMD

REDUNDANT MULTITHREADING

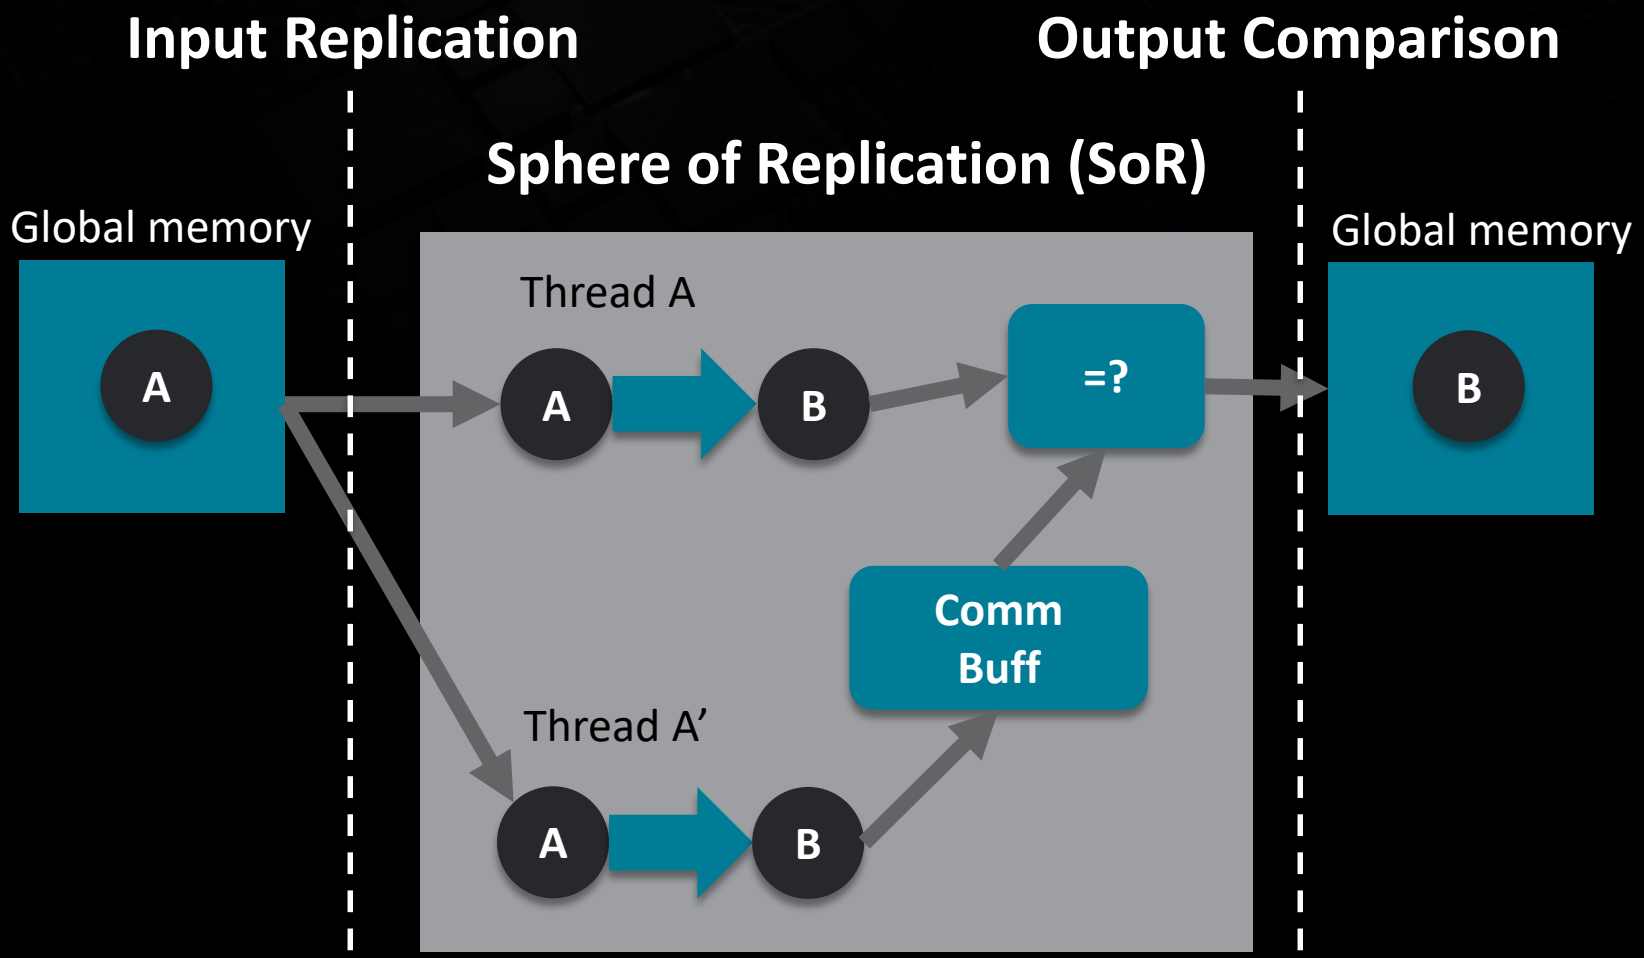

ENDNOTES: 15

REDUNDANT MULTITHREADING

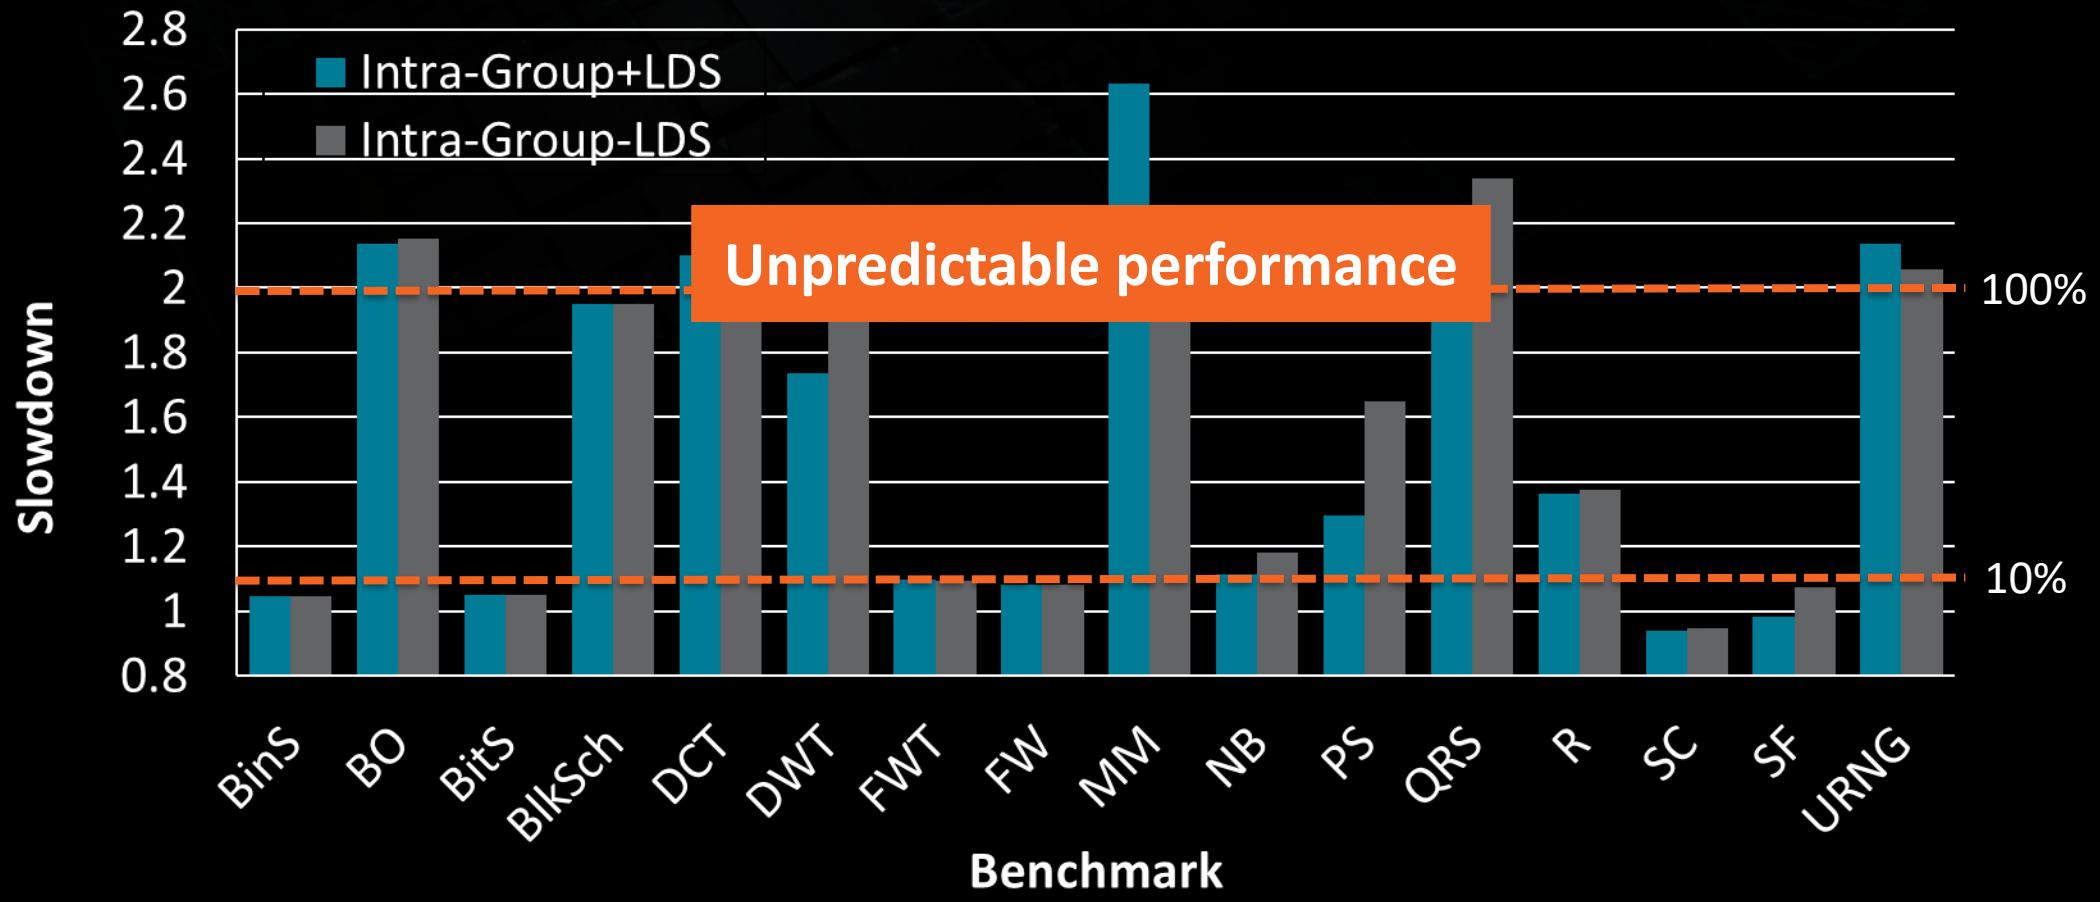

ENDNOTES: 15

AMD

REDUNDANT MULTITHREADING

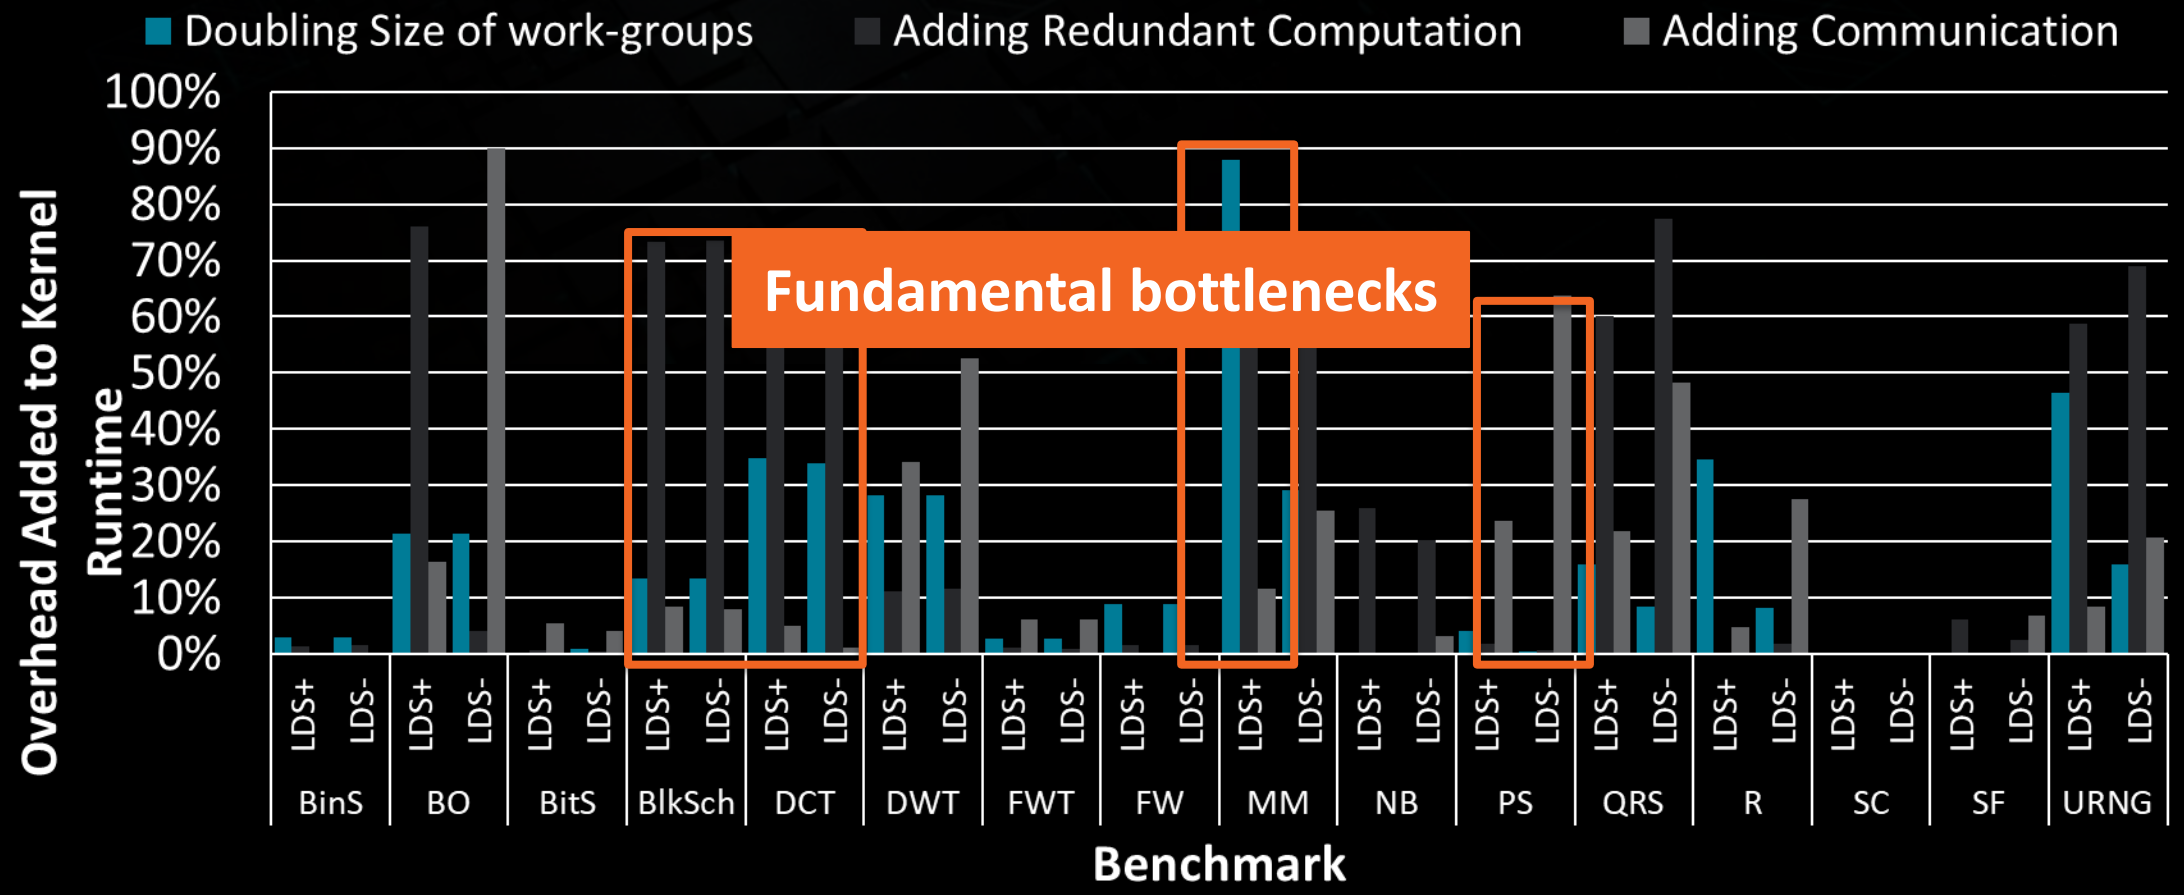

ENDNOTES: 15

ECC ANALYSIS

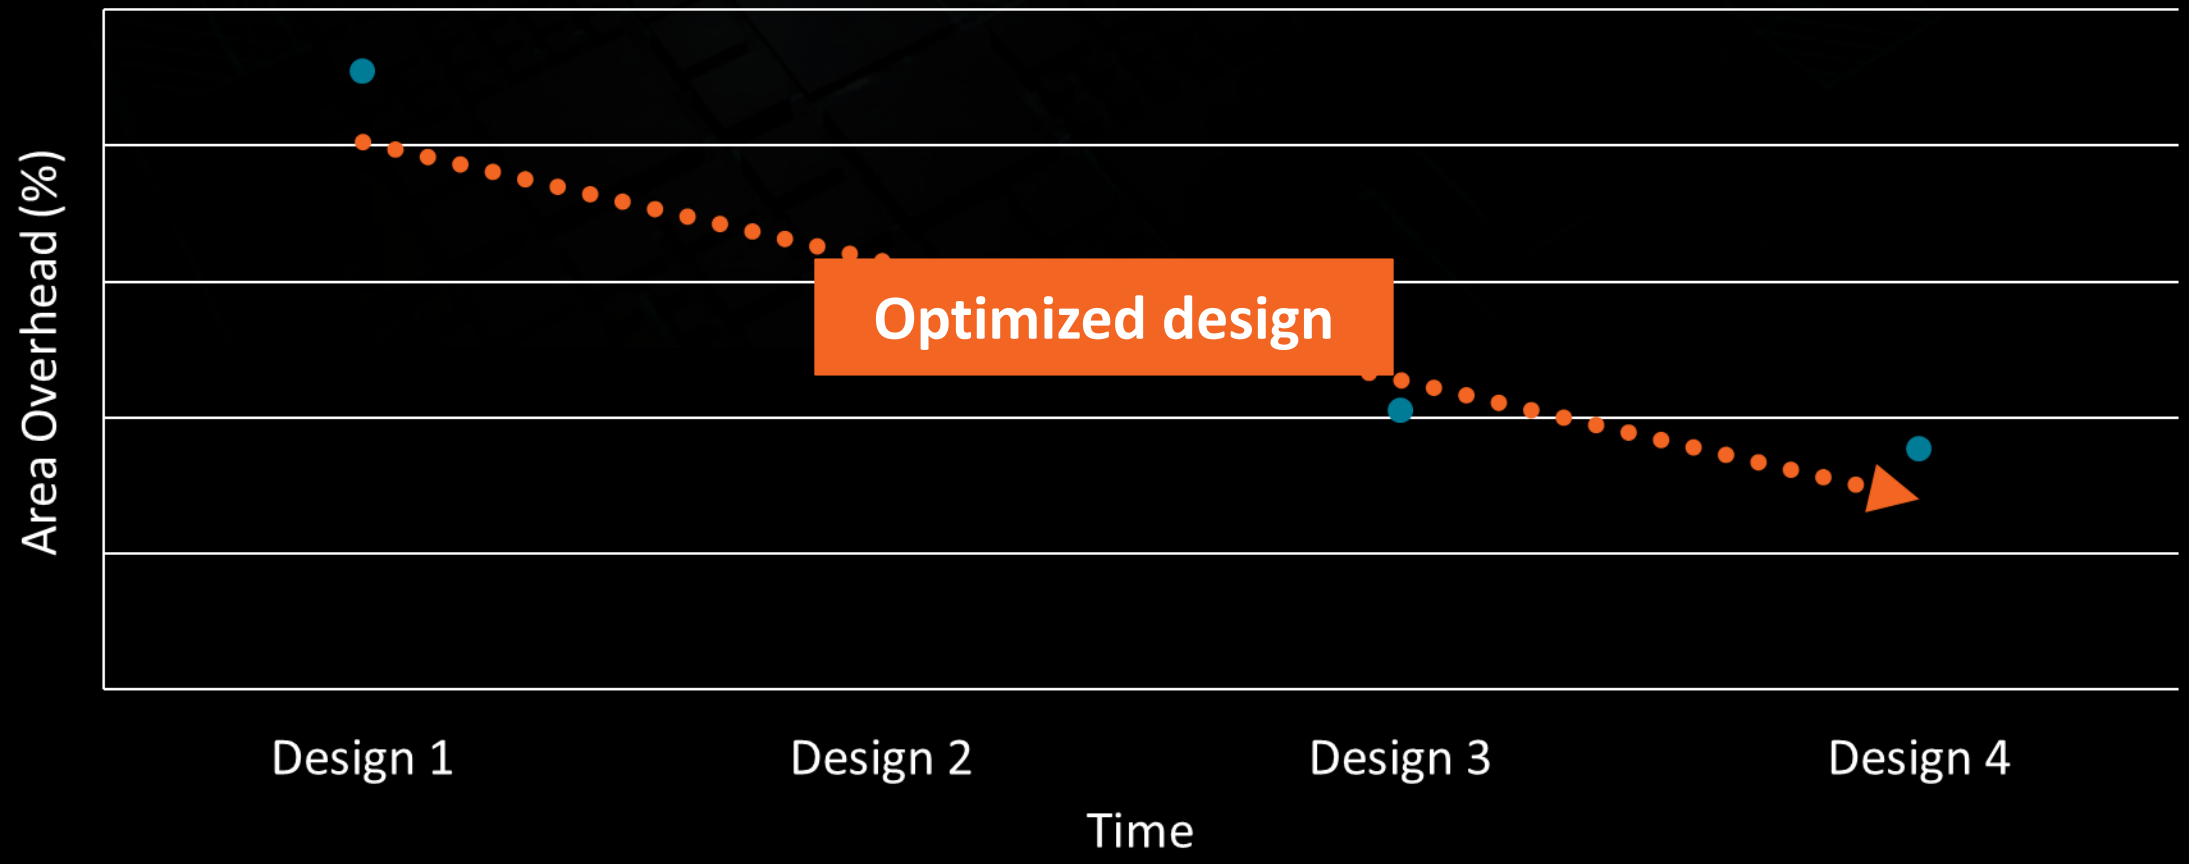

SOURCE: AMD

PRODUCT FEATURES

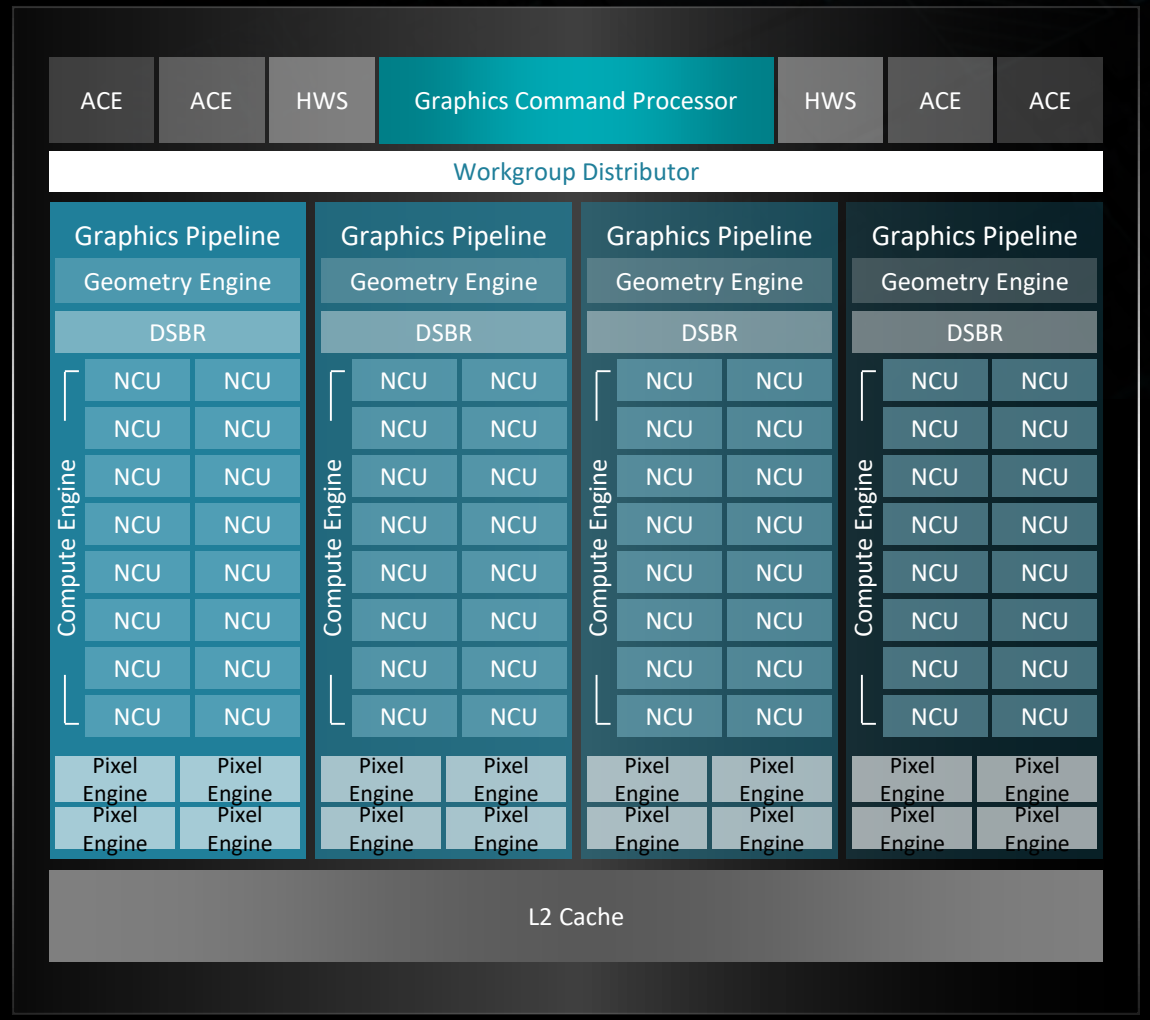

GRAPHICS ENGINE

- ECC on all important arrays
- Modest die area overhead
- Low performance overhead
- Better correction than RMT

AMD RADEON INSTINCT

SOURCE: AMD

ENTERPRISE-CLASS RAS FEATURES

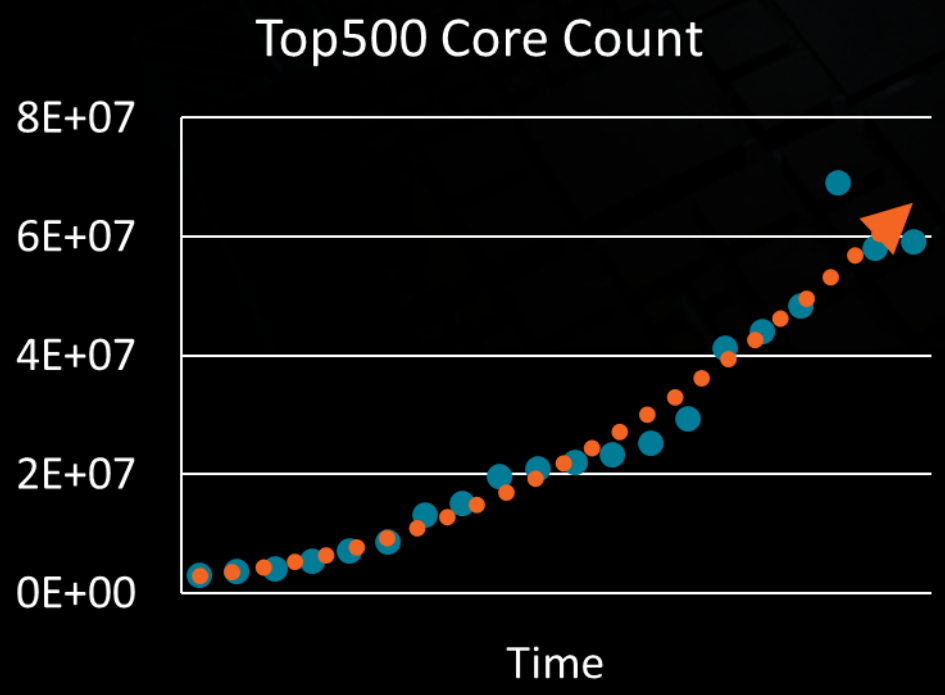

- Understand market requirements
- Adapt to technology trends
- Optimize design to meet customer needs

ENDNOTES: 1

AMD

AMD EPYC

AMD RADEON INSTINCT

RELIABLE COMPUTATION FOR THE MODERN DATACENTER

AMD

ACKNOWLEDGEMENTS

- **Mark Wilkening, Jack Wadden, Si Li, Fritz Previlon, Charu Kalra, Hyeran Jeon, Lukasz G. Szafaryn, Jaewoong Sim, Taniya Siddiqua, Daniel Lowell, Shrikanth Ganapathy, Sudhanva Gurumurthi, Steven E. Raasch, John Kalamatianos, Keith Kasprak, Bradford M. Beckmann, Alexander Lyashevsky, Mike O'Connor, Dean Liberty, Gabriel Loh,** AMD Research, Advanced Micro Devices, Inc.
- **Xun Jian, Rakesh Kumar,** University of Illinois Urbana-Champaign
- **Nathan Debardeleben, Elisabeth Moore, Qiang Guan, Sean Blanchard,** Ultrascale Systems Research Center, Los Alamos National Laboratory
- **Jon Stearley, Kurt B. Ferreira, Scott Levy,** Scalable Architectures, Sandia National Laboratories
- **Devesh Tiwari, Christian Engelmann, Saurabh Gupta,** Oak Ridge National Laboratory
- **John Shalf,** Computational Research Division, Lawrence Berkeley National Laboratory
- **David R Kaeli**, Northeastern University
- **Kevin Skadron**, University of Virginia
- **Larry Kaplan** and many others at Cray
- **David Rohr, Gvozden Neskovic, Prof. Dr. Volker Lindenstruth,** Frankfurt Institute for Advanced Studies (FIAS) / GSI Helmholtzzentrum für Schwerionenforschung
- Many others at the U.S. national labs

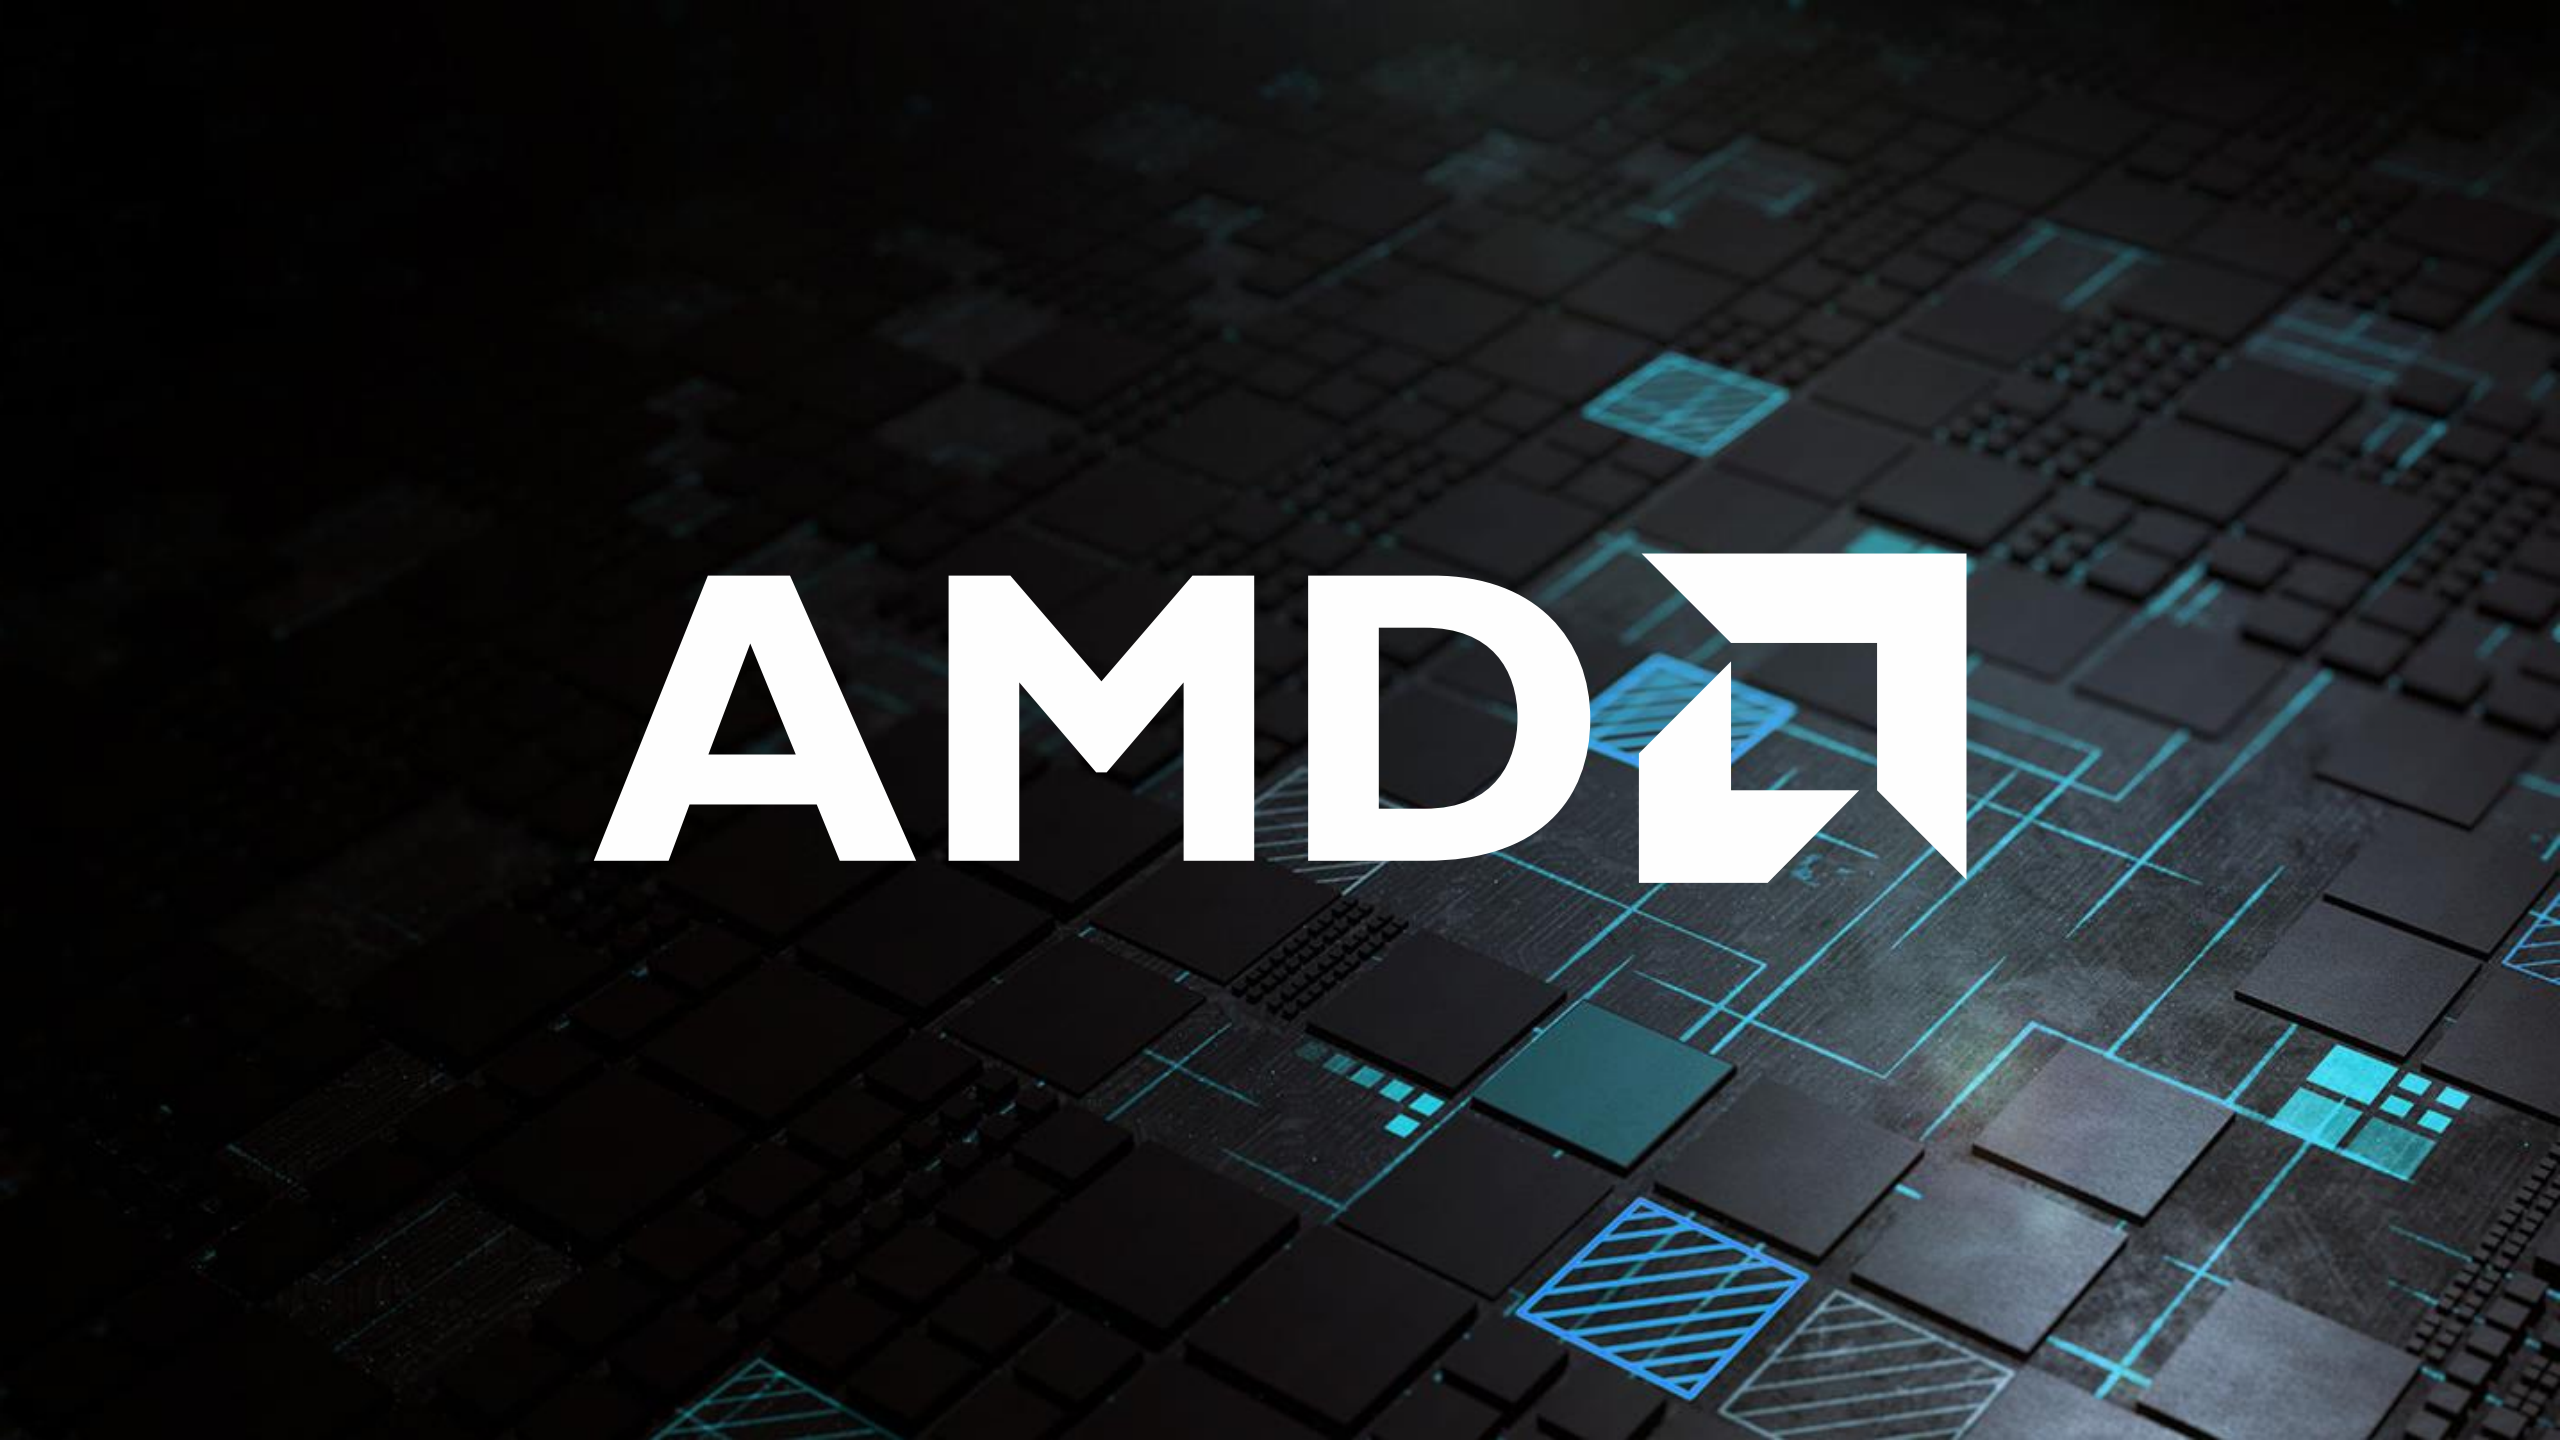
AMD

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

ENDNOTES

[1] Top500.org: https://www.top500.org/statistics/details/osfam/1. Aggregate core count for Top500 systems with Linux family operating systems.

[2] V. Sridharan and D. Liberty, A study of DRAM failures in the field, *SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, Salt Lake City, UT, 2012, pp. 1-11.

[3] V. Sridharan, J. Stearley, N. DeBardeleben, S. Blanchard and S. Gurumurthi, Feng Shui of supercomputer memory positional effects in DRAM and SRAM faults, *SC '13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, Denver, CO, 2013, pp. 1-11.

[4] V. Sridharan, N. DeBardeleben, S. Blanchard, K. B. Ferreira, J. Stearley, J. Shalf, and S. Gurumurthi. 2015. Memory Errors in Modern Systems: The Good, The Bad, and The Ugly. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems* (ASPLOS '15). ACM, New York, NY, USA, 297-310

[5] T. Siddiqua *et al*., Lifetime memory reliability data from the field, *2017 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, Cambridge, 2017, pp. 1-6.

[6] What is the difference between SDRAM, DDR1, DDR2, DDR3 and DDR4? https://www.transcend-info.com/Support/FAQ-296/. Uses DDRx and HBM2 data rates defined by the specification when released. Assumes 8-channel DDRx per socket or 4 stacks of of HBM2 per socket.

[7] J. Sim, G. H. Loh, V. Sridharan, and M. O'Connor. 2013. Resilient die-stacked DRAM caches. In *Proceedings of the 40th Annual International Symposium on Computer Architecture* (ISCA '13). ACM, New York, NY, USA, 416-427.

[8] X. Jian, V. Sridharan and R. Kumar, Parity Helix: Efficient protection for single-dimensional faults in multi-dimensional memory systems, *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Barcelona, 2016, pp. 555-567.

[9] M. Wilkening, V. Sridharan, S. Li, F. Previlon, S. Gurumurthi, and D. R. Kaeli. 2014. Calculating Architectural Vulnerability Factors for Spatial Multi-bit Transient Faults. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture* (MICRO-47). IEEE Computer Society, Washington, DC, USA, 293-305.

[10] E. Ibe, H. Taniguchi, Y. Yahagi, K. Shimbo and T. Toba, Impact of Scaling on Neutron-Induced Soft Error in SRAMs From a 250 nm to a 22 nm Design Rule, in *IEEE Transactions on Electron Devices*, vol. 57, no. 7, pp. 1527-1538, July 2010.

ENDNOTES

[11] S. Ganapathy, J. Kalamatianos, B. Beckmann, S. Raasch, Killi: Runtime Fault Classification to Deploy Low Voltage Caches without MBIST, *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Washington DC, Feb 2019.

[12] S. Ganapathy, J. Kalamatianos, K. Kasprak, and S. Raasch. 2017. On Characterizing Near-Threshold SRAM Failures in FinFET Technology. In *Proceedings of the 54th Annual Design Automation Conference 2017* (DAC '17). ACM, New York, NY, USA, Article 53, 6 pages.

[13] M. Clark, A New X86 Core for the Next Generation of Computing, *Hot Chips 2016*.

[14] H. Jeon, M. Wilkening, V. Sridharan, S. Gurumurthi, G. Loh, Architectural Vulnerability Modeling and Analysis of Integrated Graphics Processors, Workshop on Silicon Errors in Logic – System Effects (SELSE), March 2013.

[15] J. Wadden, A. Lyashevsky, S. Gurumurthi, V. Sridharan and K. Skadron, Real-world design and evaluation of compiler-managed GPU redundant multithreading, *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*, Minneapolis, MN, 2014, pp. 73-84.