RADEON™
POLARIS TECH DAY

The Polaris Architecture

Features, technologies and process

Four Principles of the New Polaris Architecture

Best Graphics Core Yet

Future-Proof Monitor Support[7]

State-of-the-Art Multimedia IP

Cool & Quiet Gaming

AMD | RADEON™

1st Generation

DirectX® 11

OpenGL 4.2

OpenCL™ 1.2

EQAA

Partially Resident Textures (PRT)

Hardware Tessellation

MLAA

Updates Over Time:

DirectX® 11.1

DirectX® 12 FL 11_1

OpenGL™ 4.5

Mantle

Vulkan™

AMD LiquidVR™

Asynchronous Compute

Virtual Super Resolution

2nd Generation

Improved compute task scheduling (DirectX® 11)

Increased geometry throughput

AMD TrueAudio

OpenCL™ 1.2

OpenGL 4.3

HSA

Updates Over Time:

AMD FreeSync™

DirectX® 11.2

DirectX® 12 FL 12_0

OpenCL™ 2.0

OpenGL™ 4.5

AMD Mantle

Vulkan™

AMD LiquidVR™

Asynchronous Compute

Virtual Super Resolution

3rd Generation

High Bandwidth Memory support

Improved tessellation performance

Delta Color Compression (DCC)

Hardware scheduler (HWS) for asynchronous compute

Native FP16/Int16 shader ops (APU)

Data parallel processing instructions

4th Generation

Improved geometry processing

AMD LiquidVR™ technology for variable resolution rendering

DCC & updated memory controller

Shader instruction pre-fetch and enhanced buffering

Real-Time and prioritized async compute with spatial and temporal scheduling

dGPU support for native FP16/Int16 Ops

AMD TrueAudio Next

H.265 Main10 decode acceleration

4K60 HEVC encode acceleration

HDMI® 2.0b

DisplayPort™ 1.3 HBR3

DisplayPort™ 1.4-HDR

HDR display support

Radeon™ RX 480 GPU at a glance

- 14nm FinFET process
- 1 Graphics Command Processor
- 4 ACE
- 2 HWS
- 36 Compute Units[8]
- 4 Geometry Processors
- 32 Pixels Output/Clock
- 144 Texture Units
- 576 32b Load/Store Units
- 2 MB L2 Cache
- 256-bit GDDR5
- AMD CrossFire™ technology
- DisplayPort™ 1.4-HDR, HDMI® 2.0b
- Video Encode/Decode acceleration
- PCIe® 3.0

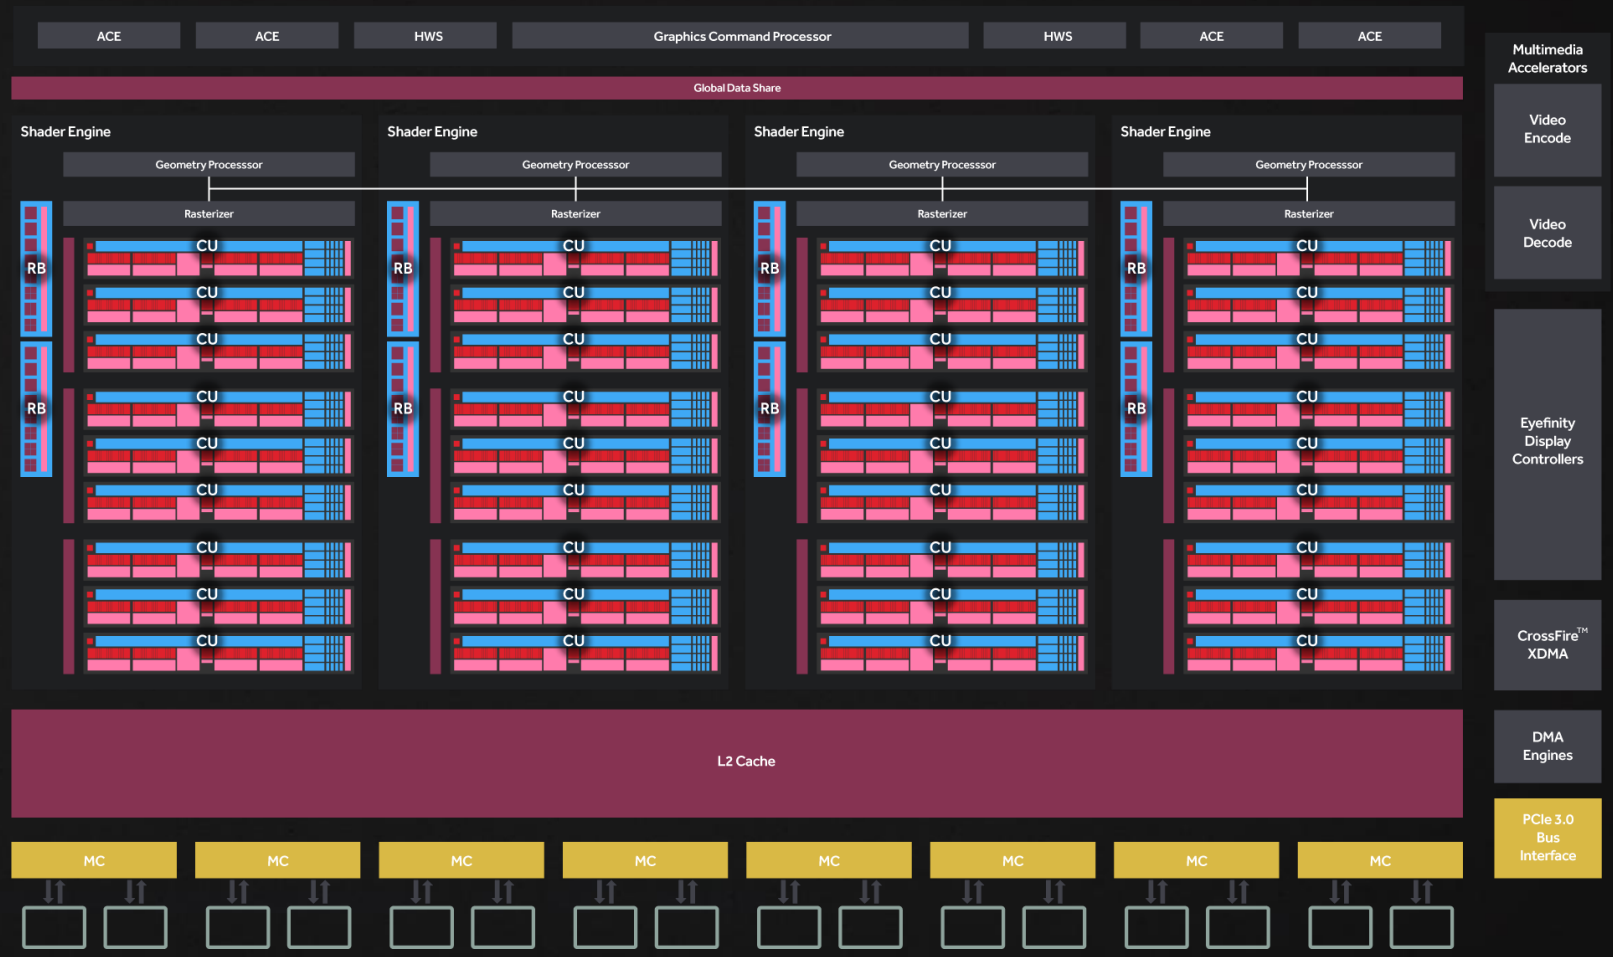

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Radeon™ RX 460 GPU at a glance

- 14nm FinFET process
- 1 Graphics Command Processor
- 4 ACE
- 2 HWS
- 16 Compute Units[8]
- 2 Geometry Processors
- 16 Pixels Output/Clock
- 64 Texture Units
- 256 32b Load/Store Units
- 1 MB L2 Cache
- 128-bit GDDR5
- AMD CrossFire™ technology
- DisplayPort™ 1.4-HDR, HDMI® 2.0b
- Video Encode/Decode acceleration
- PCIe® 3.0

4th-Gen Graphics

A Philosophy of Efficiency

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Enhanced geometry engines

- Primitive Discard Accelerator
  - Culls triangles early in the pipeline with zero area or no inclusive sample points
  - Returns grow as MSAA increases
- New Index cache for small instanced geometry

  Reduces data movement
  - Frees internal bandwidth resources
  - Maintains or improves primitive throughput during instancing

Improved shader efficiency

- Instruction prefetch
  - Improves efficiency by reducing pipeline stalls
  - Makes instruction caching more efficient
- Increased per wave instruction buffer size
  - Improve single threaded performance
- Tuned L2 cache behavior
  - Client cache request grouping
  - Improved cache and memory access efficiency
- Native FP16 and Int16 support
  - Reduced memory/register footprint and bandwidth
  - Lower power execution
  - Graphics, computer vision, data learning

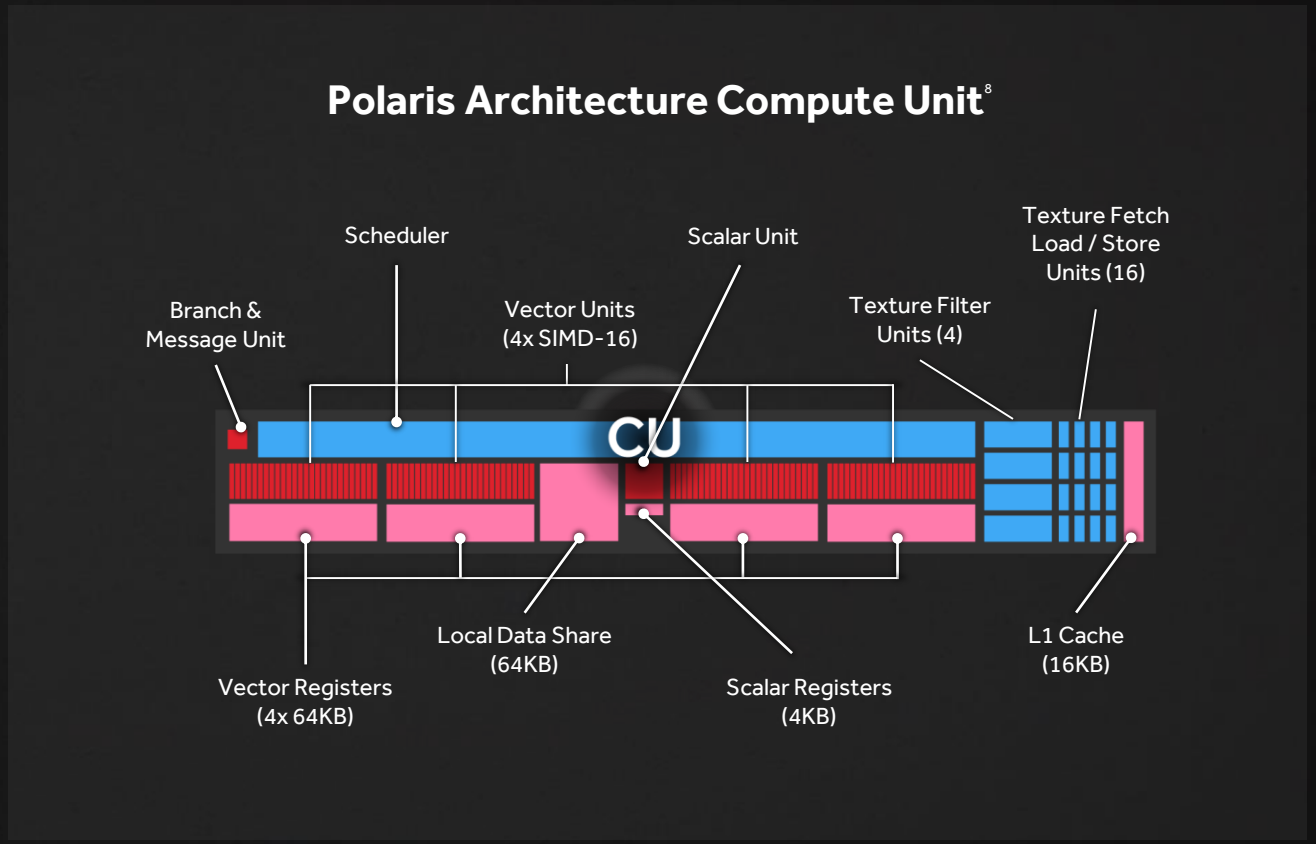

Improved shader efficiency

- Instruction prefetch
  - Improves efficiency by reducing pipeline stalls
  - Makes instruction caching more efficient
- Increased per wave instruction buffer size
  - Improve single threaded performance
- Tuned L2 cache behavior
  - Client cache request grouping
  - Improved cache and memory access efficiency
- Native FP16 and Int16 support
  - Reduced memory/register footprint and bandwidth
  - Lower power execution
  - Graphics, computer vision, data learning

UP TO

+15%

PERFORMANCE PER CU

RADEON™ RX 480 VS. RADEON™ R9 290[1]

GPUOpen shader intrinsic functions

- Direct access to inline ISA
  - Data parallel primitive access
  - Detect dynamically uniform conditions at runtime
- Enables performance-critical functions
  - Execution exactly as developers intend
- **Available now:** DirectX® 11, DirectX® 12 and Vulkan™

Console to PC

Console style programming enabled on PC to help extract more performance from the GPU

Memory & delta color compression engines

- Updated memory controller & PHY
  - Supports up to 8Gbps GDDR5 memory devices
  - Up to 256GB/s memory bandwidth
- Lossless DCC
  - Full 2/4/8 : 1 compression ratios supported
  - Driver optimized for utilization of DCC
  - Provides perf/bandwidth and power advantages
  - Extends the life, viability and efficiency of GDDR5
- Efficient balance of resources
  - Yields higher peak bandwidth

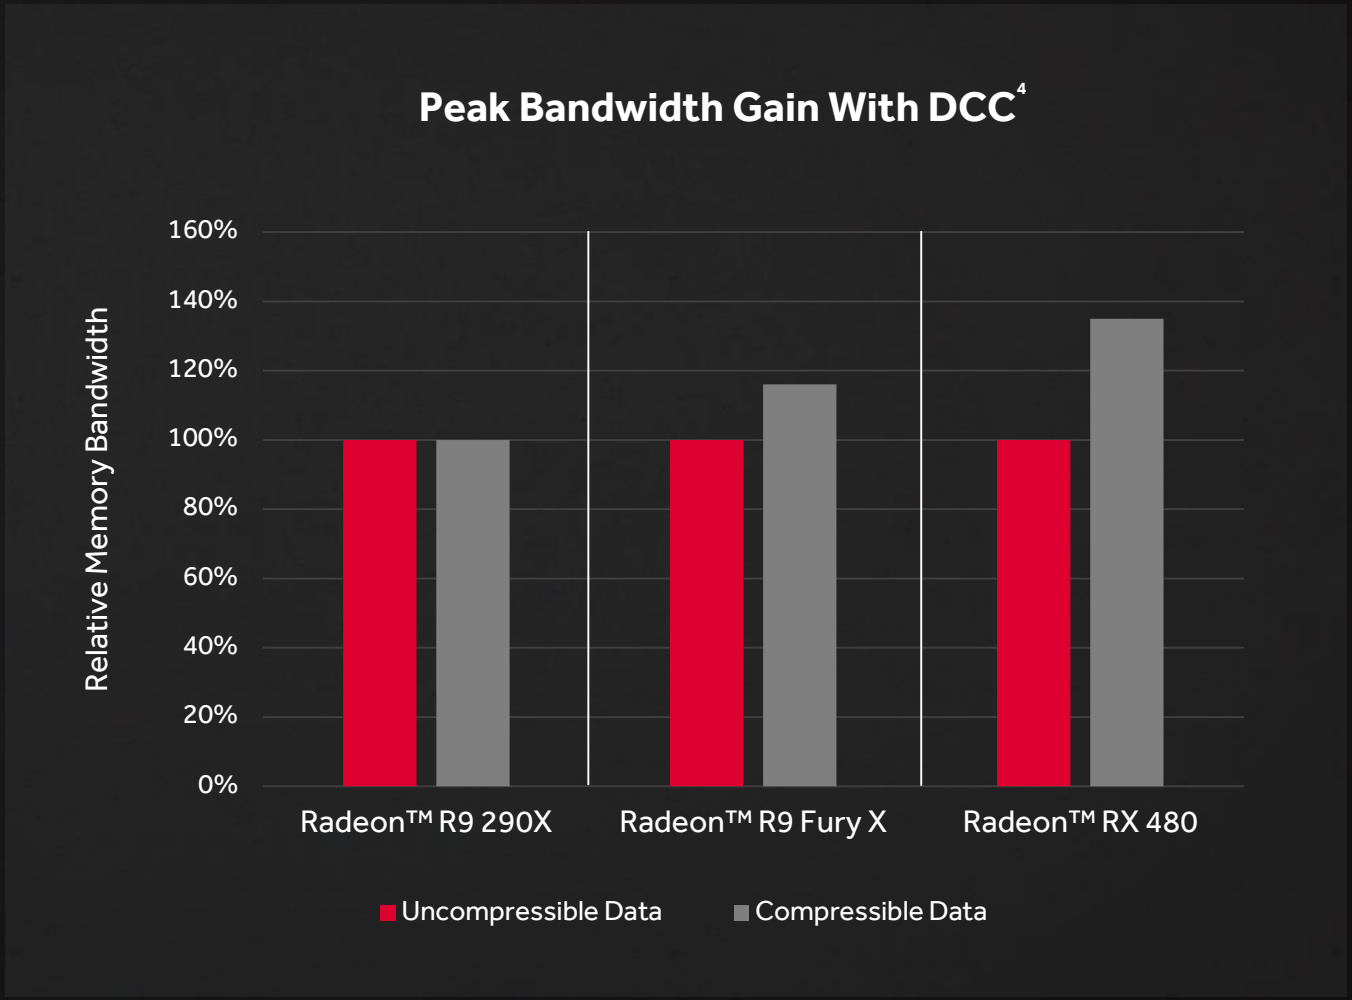

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Large L2 cache

- Doubled L2 cache
  - Reduces external memory bandwidth requirements
  - Improves power efficiency
  - Increases benefits of DCC
- Up to 40% power savings on memory transactions[5]

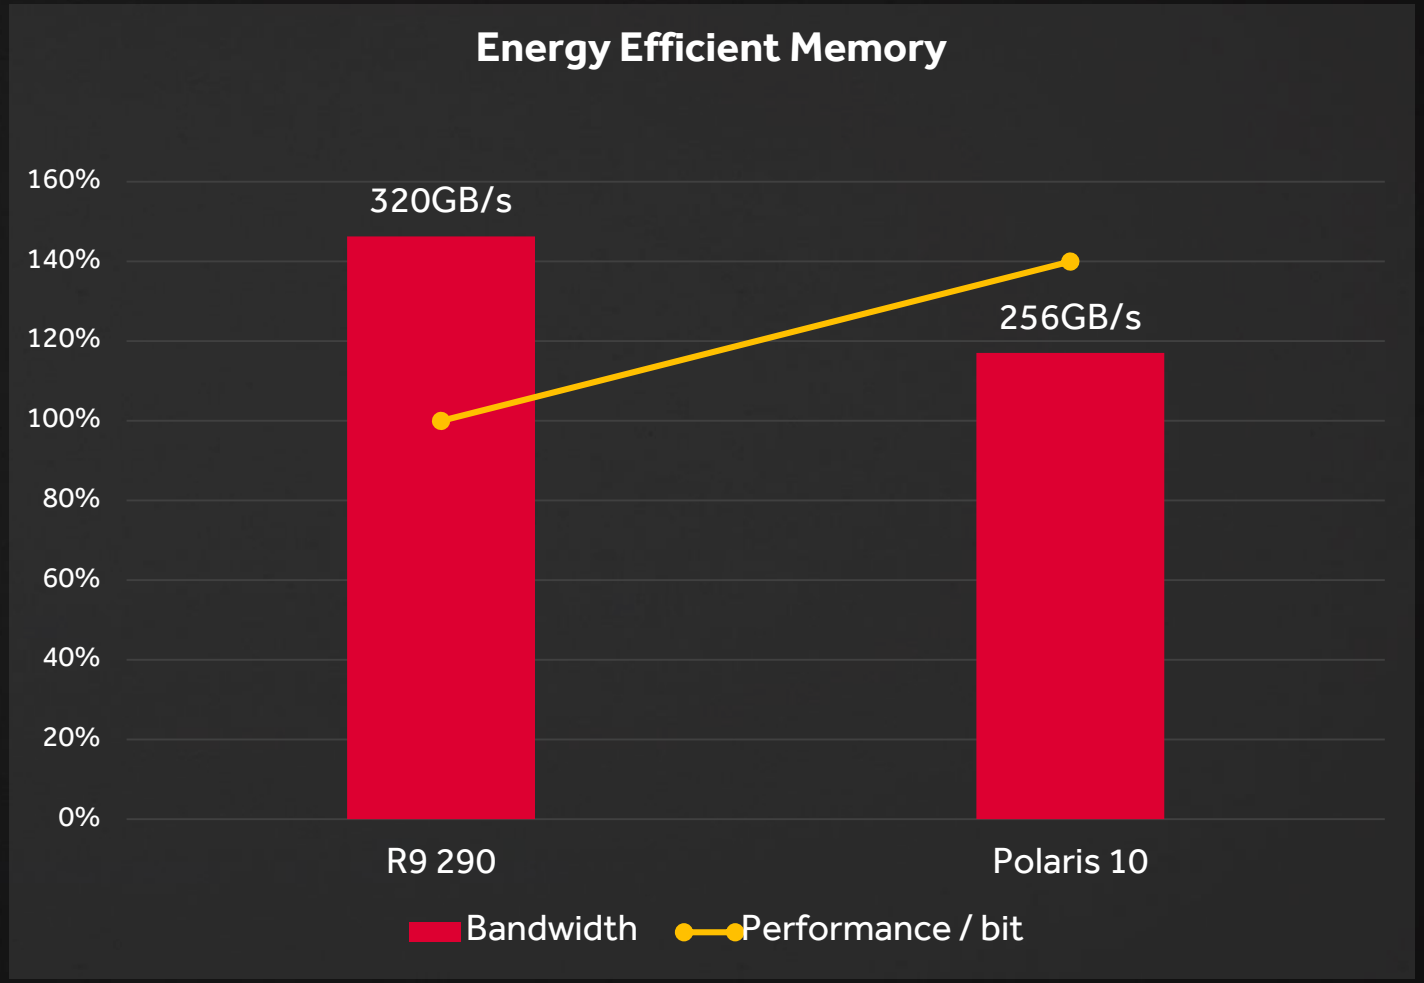

Other Graphics Architectures In High-Level APIs

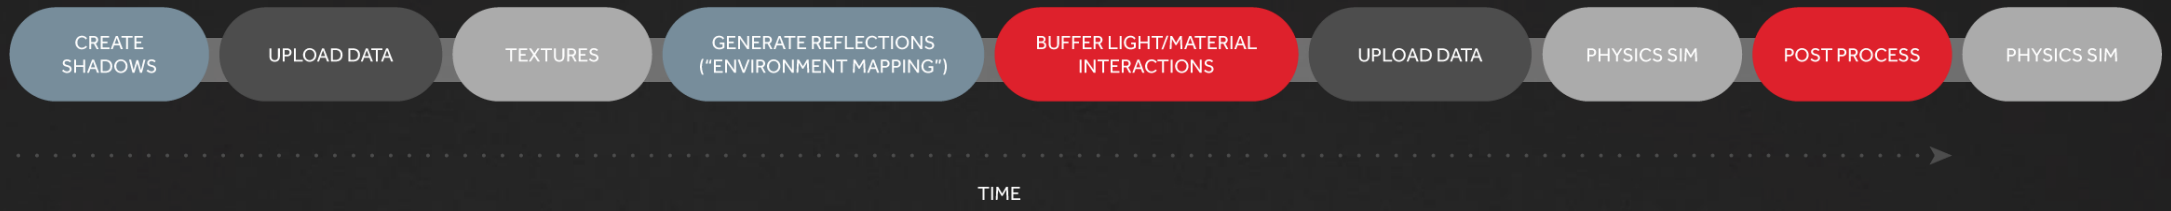

Radeon™ Graphics in DirectX® 12, Vulkan™, and VR

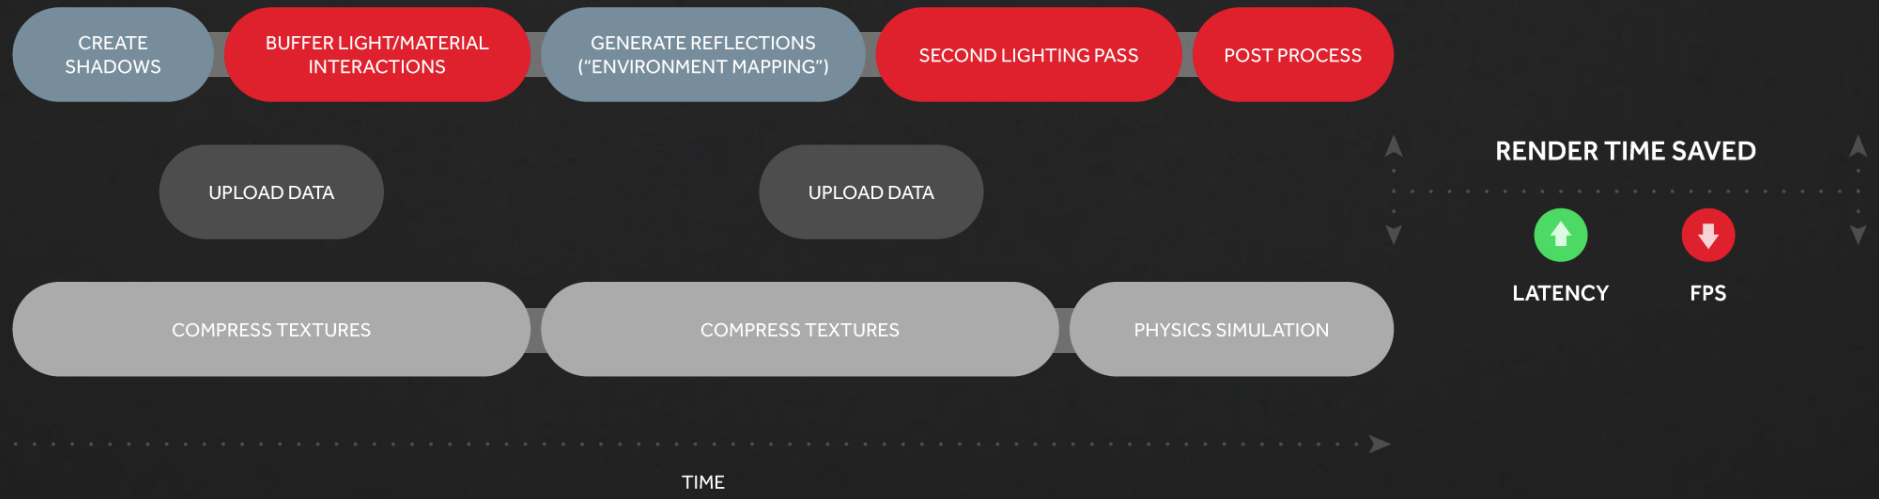

Radeon™ asynchronous compute is incredibly flexible

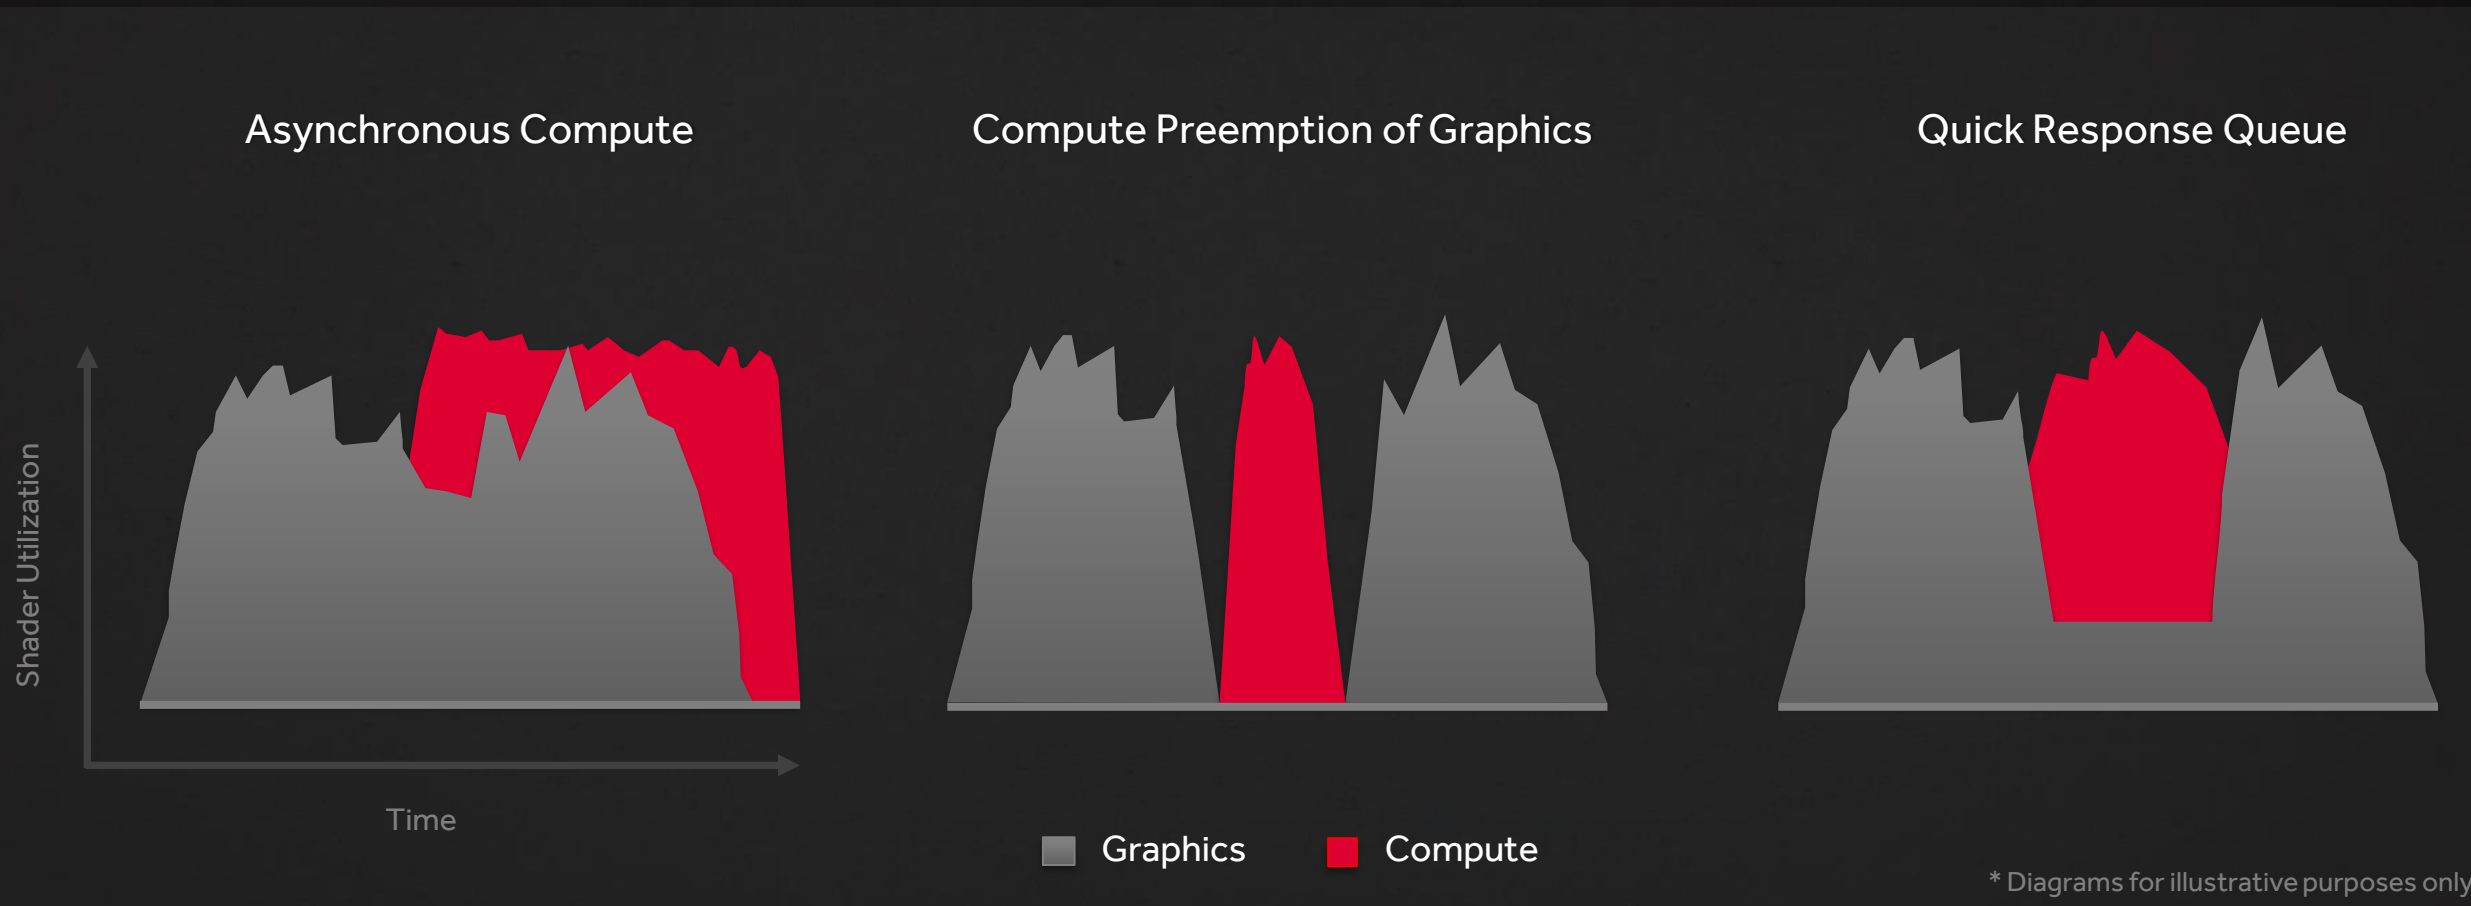

Hardware scheduler (HWS) for asynchronous compute

- CPU kernel mode driver scheduling task offload
- Real time and prioritized queues (audio/VR)
- Temporal and spatial resource management
- Concurrent task or process scheduling
- Coordination and load balancing with graphics
- Functionality can be updated with microcode

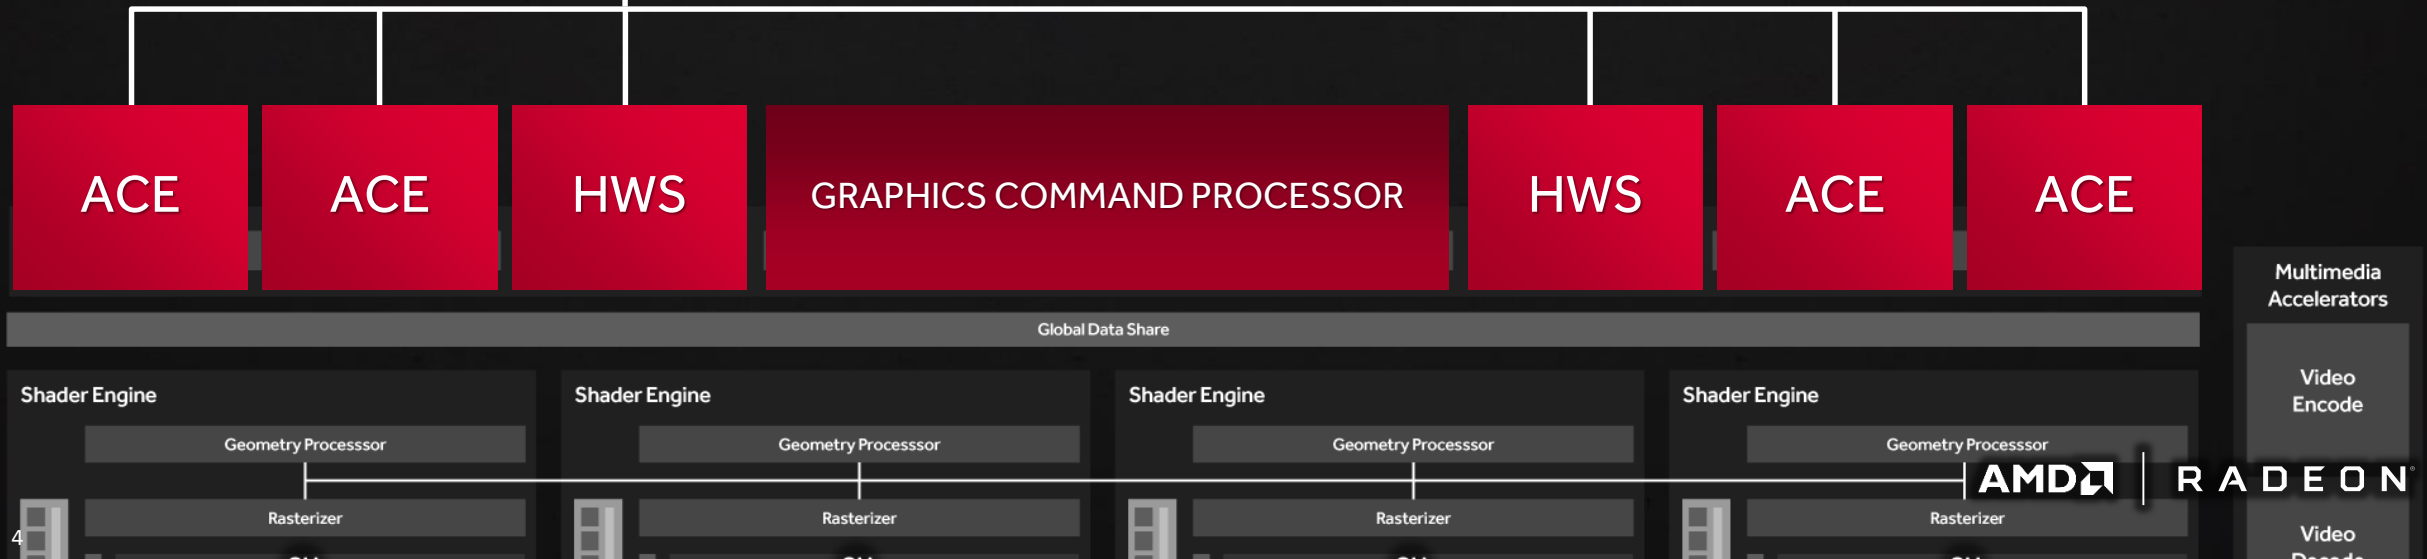

AMD | RADEON

Polaris Architecture Summary

UP TO

+15%

PERFORMANCE PER CU

RADEON™ RX 480 VS. RADEON™ R9 290[1]

UP TO

2.8x

PERFORMANCE PER WATT

RADEON™ RX 470 VS. RADEON™ R9 270X[2]

FinFET 14 Technology | Micro–Architecture Tuned for Gaming | Physical Design Optimizations | DPM Techniques

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Future-Proof Display Support

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

DisplayPort™ 1.3 HBR3 and DisplayPort™ 1.4-HDR

- Polaris is ready for DisplayPort™ 1.3 HBR3
  - Existing cables and connectors utilized
  - Up to 32.4Gbps of bandwidth
  - +80% bandwidth over HDMI® 2.0b
- Polaris is ready for DisplayPort™ 1.4-HDR
  - Up to 10-bit 4K96
  - CTA-861.3 HDR metadata transport
  - Rec.2020 color space support
  - SMPTE 2084 EOTF

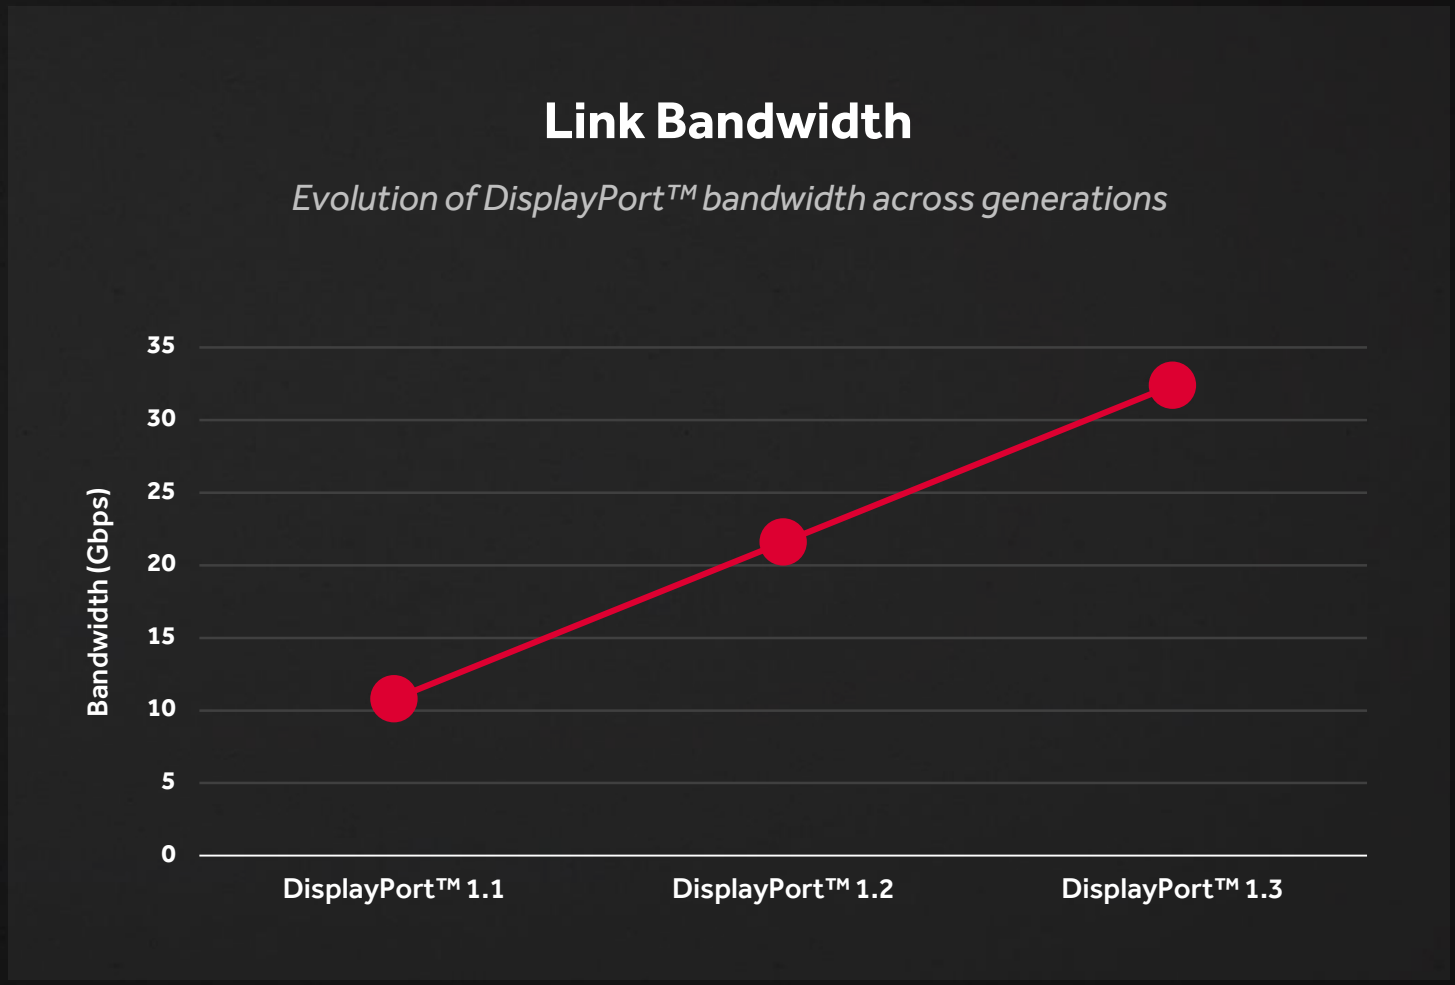

DisplayPort™ 1.3 is a major leap for 4K FreeSync™

- Polaris can enable 4:4:4 4K displays at up to 120Hz
- We anticipate DRR-ready 4K120 panels by 4Q2016
- DRR 4K120 panels will enable 4K AMD FreeSync™ monitors compatible with 30-120 FPS coverage

DisplayPort™ 1.3 enables single cable 5K60

- Radeon™ RX 400 Series graphics will enable single-cable 4:4:4 5K60 resolution
- 5K is 78% more pixels than 3840x2160 (4K)
- 210ppi of 5K/28" provides incredible sharpness and working space vs 142ppi of 4k/28"
- We anticipate displays ready for single-cable 5K by year end

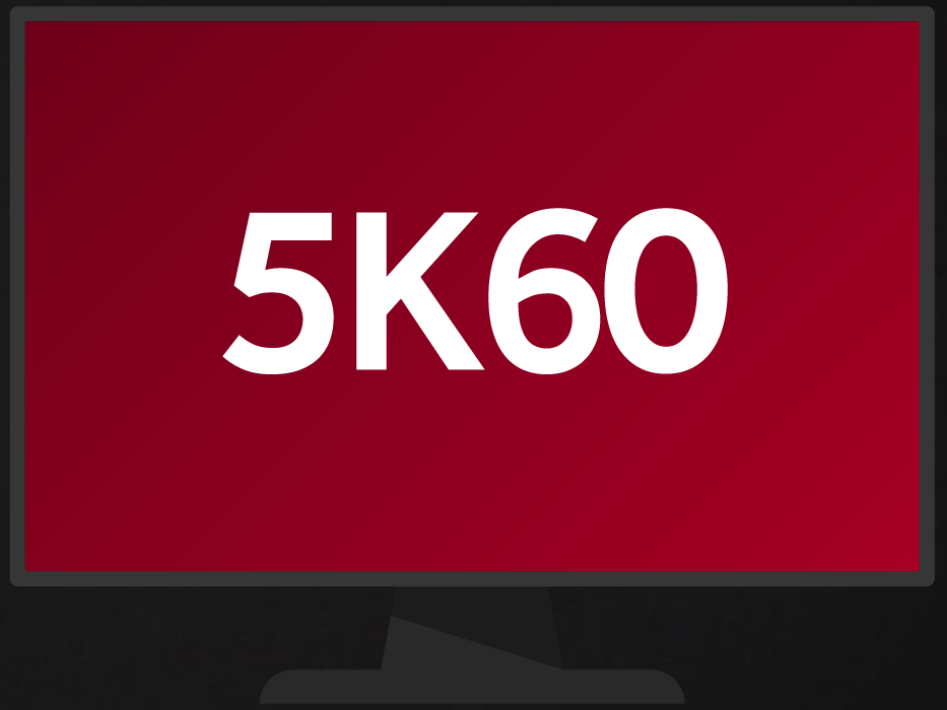

DisplayPort™ 1.3 is future-proof for gaming monitors[7]

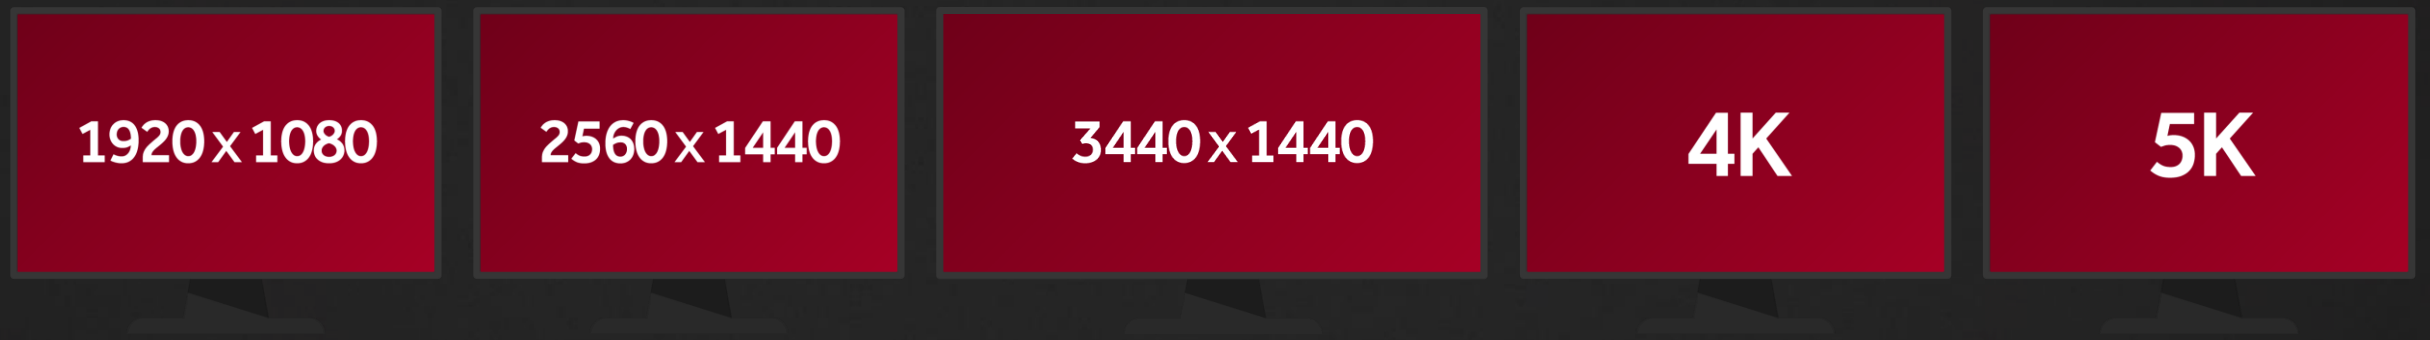

1920 x 1080	2560 x 1440	3440 x 1440	4K	5K
Up to **240Hz SDR**	Up to **240Hz SDR**	Up to **190Hz SDR**	Up to **120Hz SDR**	Up to **60Hz SDR**
Up to **240Hz HDR**	Up to **200Hz HDR**	Up to **155Hz HDR**	Up to **96Hz HDR**	

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

The Polaris architecture is ready for home theater PCs

4K60

HDMI™
HIGH-DEFINITION MULTIMEDIA INTERFACE

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

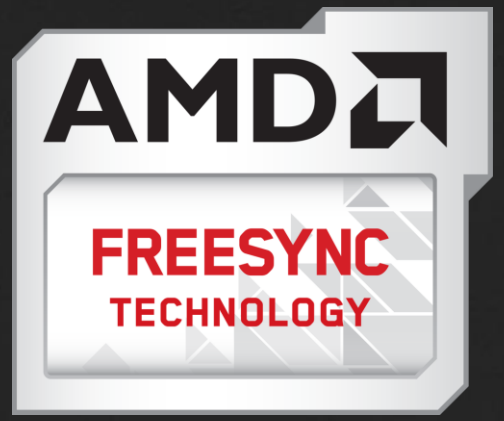

AMD FreeSync™ Technology over HDMI® 2.0b

Supported on 28 displays as of May 28th, 2016

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Acer | LG | Mstar Semiconductor | Novatek | Realtek | Samsung

AMD FreeSync™ Technology over HDMI® 2.0b

Technology Partners

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Polaris enables better pixels with HDR

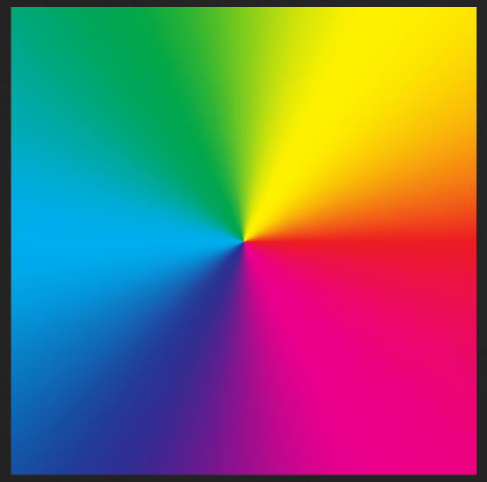

Larger Color Spaces

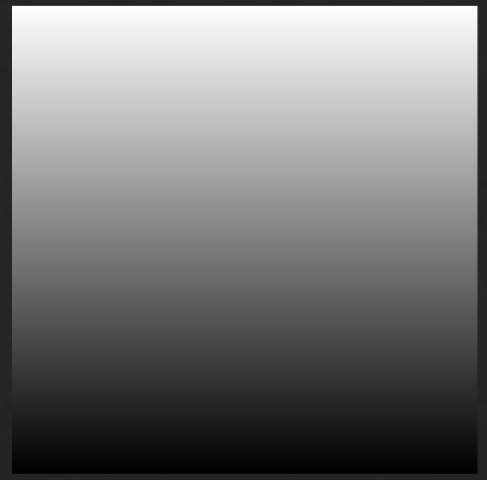

Higher Contrast Ratios

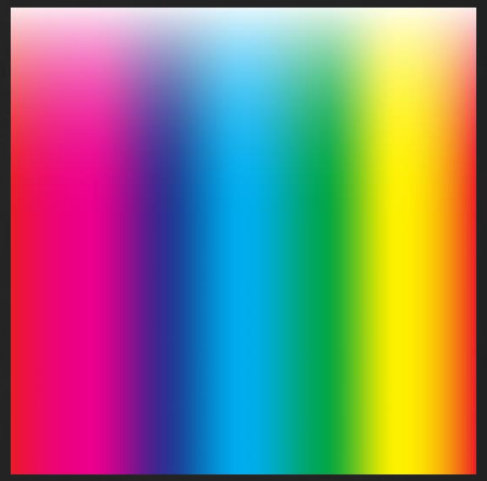

Higher Peak Luminance

HDR drives more colors

- Typical monitors are sRGB (blue triangle)
- Short-term HDR goal is to reach P3 Digital Cinema (yellow triangle)
- Long-term HDR goal is to reach Rec. 2020 (red triangle)
  - 2X visible colors vs. sRGB
  - 75% coverage of the human visual spectrum
  - Red/gold/cyan notably improved

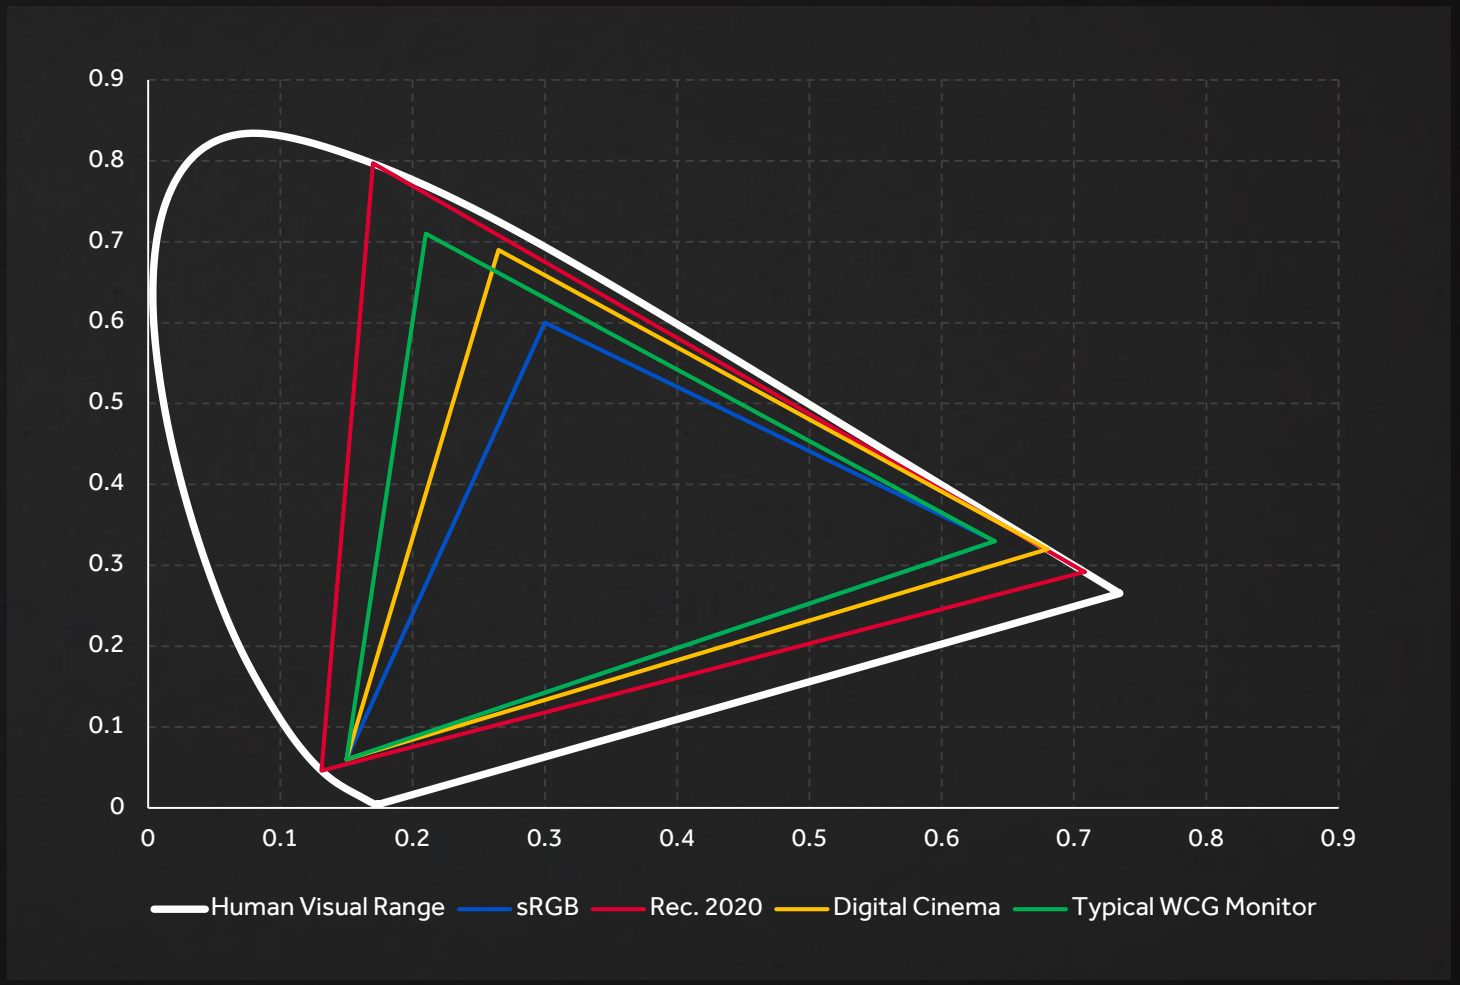

HDR replicates the human vision system

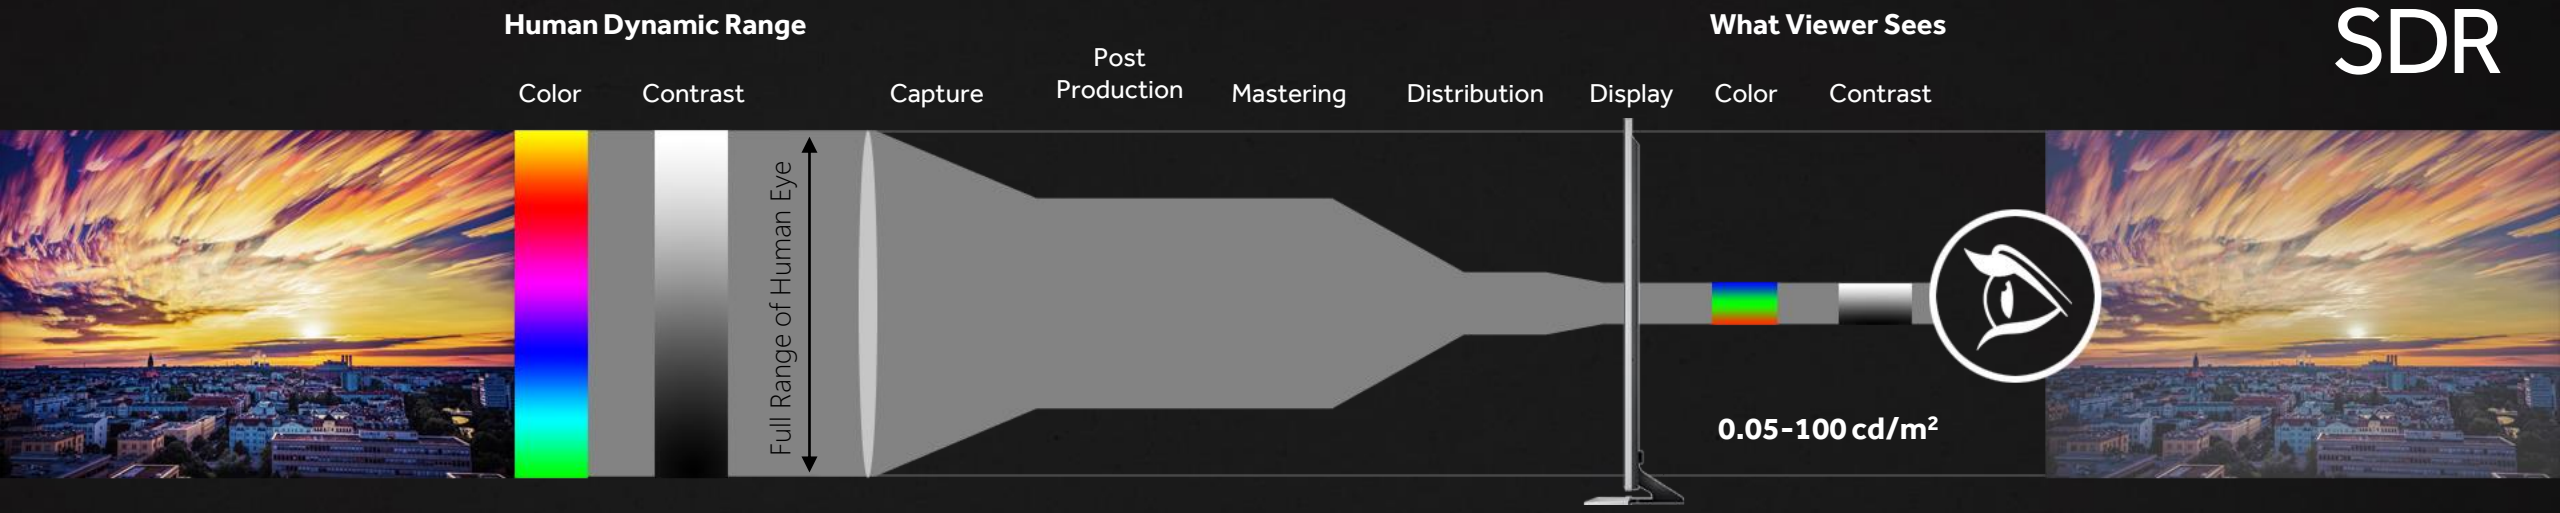

Under embargo until June 29, 2016 at 9 a.m. EST. SDR vs. HDR imagery simulated for illustrative purposes only.

AMD | RADEON™

HDR drives higher luminance

Sunlight	1,600,000,000 nits
Fluorescent light	10,000 nits
Highlights	1,000-10,000 nits
White Range	250-1,000 nits
Typical Objects	1-250 nits
Shadow Details	0.01-1 nit
Ultra Blacks	0-0.01 nit

** Image courtesy "Russellstreet" under CC BY-SA 2.0 license from Flickr. Scene comparison simulated.*

HDR is better at translating bits to photons

- Electro-Optical Transfer Functions are how we convert digital signals to light
- Today's consumer content is designed around CRT television EOTFs
  - Based on work started in 1930s
  - Only considers up to 100 nits of brightness
  - Wastes encoding bits and crushes black levels
- 10-bit ST 2084 is for HDR
  - Designed for the next 20 years of displays
  - Consistent encoding 0-10,000 nits
  - Dramatically improves detail/efficiency in black levels

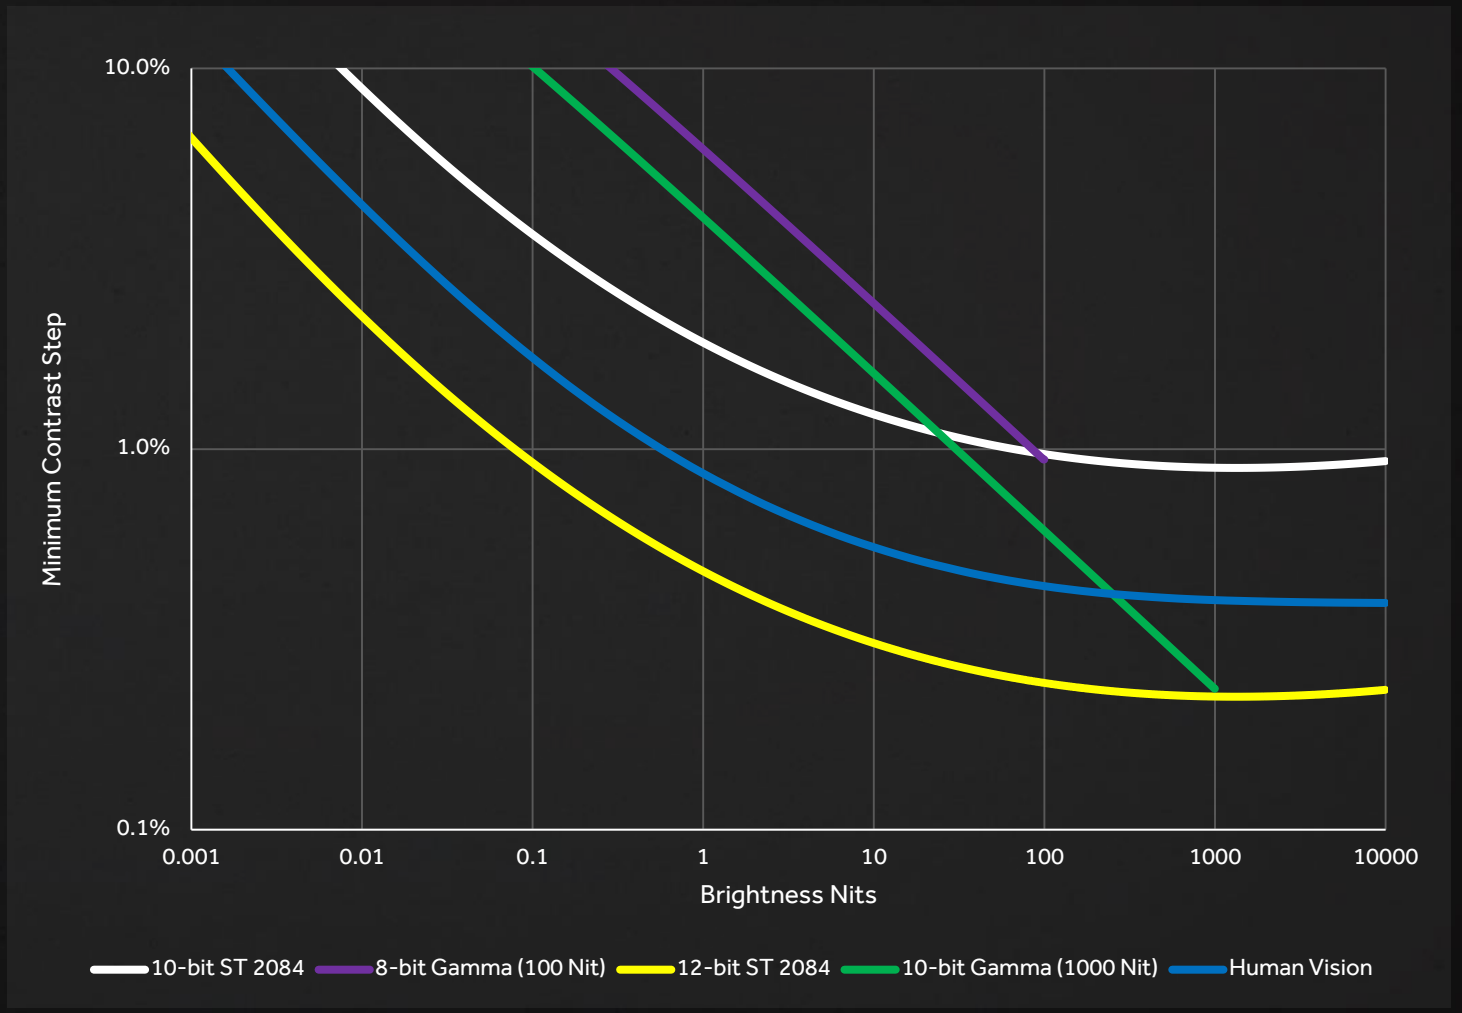

Polaris architecture display pipe

- Polaris is ready for 10-bit and 12-bit HDR
  - First-gen HDR displays will be 10-bit
  - 12-bit will exceed human visual acuity
- Display color processing engine
  - Programmable gamut remapping
  - Programmable gamma control
  - Floating point processing
  - Enables 1:1 mapping with display's native capabilities
  - Eliminates in-display remapping latency

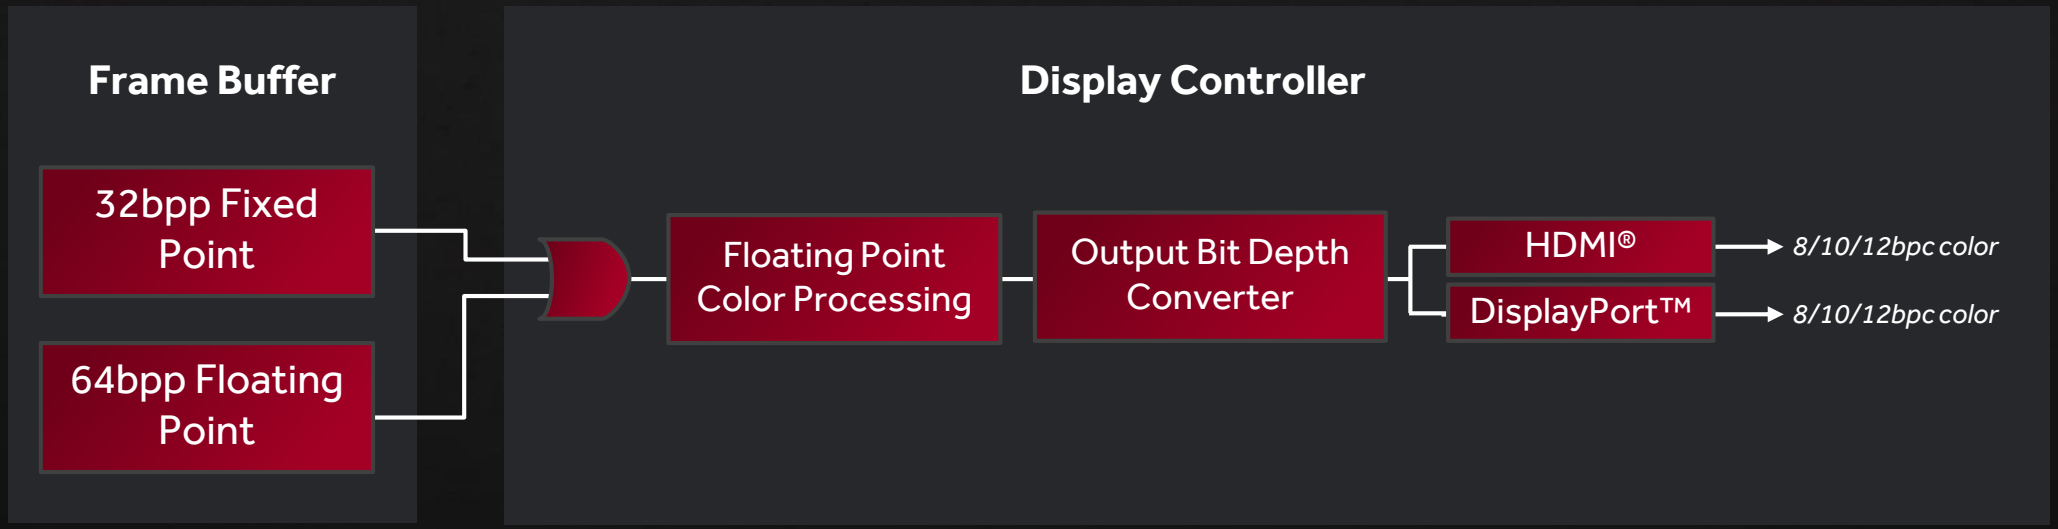

Radeon™ Photon SDK for HDR gaming

- Tonemapping should be done by the game engine
  - Tonemapping in the display adds significant input lag
- We poll color/contrast/nits from the display
  - Polling fed back to engine tonemapper
  - Output from GPU optimized for the connected display
- Switching to HDR tonemapping not a large hurdle
  - Game engines already do HDRR tonemapped to SDR
- Available for developers now:
  - DirectX® 11 HDR driver support available for developers
  - DirectX® 12 support planned in a future driver update
  - HDR Video playback support available for developers

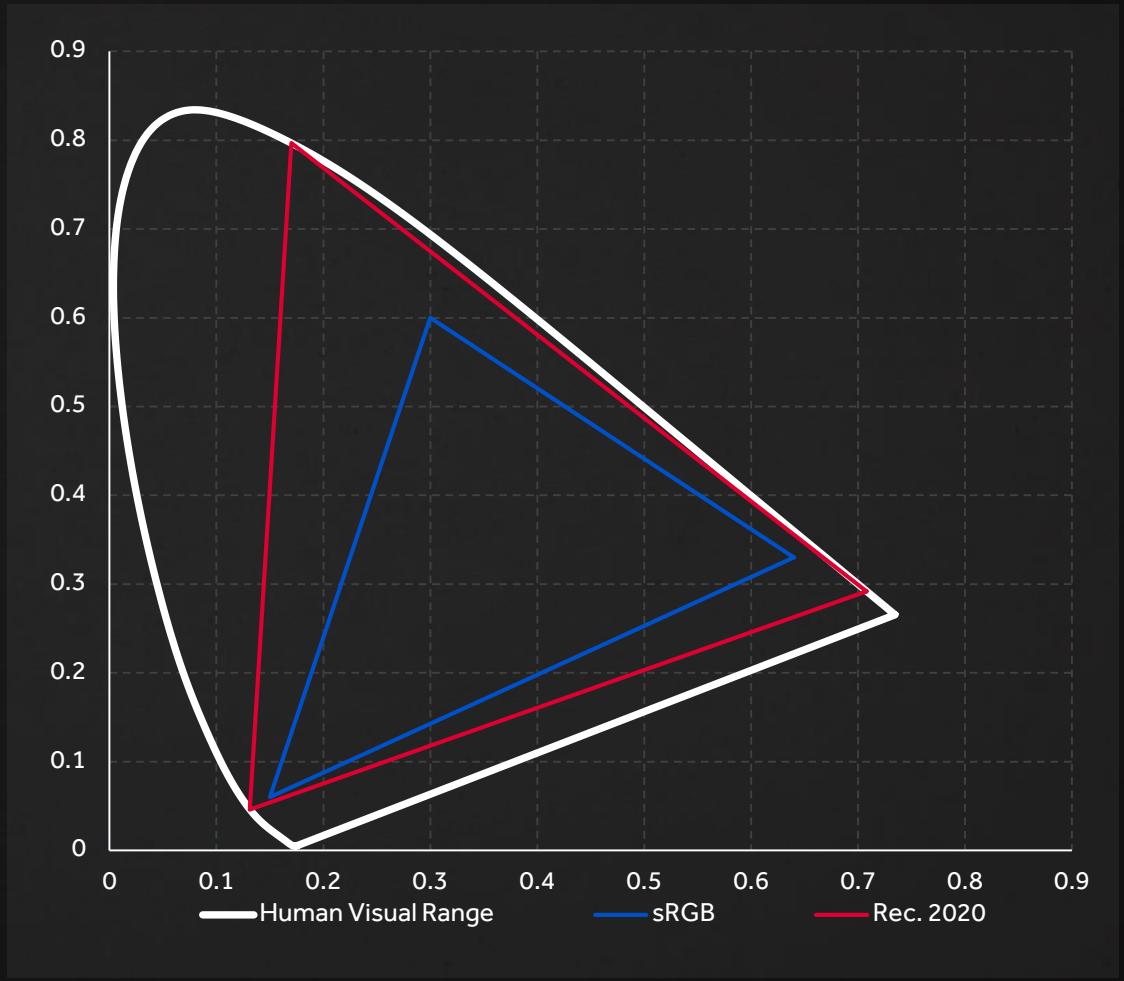

Polaris architecture HDR display support

HDMI® 2.0b (with HDCP 2.2)	**1920x1080 @ 192Hz** **2560x1440 @ 96Hz** **3840x2160 @ 60Hz (4:2:2)**
DisplayPort™ 1.4-HDR (with HDCP 2.2)	**1920x1080 @ 240Hz** **2560x1440 @ 192Hz** **3840x2160 @ 96Hz**

Accelerated Video Encoding & Decoding

A history of AMD's video encode accelerator

Radeon™ HD 7970
H.264 up to 1080p75

Radeon™ R9 290X
H.264 up to 1080p75

Radeon™ R9 285
H.264 up to 4K60

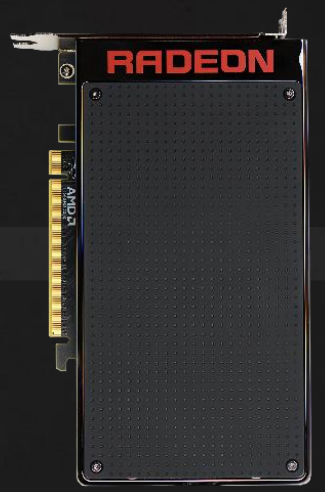

Radeon™ R9 Fury X
H.264 up to 4K30

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Video encode acceleration with Polaris

H.264	HEVC*
1080p120	1080p240
1440p60	1440p120
4K30	4K60

Under embargo until June 29, 2016 at 9 a.m. EST.

* HEVC acceleration is subject to inclusion/installation of HEVC-compatible applications.

Improving stream quality with Polaris

- Fast HQ encode is a tough problem
  - Rapid scene changes
  - Non-uniform complexity distribution
- Achieving fast quality with 2-pass encoding
  - Downscaled pre-encode guides rate control
  - Picture-level analysis for efficient budget allocation
  - Macroblock-level analysis for efficient QP selection
- 2-pass encoding yields up to 1dB Y-PSNR gain
  - Works with H.264 and HEVC content

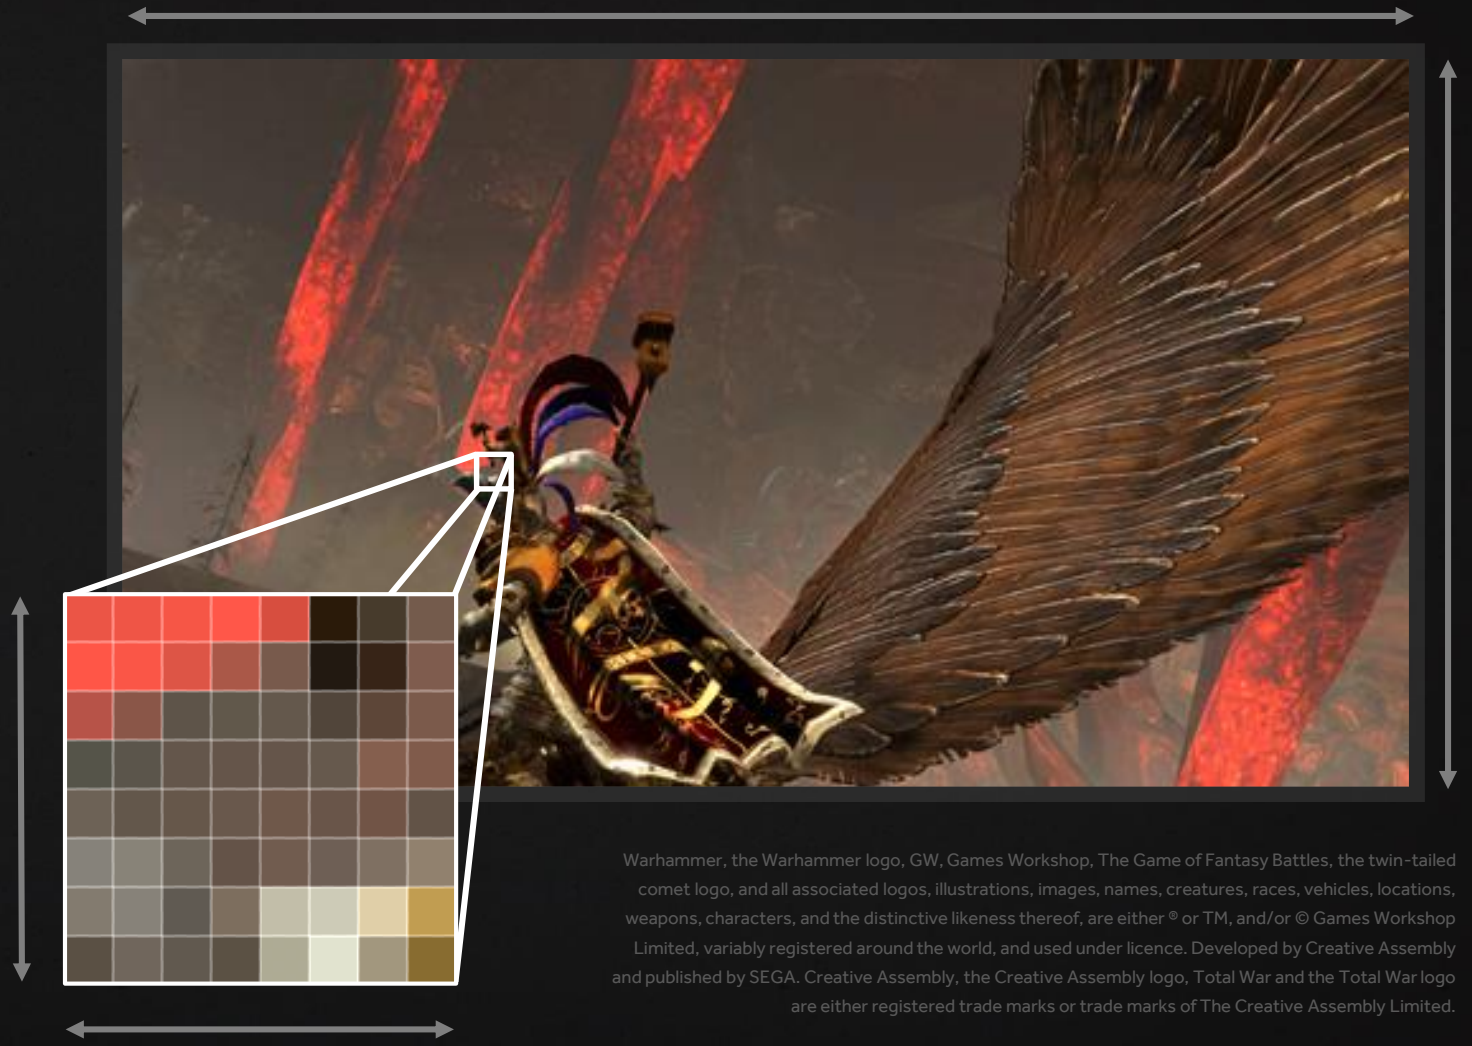


1-Pass Encoding
2-Pass Encoding
AMD
RADEON

Leveraging the Polaris architecture encode accelerator for games

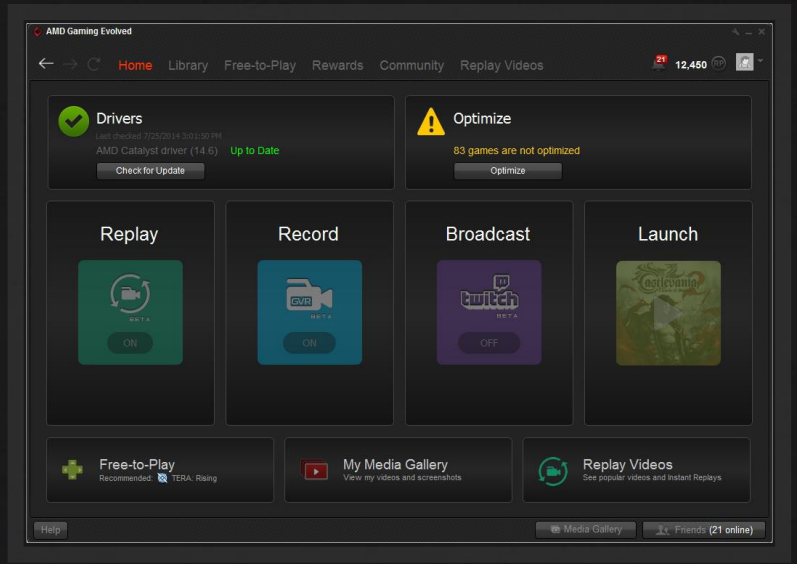

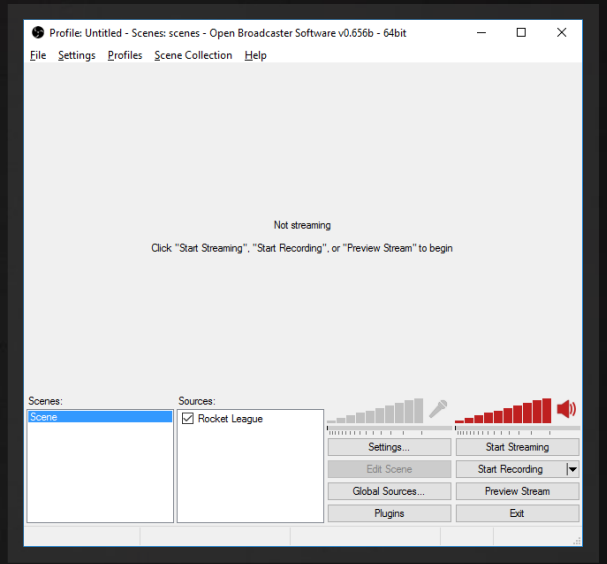

PLAYS.TV

GAMING EVOLVED CLIENT
POWERED BY raptr

Open Broadcaster Software

Under embargo until June 29, 2016 at 9 a.m. EST.

* Use of third party logos is for informational purposes only and no endorsement of or by AMD is intended or implied.

AMD | RADEON™

A history of AMD's video decode accelerator

Radeon™ HD 7970
H.264 up to 1080p60
VC1 up to 1080p60
MP4-P2 up to 1080p60

Radeon™ R9 290X
H.264 up to 1080p60
VC1 up to 1080p60
MP4-P2 up to 1080p60

Radeon™ R9 285
H.264 up to 4K120
VC1 up to 1080p60
MP4-P2 up to 1080p60

Radeon™ R9 Fury X
H.264 up to 4K120
Main HEVC up to 4K60
VC1 up to 1080p60
MP4-P2 up to 1080p60

Our most versatile decode accelerator yet

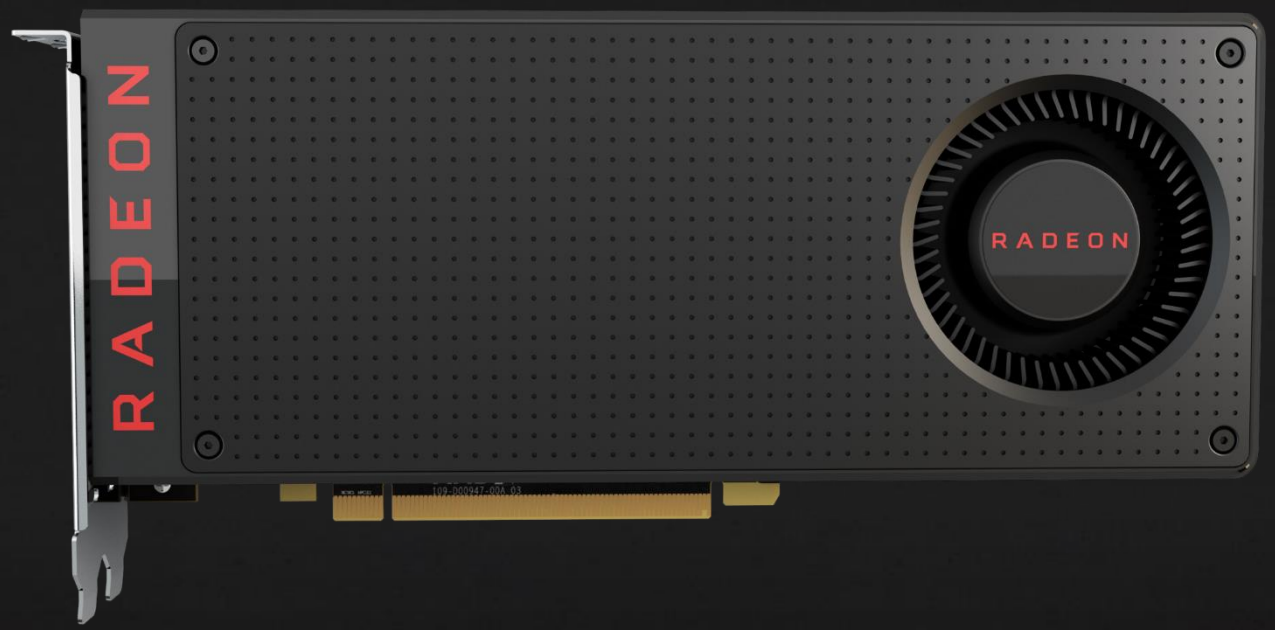

HEVC	VP9	MJPEG	H.264	MP4-P2	VC1
Up to 4K60 Main-10	Up to 4K[9]	Up to 4K30	Up to 4K120	Up to 1080p60	Up to 1080p60

The Polaris architecture accelerates the best content

Beautiful HDR Movies
(10-bit HEVC)

NETFLIX amazon

Your Favorite Streamers
(H.264)

twitch

More Cat Videos
(VP9)

YouTube

Razor-Sharp Video Calling
(HEVC & MJPEG)

skype™

* Logos for illustrative purposes only; no endorsement implied.

Optimizing FinFET 14

It's been a long road

90nm	80nm	65nm	55nm	40nm	28nm	FinFET
2005	2006	2007	2007	2009	2011	2016

Drive to lower power consumption

- Transistors per area have doubled effectively every 2 years
  - Sub-nm range, static leakage/area also doubling
  - Voltage reductions from technology flattened
- Many PD techniques can reduce leakage; mitigate poor scaling
  - Power gating
  - Voltage islands
  - Back bias or advanced circuits
- Unfortunately: a poor substitute for native process benefits

FinFET 14 performance vs. 28nm planar

- FinFET offers product level performance advantage over 28A planar
  - Significant variation reduction and leakage power improvement
- FinFETs have fundamentally lower variation than planar bulk

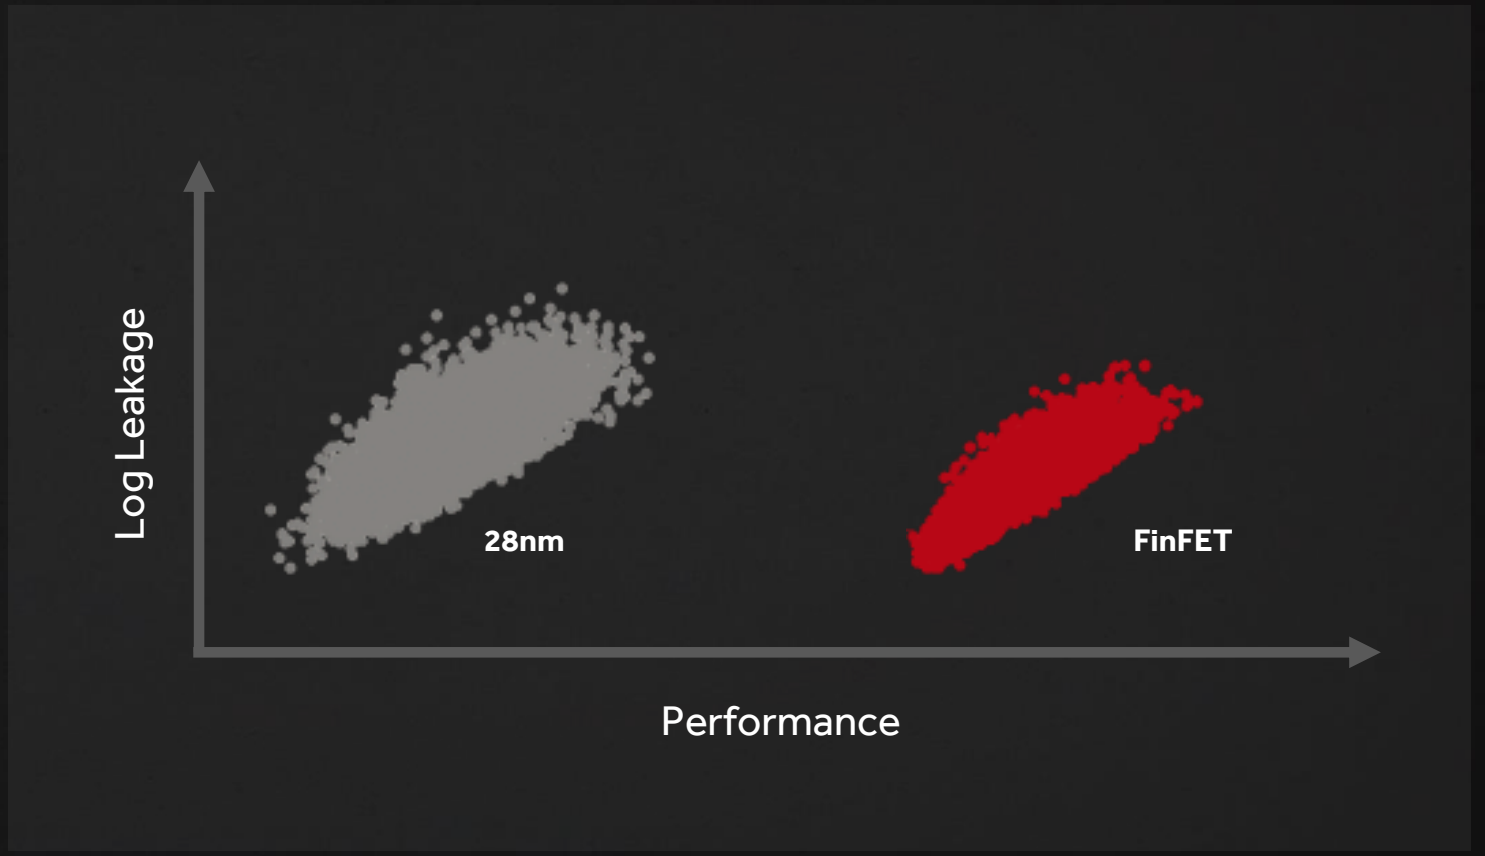

* Diagram for illustrative purposes.

FinFET 14 has better variation

- Lower variation improves product performance and power
  - Performance set by the slowest device and the required voltage
  - Power set by the leakiest device
- FinFETs have fundamentally lower variation than planar bulk

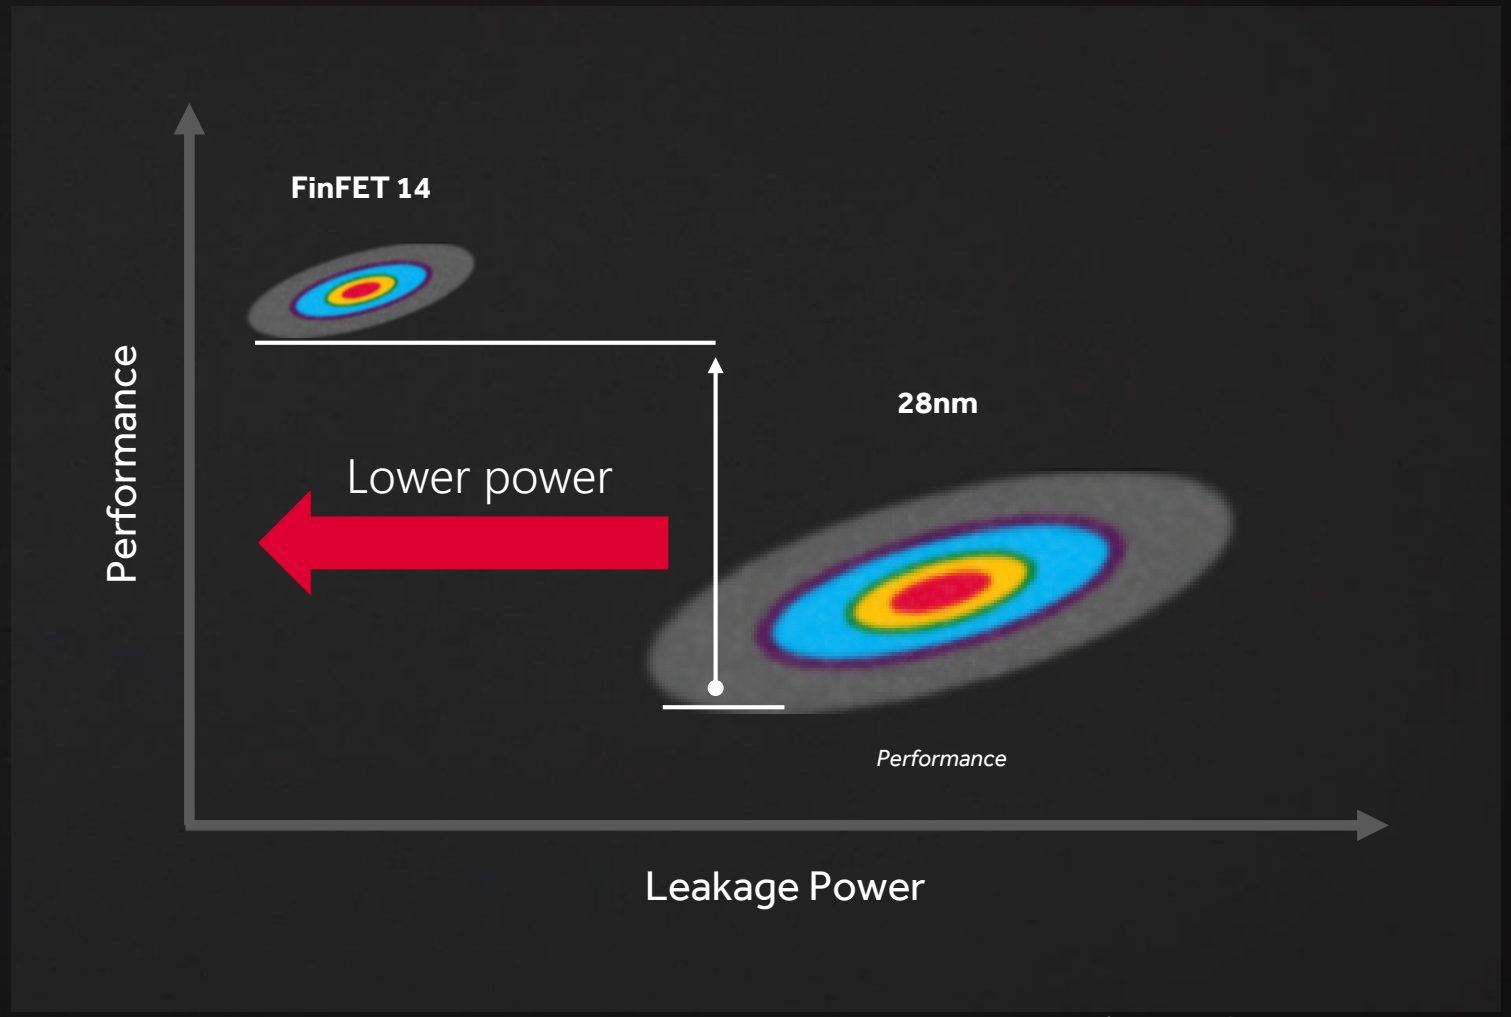

* Diagram for illustrative purposes.

FinFET performance/power benefit

- Bottom line
  - FinFET devices provide fundamental performance / power benefit over planar

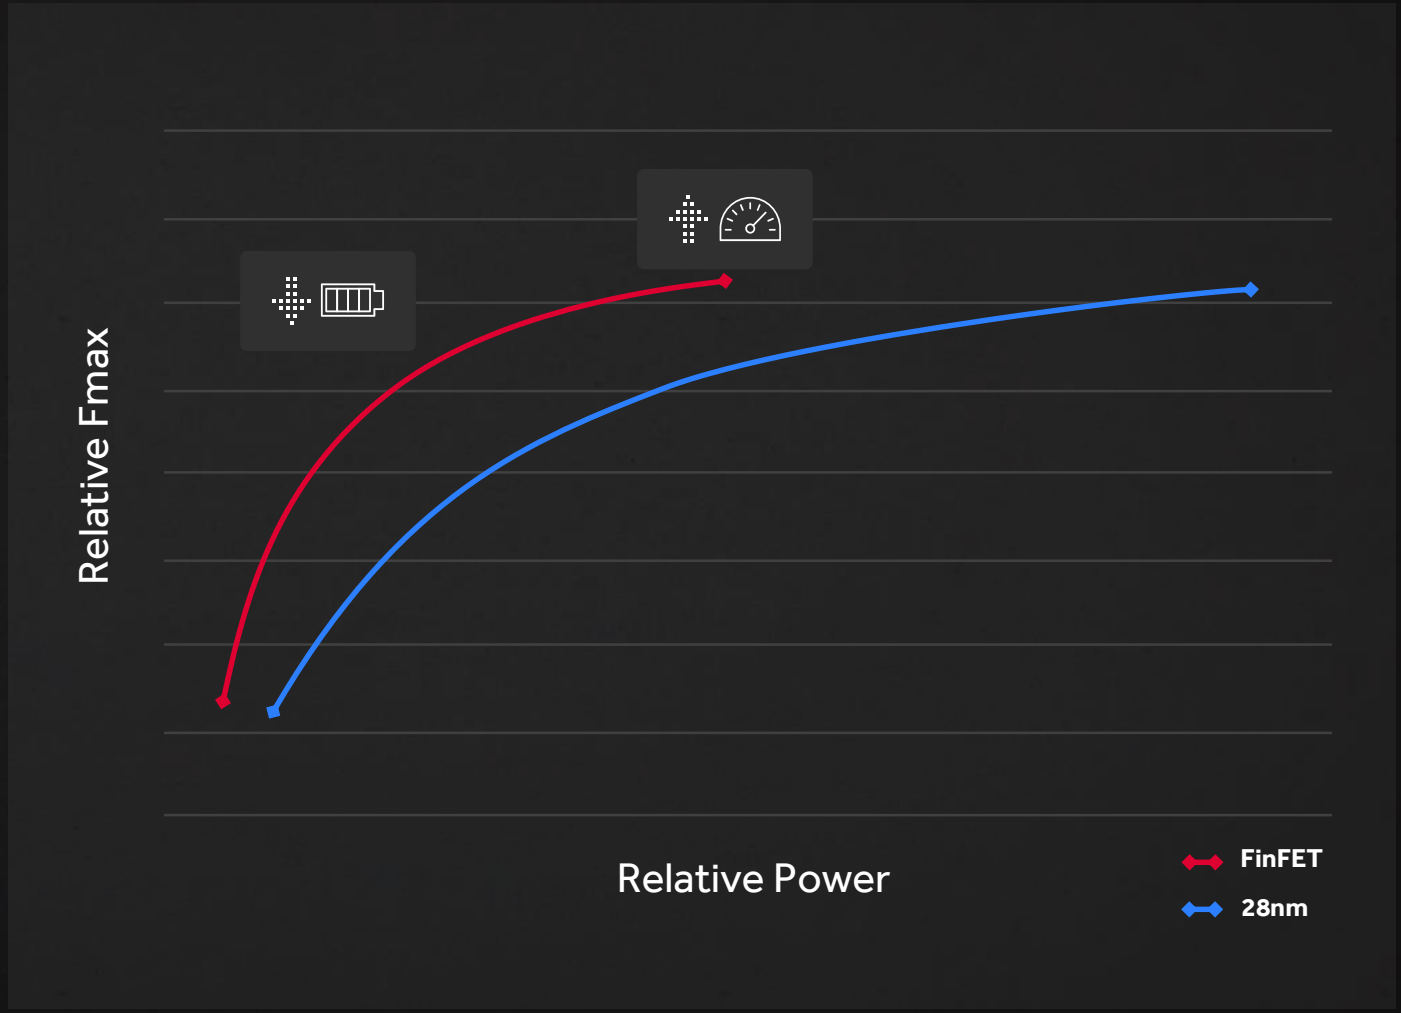

* Diagram for illustrative purposes.

Summarizing FinFET 14

- FinFETs are a very significant step
  - Reduced operating voltages
  - Reduced leakage
  - Fundamental improvement in power efficiency

- FinFETs enable new product categories
  - Thin & light gaming notebooks
  - Small form factor desktops
  - Discrete cards with fewer power connectors

Design Approaches for Power Efficiency

Adaptive clocking

- All GPUs operate at low voltages with high current
  - Challenging to deliver quality voltage from packages and power supplies
- Voltage variations can be 10%-15% of the nominal value
  - Average voltage must be raised to cover this variation, which wastes a lot of power
- AMD's adaptive clocking recovers that waste with up to a 25% power reduction

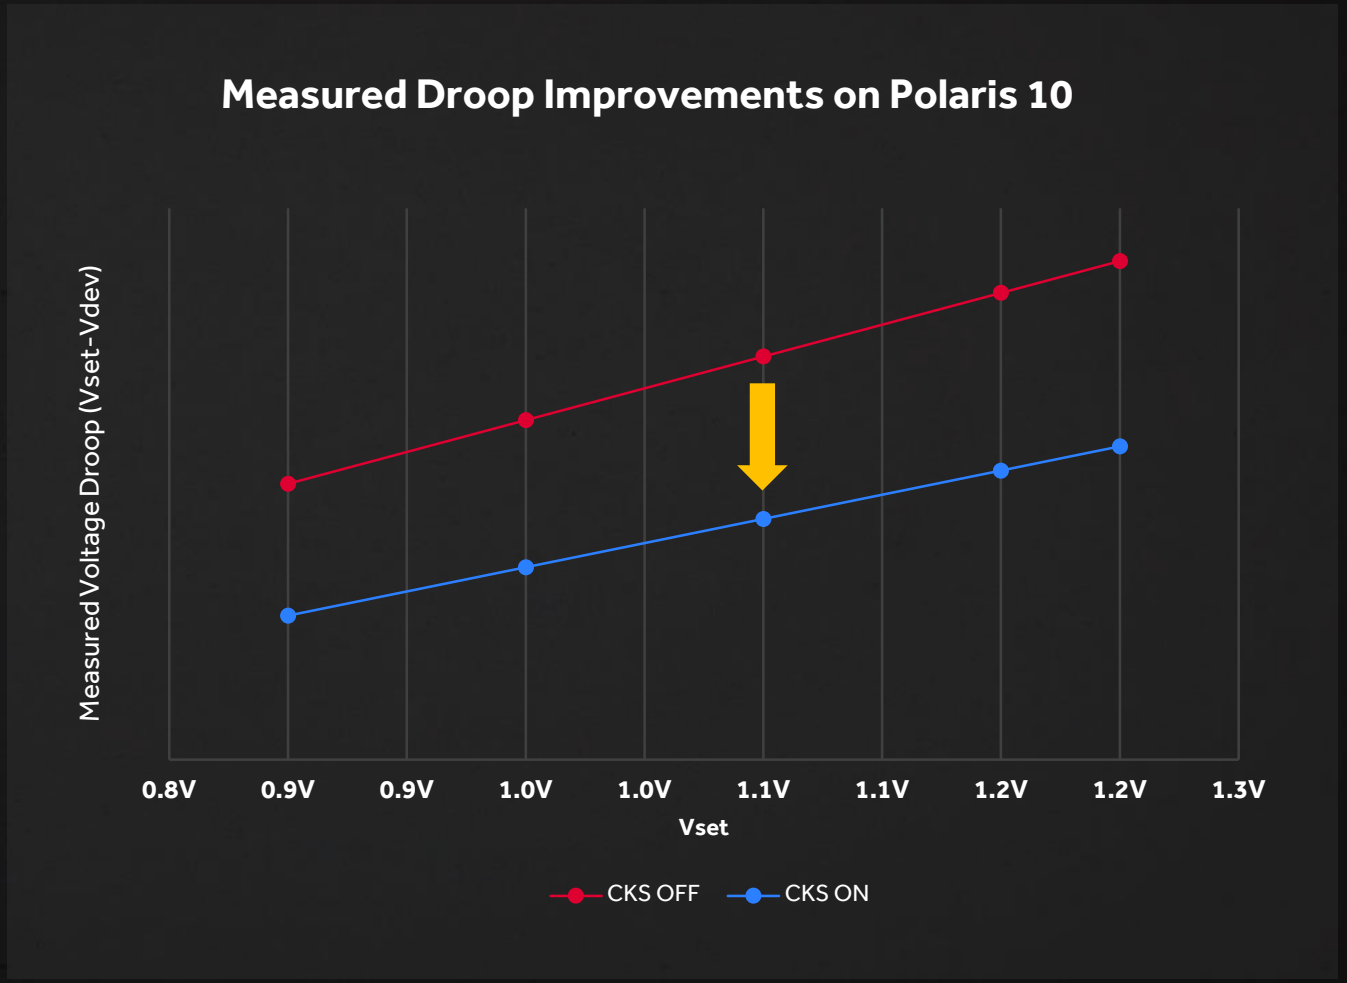

Illustrating adaptive clocking

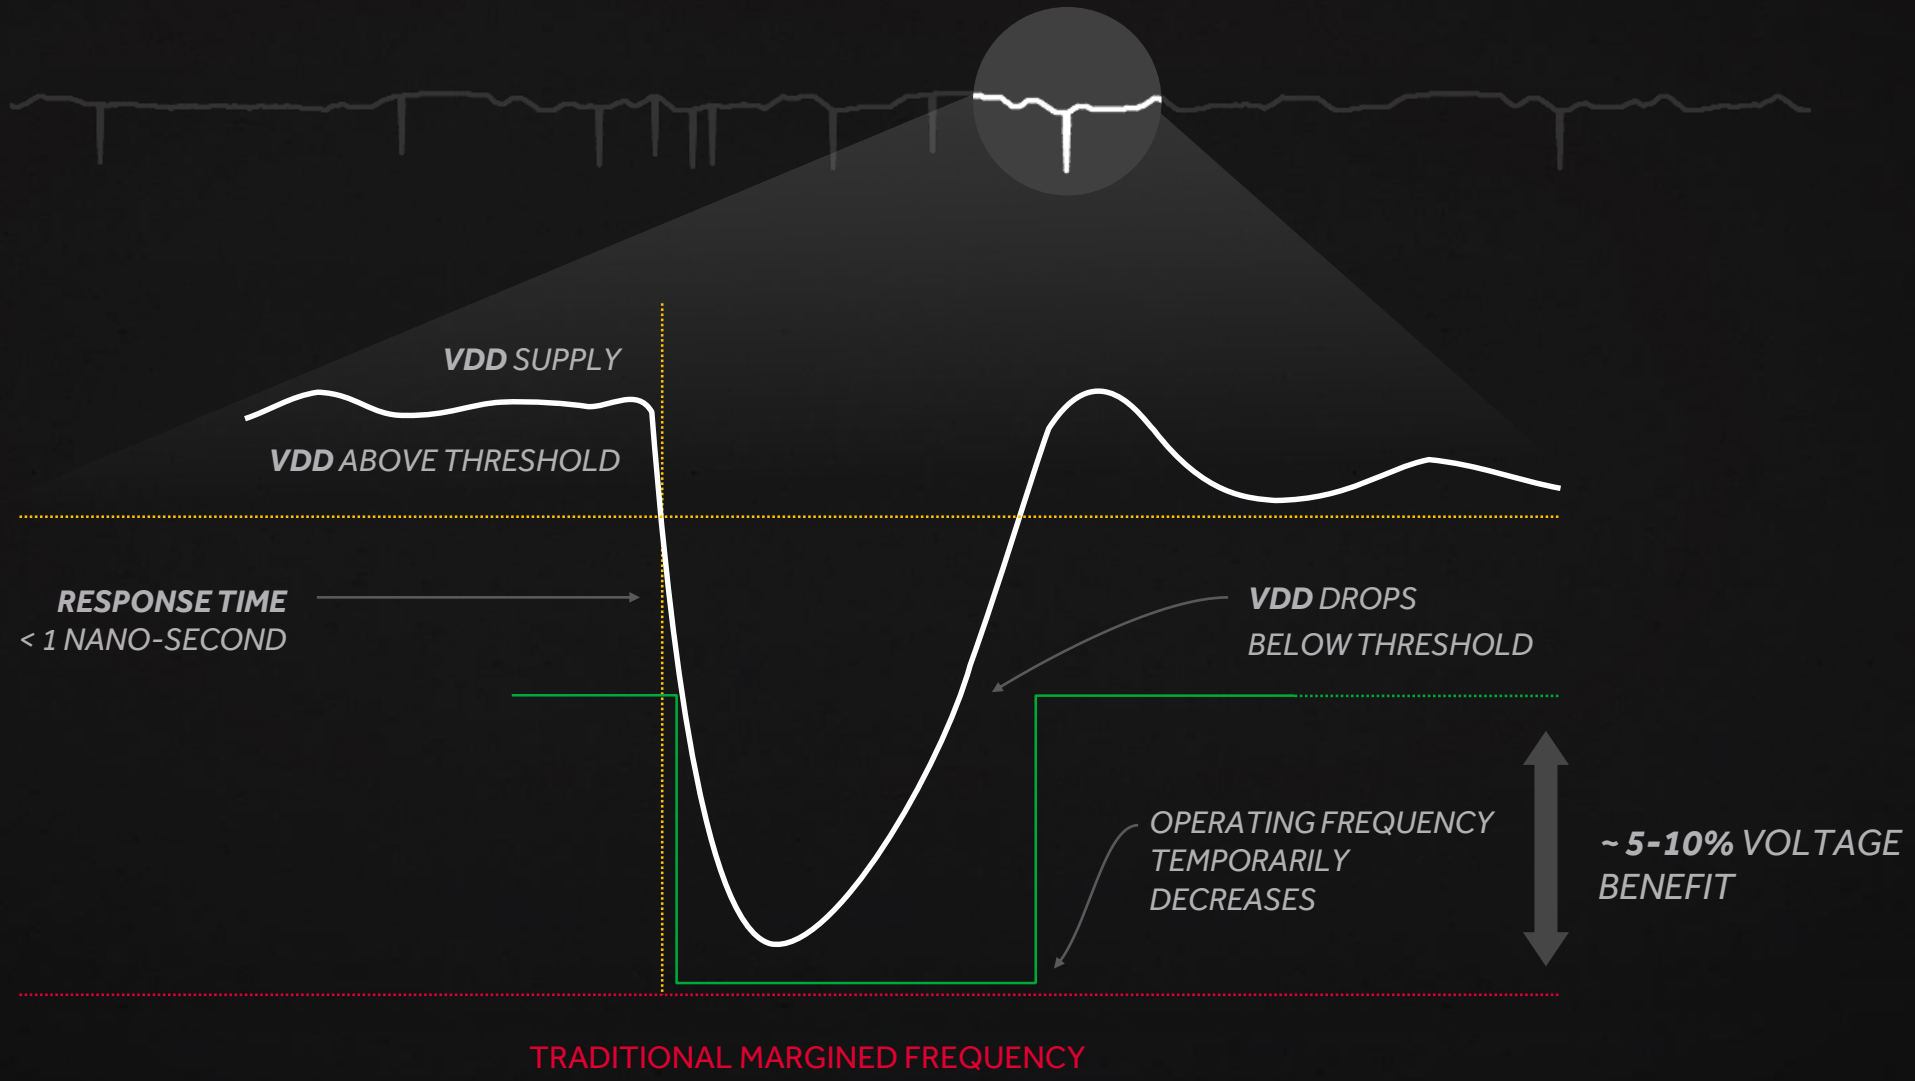

Adaptive voltage & frequency scaling (AVFS)

- Reliably extract the true silicon speed capability of the GPU
  - Including: part-to-part processing, temperature and power delivery
  - How: add a frequency sensor to existing power and temp sensors
- Enables accurate setting of the optimal operating point
  - Optimal energy efficiency across the entire ASIC process, voltage, and temperature

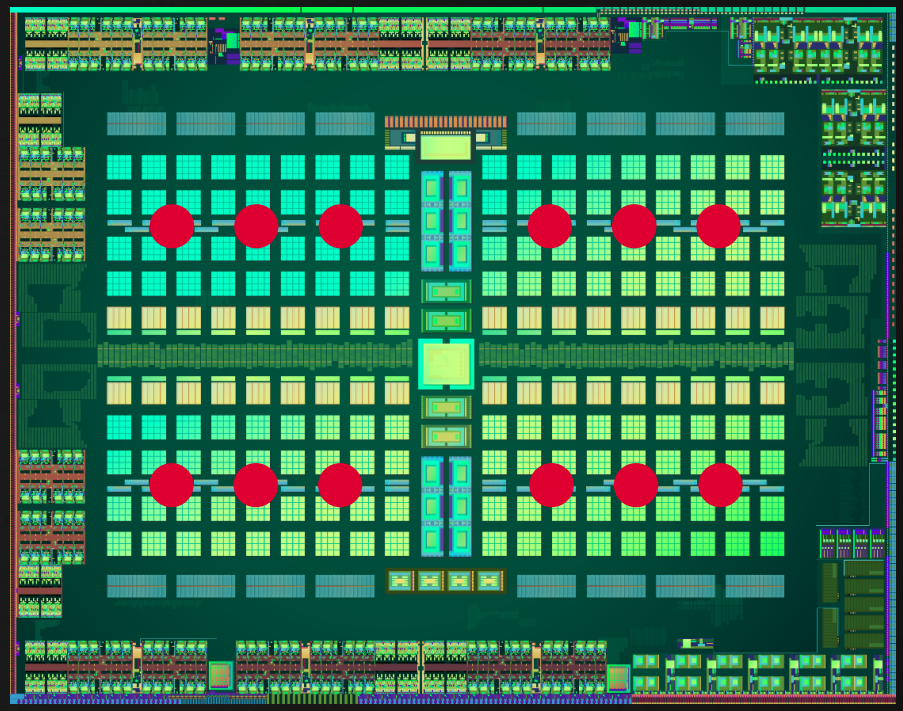

Individual speed and voltage sensors distributed around the graphics core

Illustrating AVFS on the Polaris architecture

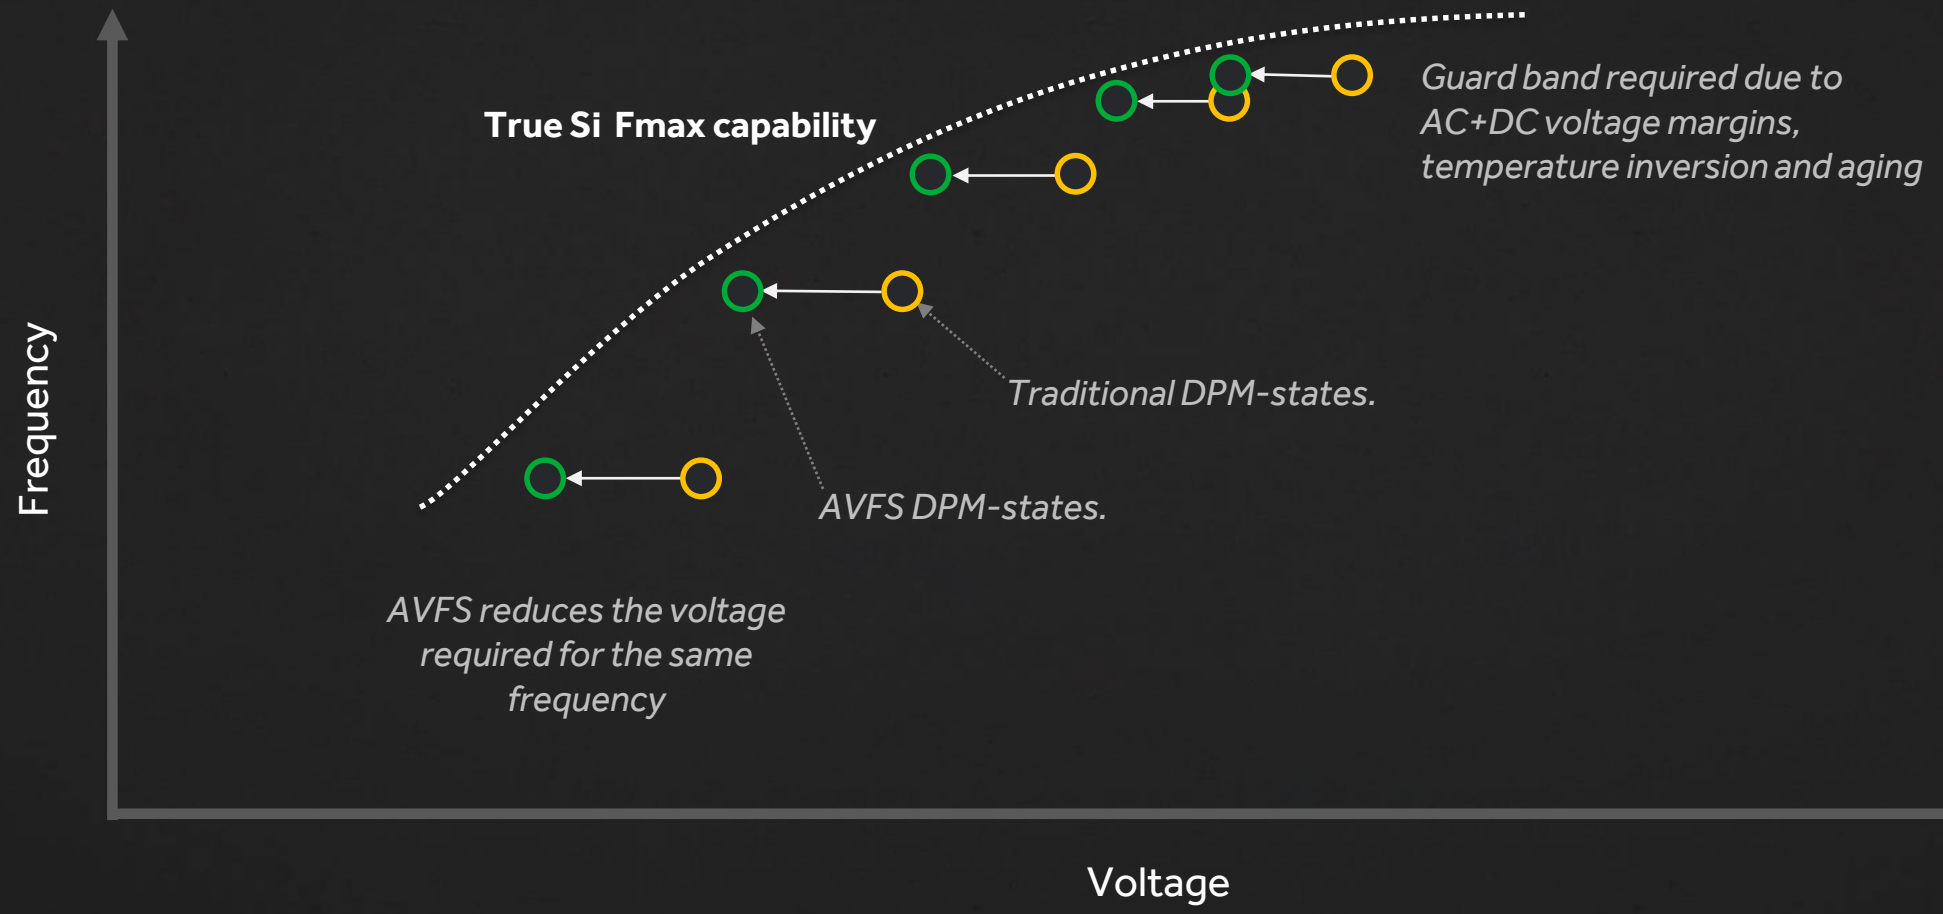

* Diagram for illustrative purposes.

Boot time power supply calibration (BTC)

- Run voltage analysis code on a tester when the part is tested and binned
  - Log the voltage as seen by the integrated power supply monitors
- Run the same code on PC boot and observe voltage
  - Dial board vregs to deliver the same voltage as observed on the tester
- Helps eliminate waste power usually spent accommodating board/system variation

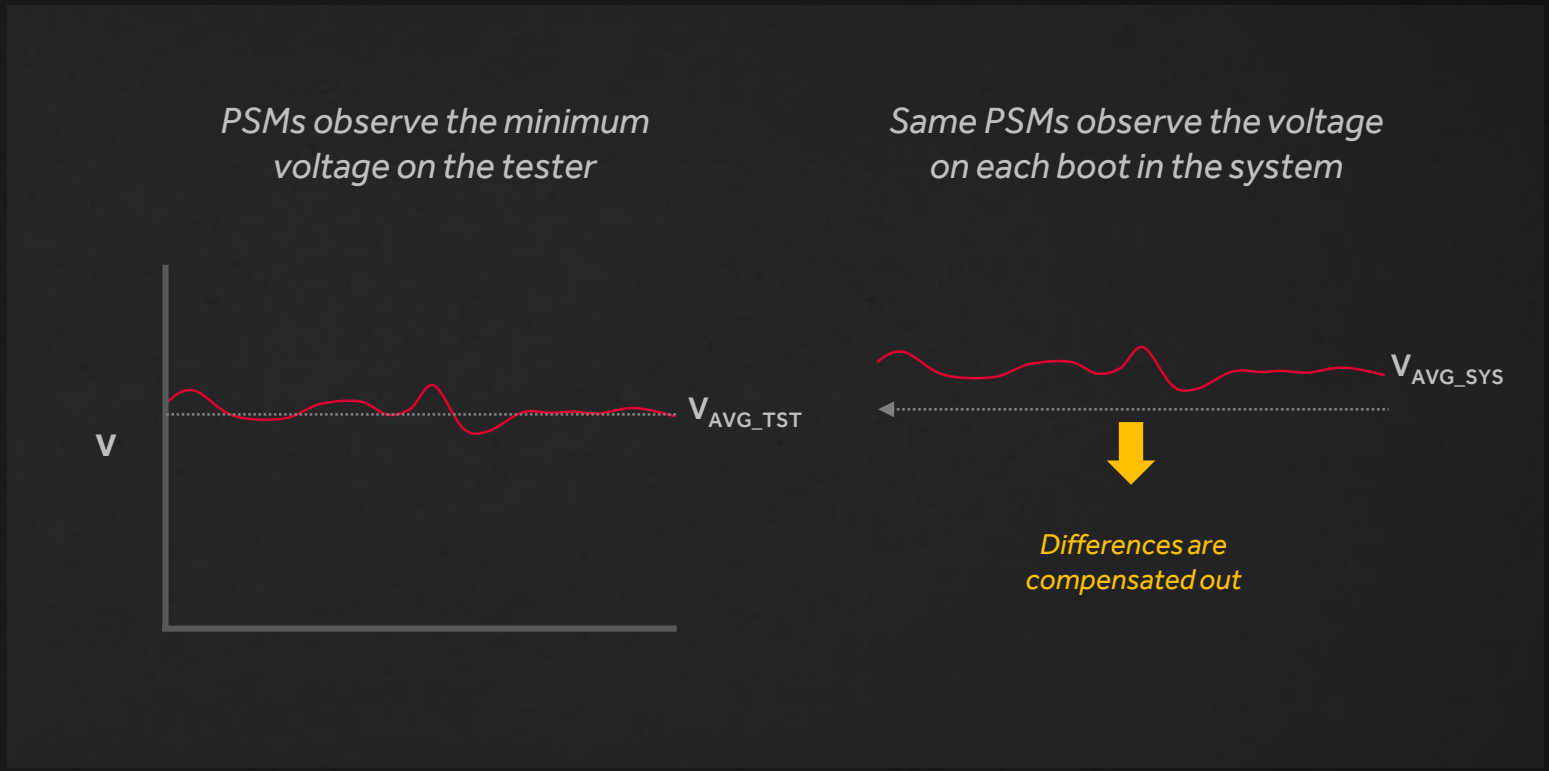

* Diagram for illustrative purposes.

Adaptive aging compensation

- GPUs require 2-3% clockspeed margin to accommodate transistor aging
- Other aspects of the system also exhibit aging (e.g. lower voltage from system)
- We want our parts to self-calibrate and adapt to changes over time – good or bad
  - More robust operation over time
  - Out-of-box performance is improved

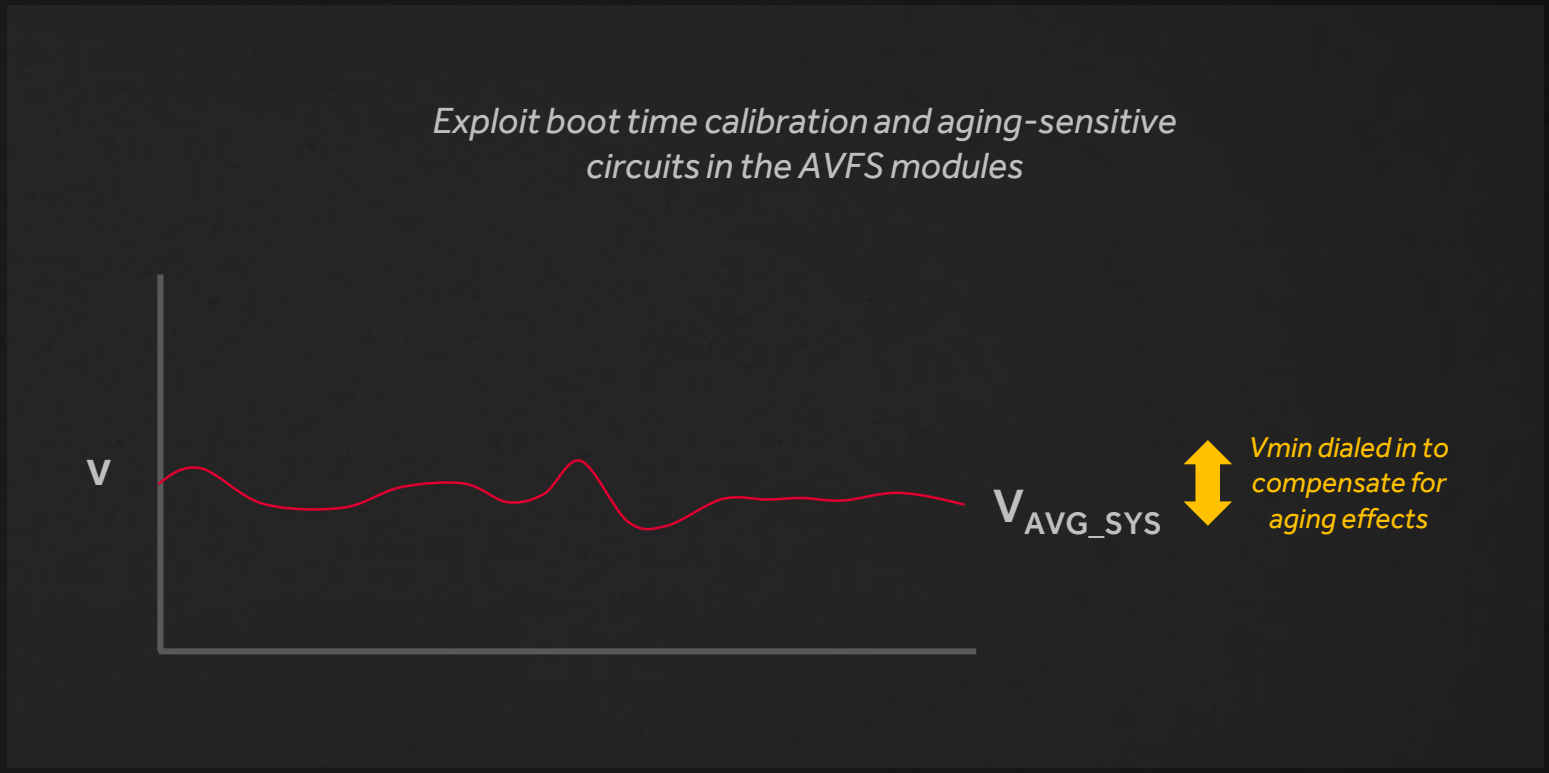

* Diagram for illustrative purposes.

Multi-bit flip-flop (MBFF)

- Specially-designed banks (MBFFs) that share circuitry
  - Critical clock sequence storage elements
- There are approx. 21 million flip-flops in Polaris 10
  - They utilize about 15% of ASIC TDP
- Save significant internal and sClk power
  - 4-5% TDP savings

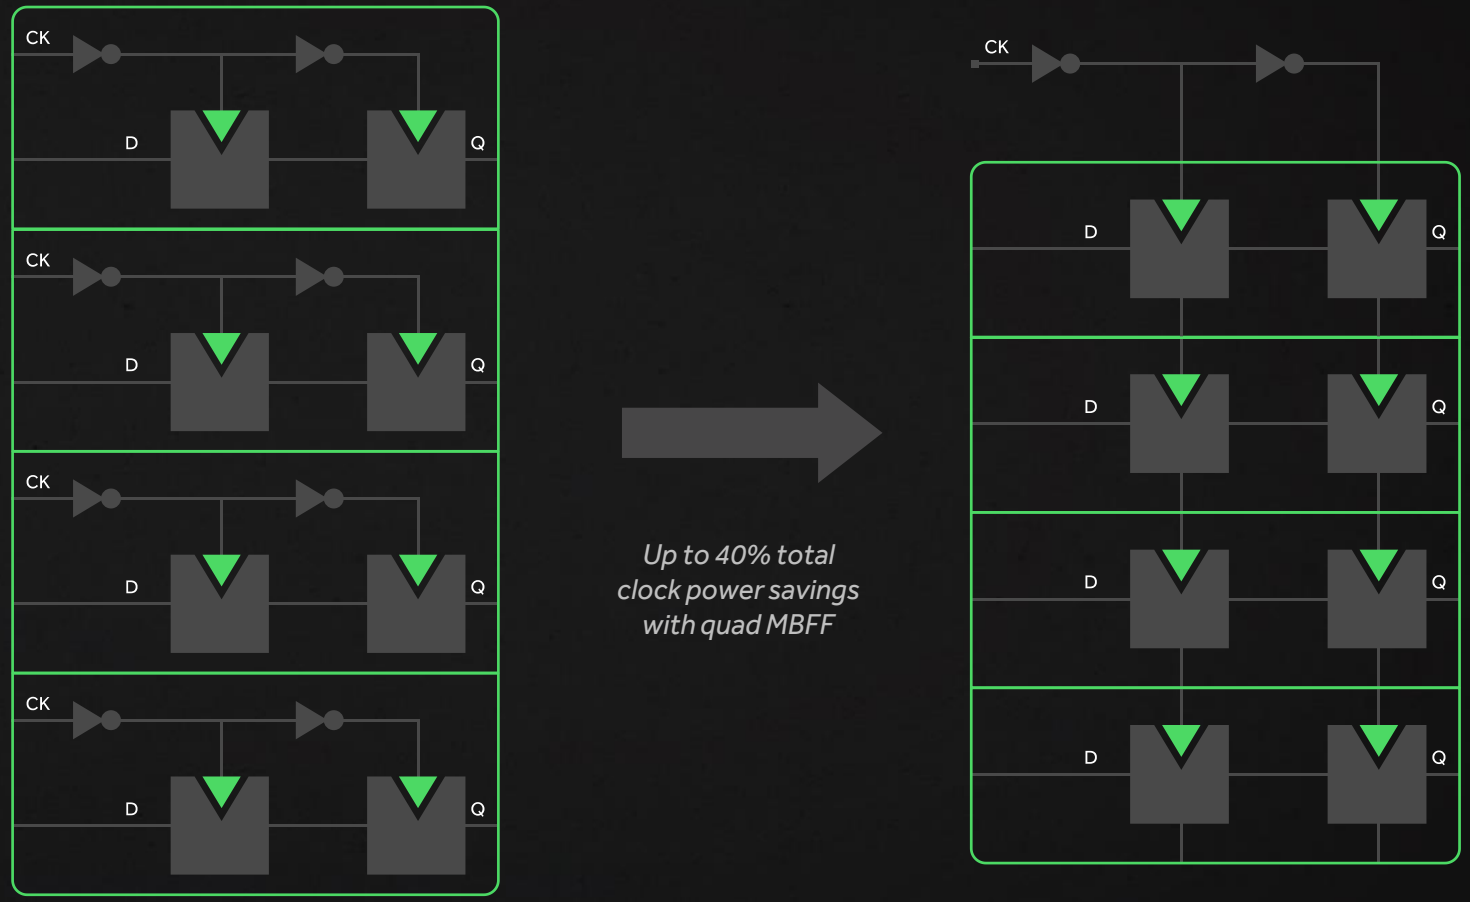

Energy efficient memory interface

- Reducing memory bandwidth dependence:
  - Double L2 cache
  - Improve L2 utilization and efficiency
  - Improve LDCC algorithms and usage
- Simultaneously reduce energy / bit
  - FinFETs help
  - Optimized interface design
  - Power gating, clock gating

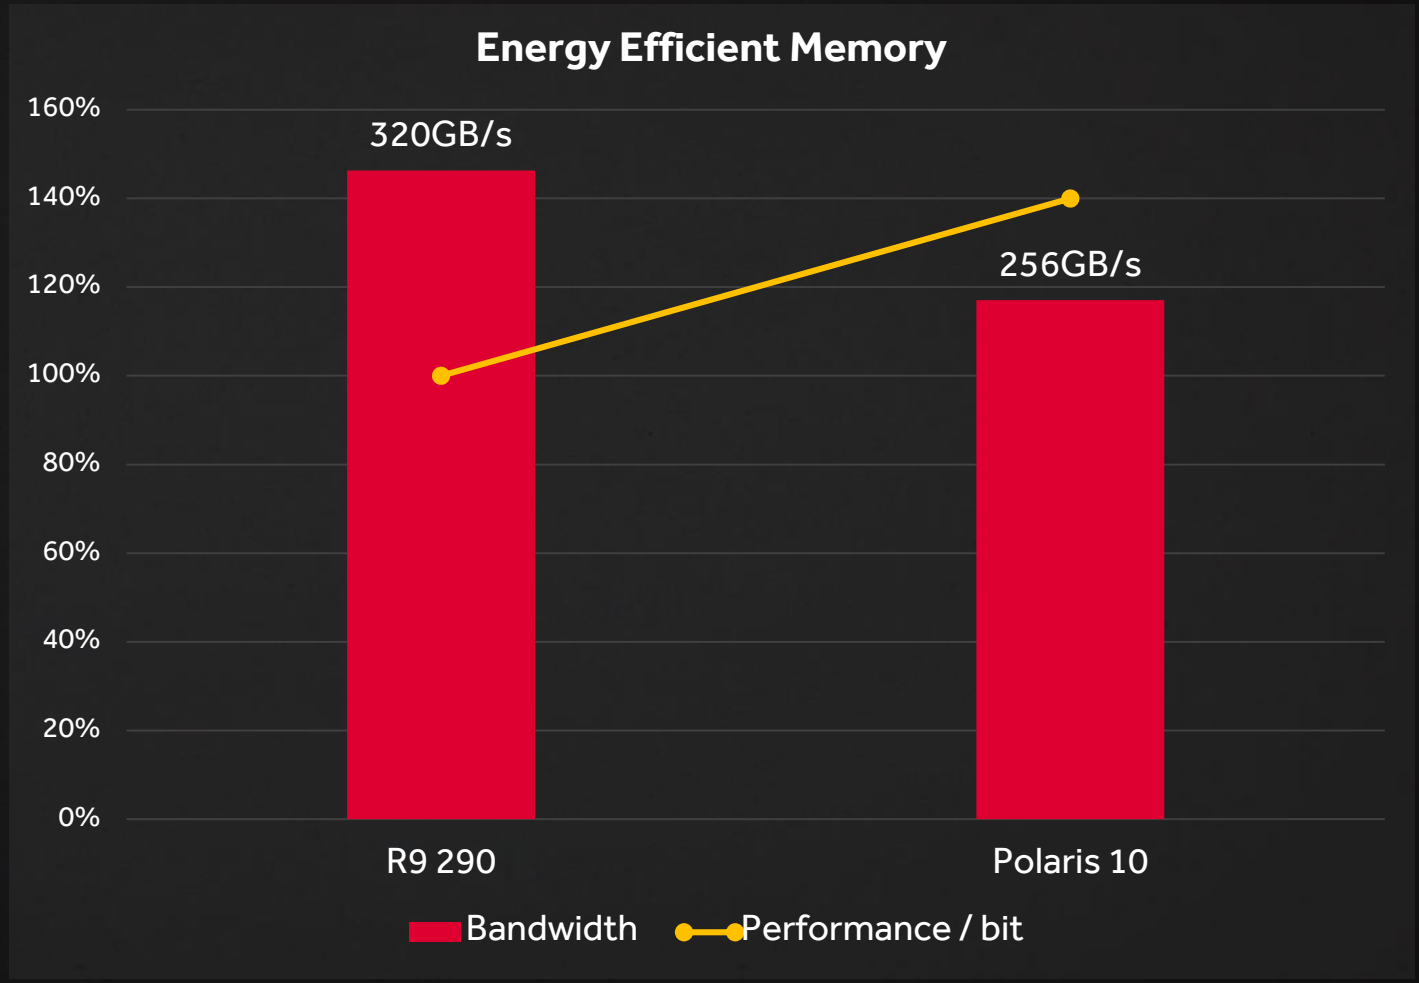

Energy efficient memory interface

- Reducing memory bandwidth dependence:
  - Double L2 cache
  - Improve L2 utilization and efficiency
  - Improve LDCC algorithms and usage
- Simultaneously reduce energy / bit
  - FinFETs help
  - Optimized interface design
  - Power gating, clock gating

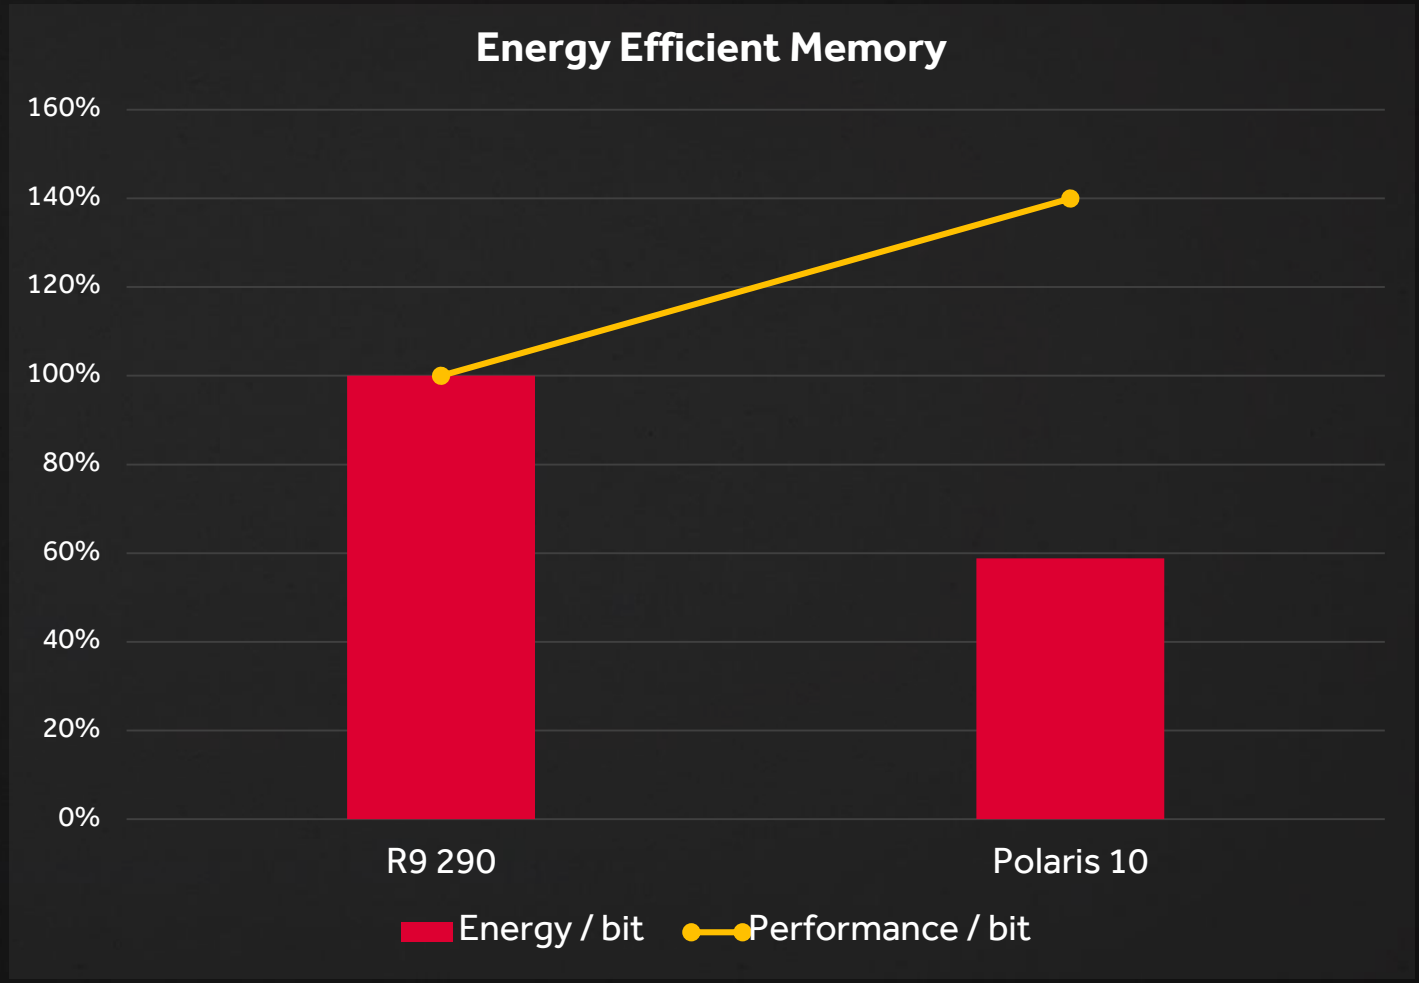

Energy efficient memory interface

- Lower Energy per bit
- Better performance per bit

- Outcome:
  - Memory bandwidth needs reduced
  - Total memory interface power reduced by up to 58%[5]

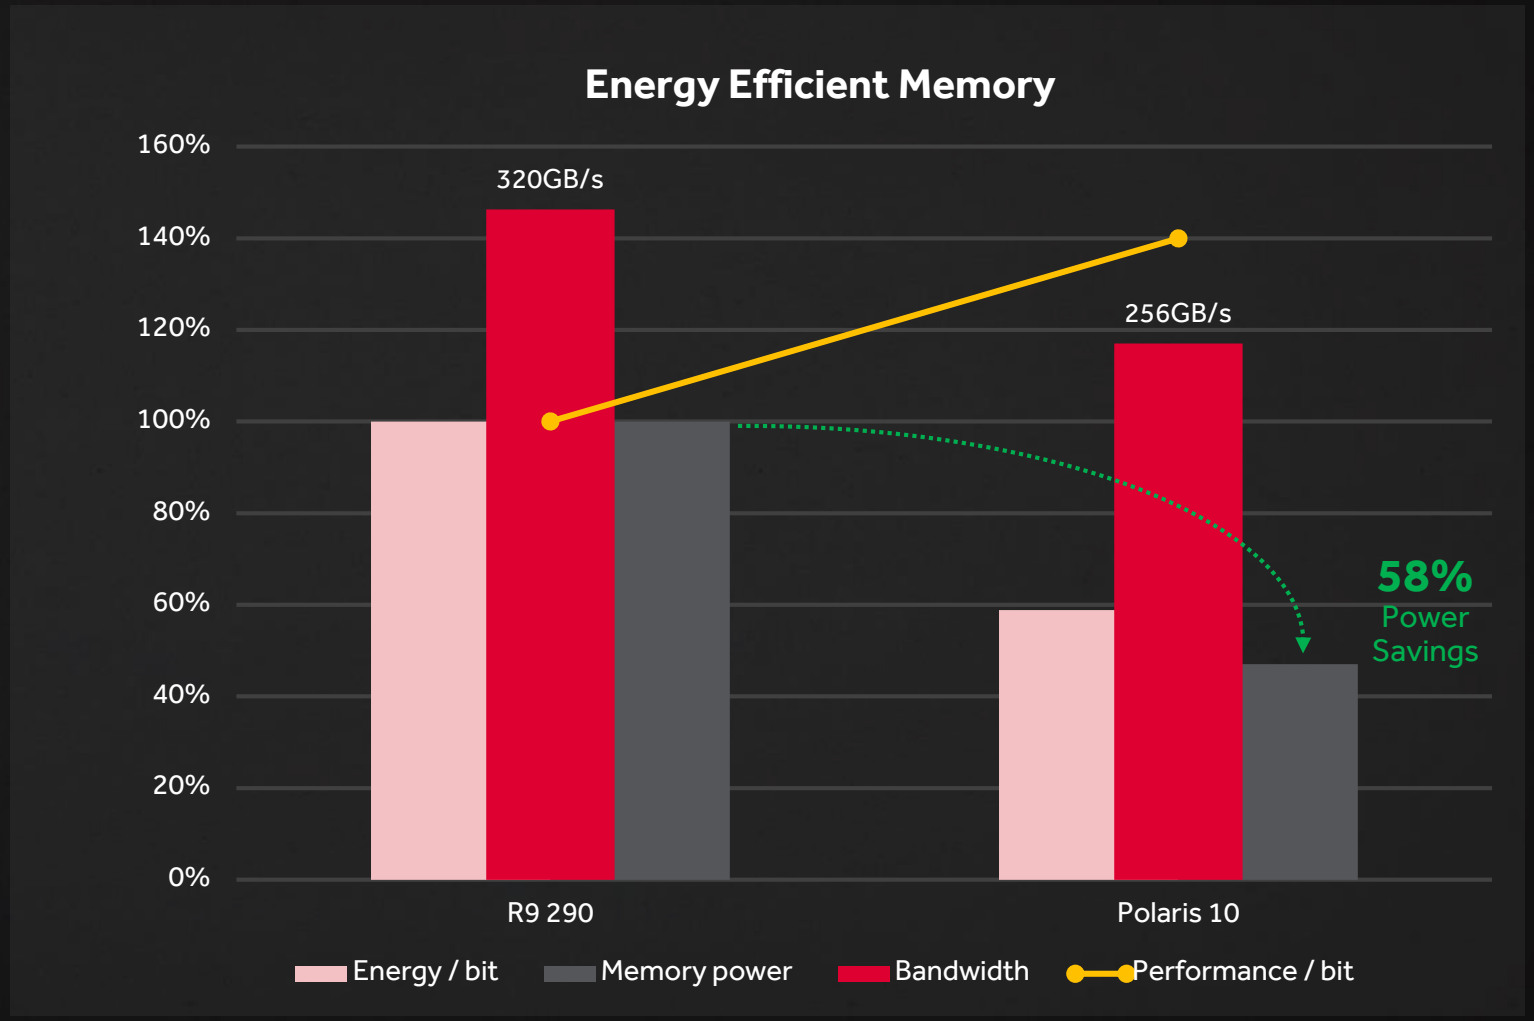

14nm FinFET optimized by AMD

UP TO

1.7x

Performance / Watt

With FinFET 14[6]

UP TO

2.8x

Performance / Watt

With AMD technologies[2]

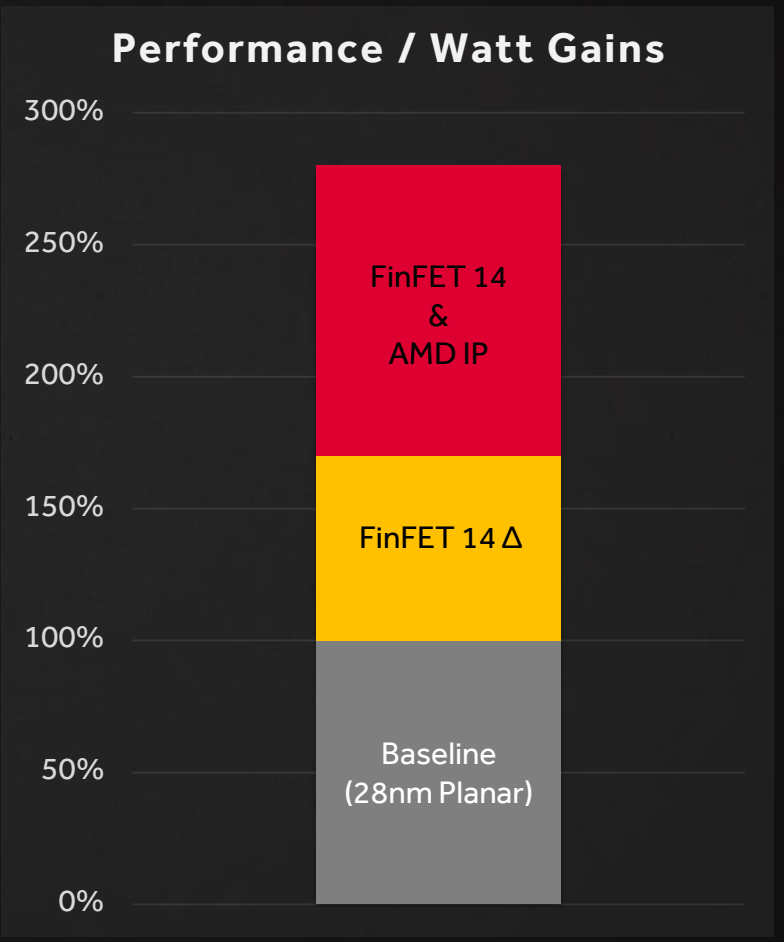

Under embargo until June 29, 2016 at 9 a.m. EST.

AMD | RADEON™

Radeon™ RX 460 graphics

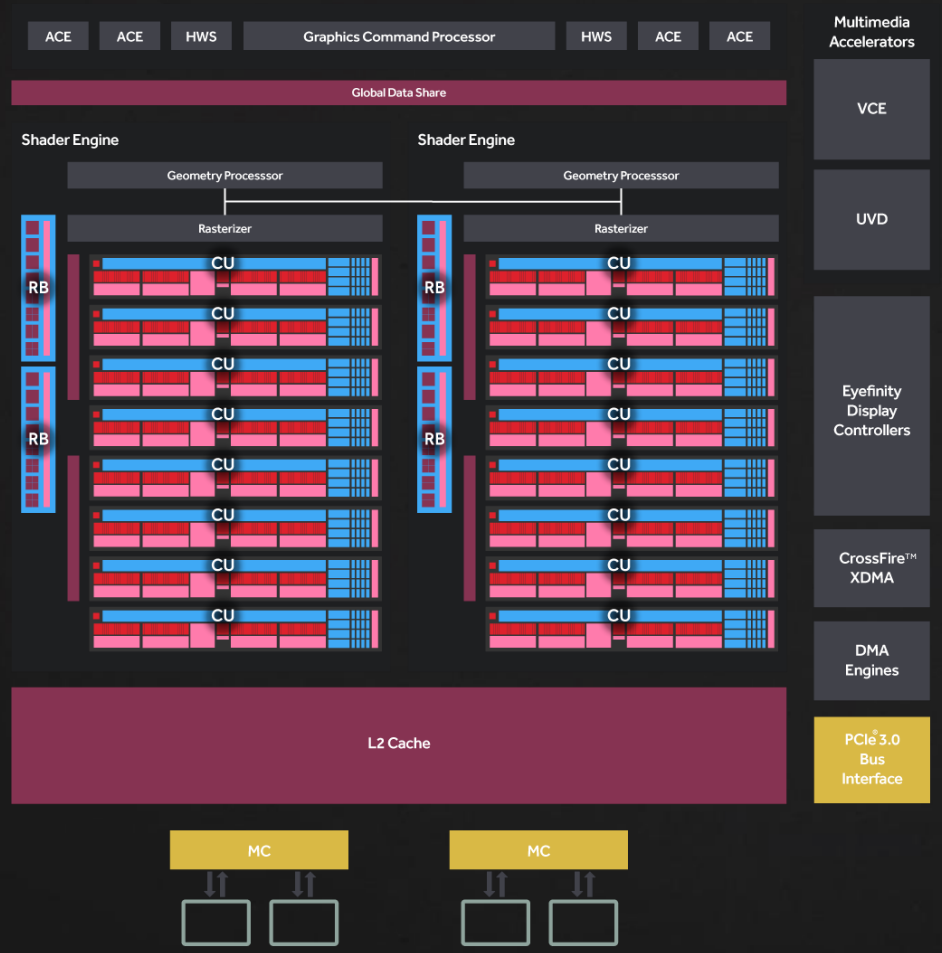

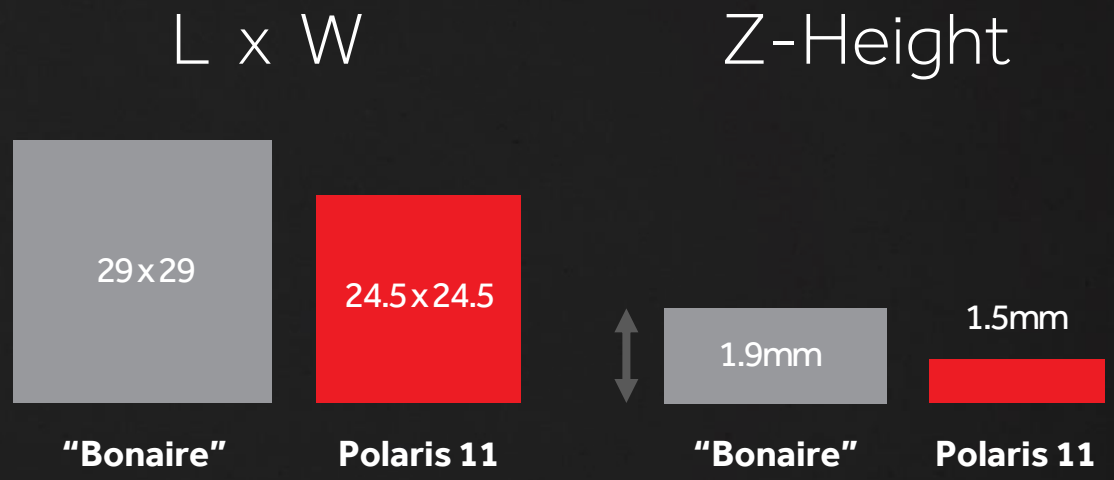

Summarizing the Polaris Architecture

- Our best graphics core yet
  - Up to +15% generational performance per compute unit[1]
- Future-proof display support[7]
  - Ready for 10-bit and 12-bit HDR monitors
  - Ready for next-gen high density and high refresh gaming monitors
- State-of-the-art multimedia support
  - CPU-free streaming and recording of your favorite games
  - Ready for HDR streaming video services
  - New 2-pass video encoding
- Cool and quiet gameplay
  - FinFET 14 optimized by AMD technologies
  - Up to 2.8X generational perf/W uplift[2]

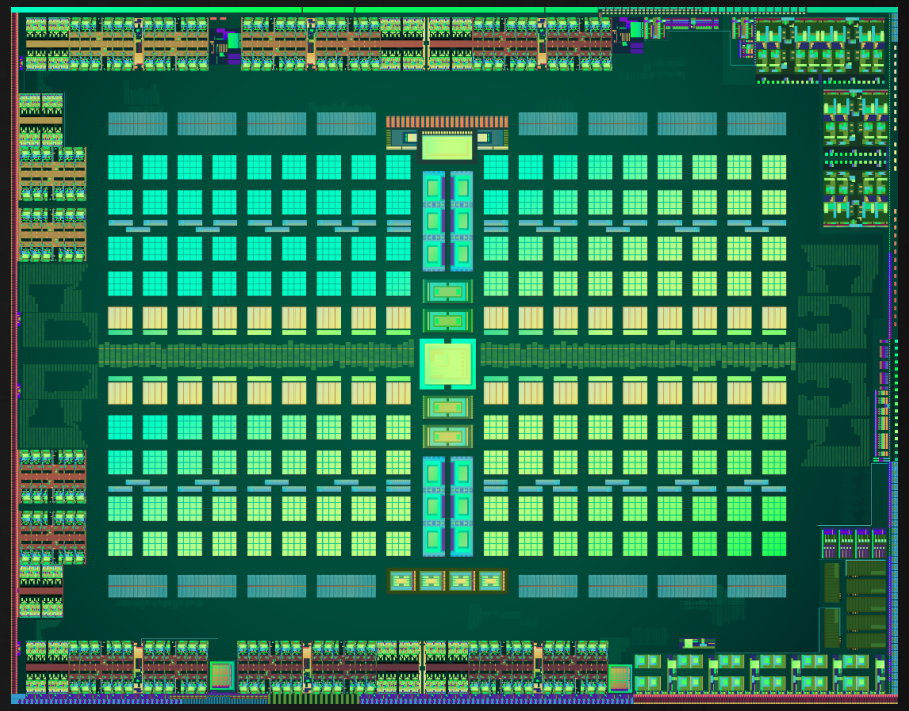

Polaris 10 GPU

Footnotes

1. Testing conducted by AMD performance labs as of May 18, 2016 on the Radeon RX 480 and Radeon R9 290 on a test system comprising Intel Core i7-5960X, 16GB DDR4-2666, Gigabyte X99-UD4, Windows 10 x64 (build 10586), Radeon Software Crimson Edition 16.5.2using Ashes of Singularity, GTA V, Project Cars, Witcher, and Assassin's Creed Syndicate, All games tested at 1440p. Radeon RX 480 graphics (150W TGP/36 CU) vs. Radeon R9 290 graphics (275W TGP/40 CU) scores as follows: Ashes of the Singularity (44.19 FPS vs 46 FPS); GTA V (66.23 FPS vs. 66.44 FPS); Project Cars (48.99 FPS vs. 45.99 FPS); Witcher 3 (50.78 FPS vs. 50.13 FPS); Assassin's Creed Syndicate (50.51 FPS vs. 45.78 FPS). Average FPS of above game scores: 52.14 (Radeon RX 480) vs. 50.06 (Radeon R9 290). Discrete AMD Radeon™ GPUs and AMD FirePro™ GPUs based on the Graphics Core Next architecture consist of multiple discrete execution engines known as a Compute Unit ("CU"). Each CU contains 64 shaders ("Stream Processors") working in unison (GD-78). CU efficiency formula = average FPS/# of CUs. Test results are not average and may vary. RX-4
2. Testing conducted by AMD Performance Labs as of May 10, 2016 on the AMD Radeon™ RX 470 (110w) and AMD Radeon™ R9 270X (180w), on a test system comprising i7 5960X @ 3.0 GHz 16GB memory, AMD Radeon Software driver 16.20 and Windows 10. Using 3DMark Fire Strike preset 1080p the scores were 9090 and 5787 respectively. Using Ashes of the Singularity 1080P High, the scores were 46 fps and 28.1 fps respectively. Using Hitman 1080p High, the scores were 60 fps and 27.6 fps respectively. Using Overwatch 1080p Max settings, the scores were 121 fps and 76 fps respectively. Using Performance/Board power, the resulting average across the 4 different titles was a perf per watt of 2.8X vs the Radeon R9 270X. Test results are not average and may vary. RX-6
3. Based on AMD internal small prim filter test. Primitive assembly rates with prim filter ON vs. OFF: 18 tri/px (3.947 vs. 1.255), 32 tri/px (3.901 vs. 1.773), 50 tri/px (3.760 vs. 1.402), 72 tri/px (3.303 vs. 1.187), 98 tri/px (3.928 vs. 1.171), 128 tri/px (3.870 vs. 1.111). System configuration: Radeon™ RX 480, Core i7-6700K, 16GB DDR4-2666, Windows 10 x64, Radeon™ Software 16.5.2.
4. Based on AMD internal memory bandwidth test. Radeon™ R9 290X: 263GB/s peak memory bandwidth. Radeon™ R9 Fury: 333 peak GB/s without DCC vs. 387 peak GB/s with DCC. Radeon™ RX 480: 186 peak GB/s without DCC vs. 251 peak GB/s with DCC. System configuration: Core i7-6700K, 16GB DDR4-2666, Windows 10 x64, Radeon™ Software 16.5.2.
5. Based on measurements of total memory interface power in watts conducted by the AMD performance labs as of 5/21/2016. System configuration: Radeon™ R9 290 vs. Radeon™ RX 480, Core i7-5960X, Gigabyte GA-X99-UD7, 16GB DDR4-2666, Windows 10 x64, Radeon™ Software 16.5.2.
6. Based on AMD internal data generated in AMD performance labs as of May 2016, measurements of capacitance, voltage frequency, leakage and power data show up to 1.7x performance/watt on 14nm vs 28nm FINFET technology. Final performance/watt results on AMD products using 14nm FinFET technology may vary and will depend on various factors including but not limited to clock speed, voltage, and various AMD proprietary technologies. RX-17
7. Statement of "future-proof" refers to support of current and upcoming technology standards including 14nm FinFET process technology, DirectX®12 and Vulkan™ API support, new display technology, and experiences such as VR. "Future-proof" statement is not meant to serve as a warranty or indicate that users will never have to upgrade their graphics technology again. Support of current and upcoming technology standards described above has the potential to reduce frequency of graphics upgrades for some users.
8. One compute unit on an AMD Radeon™ GPU is equal to 64 stream processors.
9. Hardware-accelerated VP9 decode support is planned for enablement in a future Radeon™ Software driver release.

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

HDMI™
HIGH-DEFINITION MULTIMEDIA INTERFACE