


intel.

**TECH.**  
tour.TW

# Xe2 and Lunar Lake GPU Deep Dive

**TAP**

Intel Fellow





# Xe2





**Xe**

First time scaling the engine

2 years of software effort

Efficiency



**Xe2**


Higher utilization

Improved work distribution

Less SW overhead

# Xe2


Improving IP  
Performance  
Efficiency



Relative IP performance (Higher is better)  
Normalized to configuration and clock frequency

See backup for workloads and configurations. Results may vary.

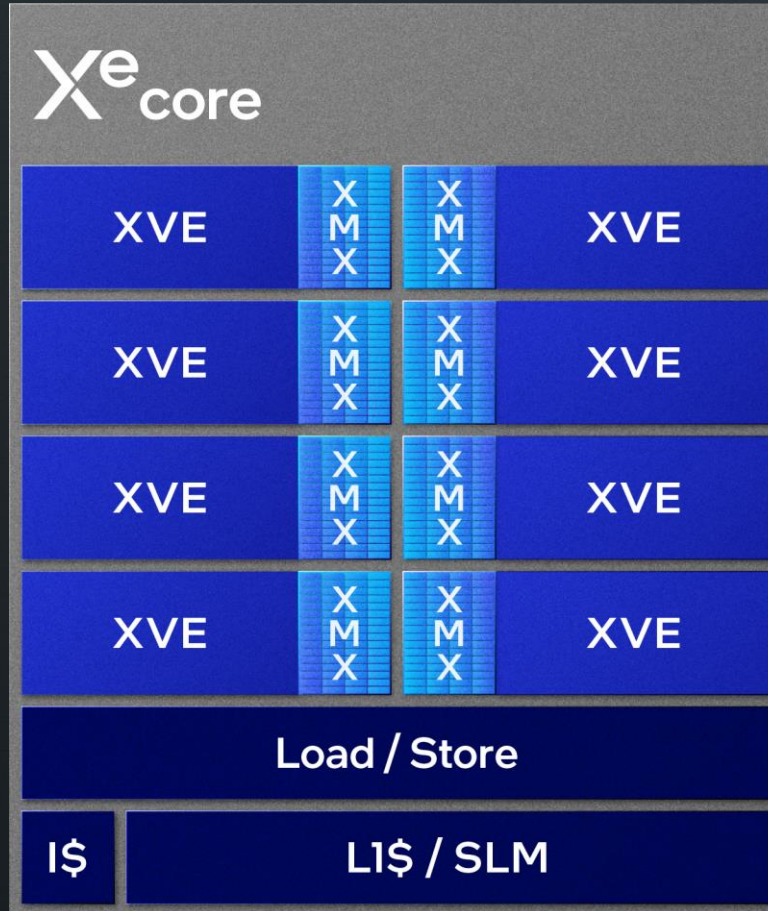
# Xe2 Architecture Scalability



2<sup>nd</sup> GEN

# Xe core

Compute resources repartitioned in native SIMD16 engines for increased efficiency



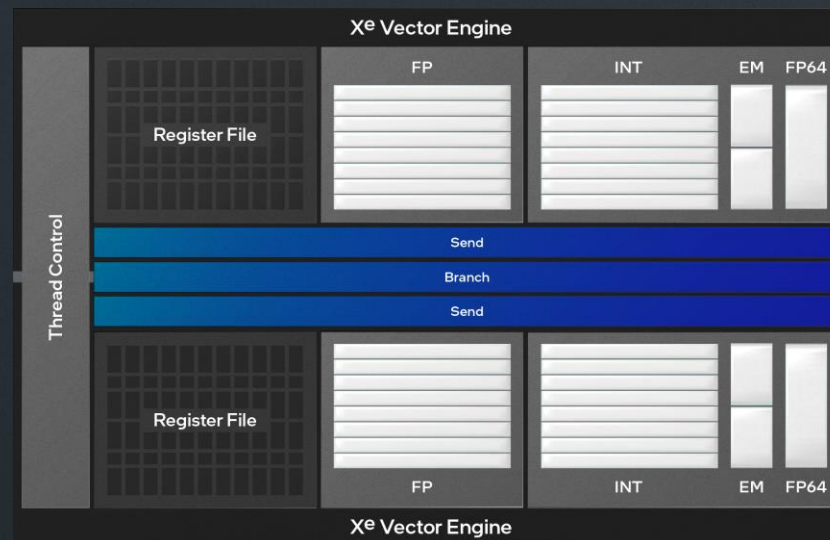
8 512-bit Vector Engines

8 2048-bit XMX Engines

64b atomic ops support

192KB Shared L1\$ / SLM

# New Vector Engine



## SIMD16 native ALUs

Support for SIMD16 and SIMD32 ops

## Xe Matrix Extensions

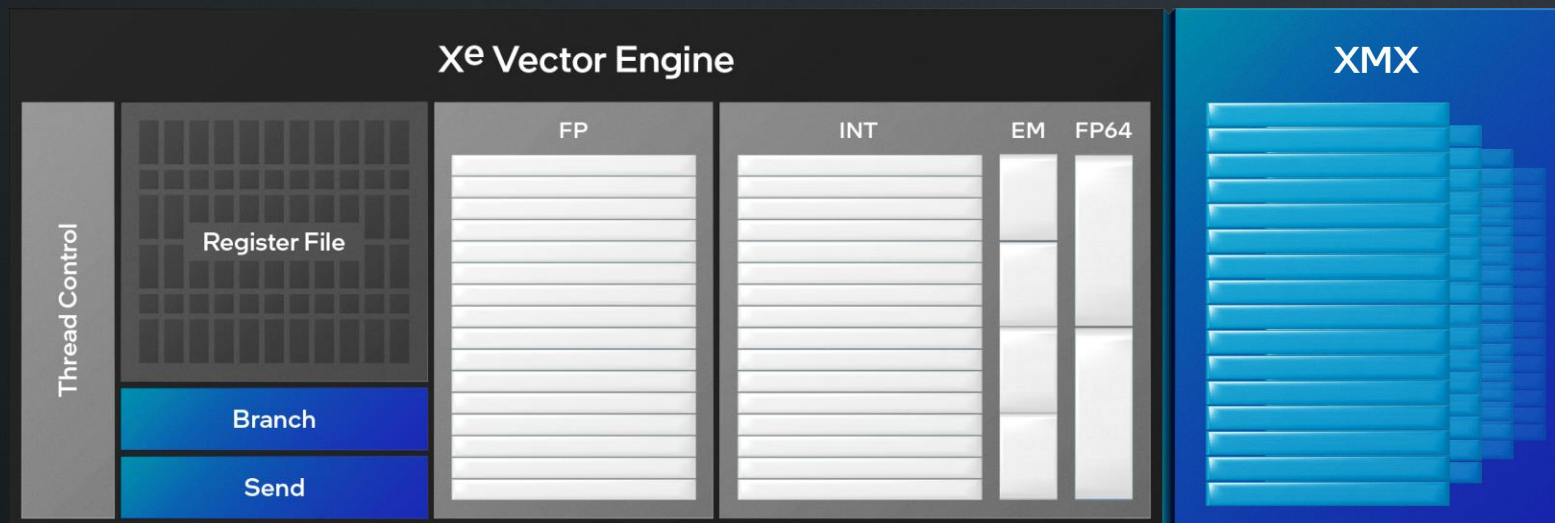
Support for INT2, INT4, INT8, FP16, BF16

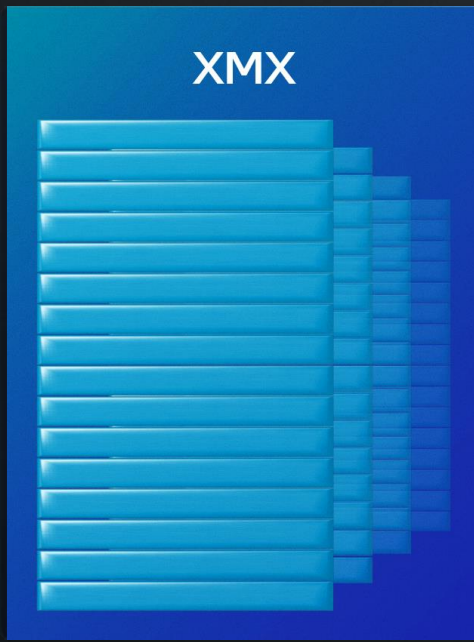
## Extended Math and FP64

Transcendentals: SIN, COS, LOG, EXP...

## 3-way co-issue


FP + INT/EM + XMX






**INT8**

16x4x4x2



**FP16**

16x2x4x2



New

# X<sup>e</sup> Matrix Extension Engines

**FP16 2048** OPS/clock

**INT8 4096** OPS/clock



# Key Peak Metrics for 2<sup>nd</sup> Gen X<sup>e</sup>-core

	Number of XVE	SIMD width	MAC/lane	Depth	Ops/MAC	Ops/clock
FP32	8	16	1	1	2	256
FP16	8	16	2	1	2	512
DP4a INT8	8	16	4	1	2	1024
XMV FP16 / BF16	8	16	2	4	2	2048
XMV INT8	8	16	4	4	2	4096
XMV INT4 / INT2	8	16	8	4	2	8192

# New Render Slice

Deep micro and macro **analysis** of all graphics acceleration functions

**Optimized** to reduce latency, remove stalls and improve HW/SW handshake



# New Render Slice

## Command Front End

### Render Slice



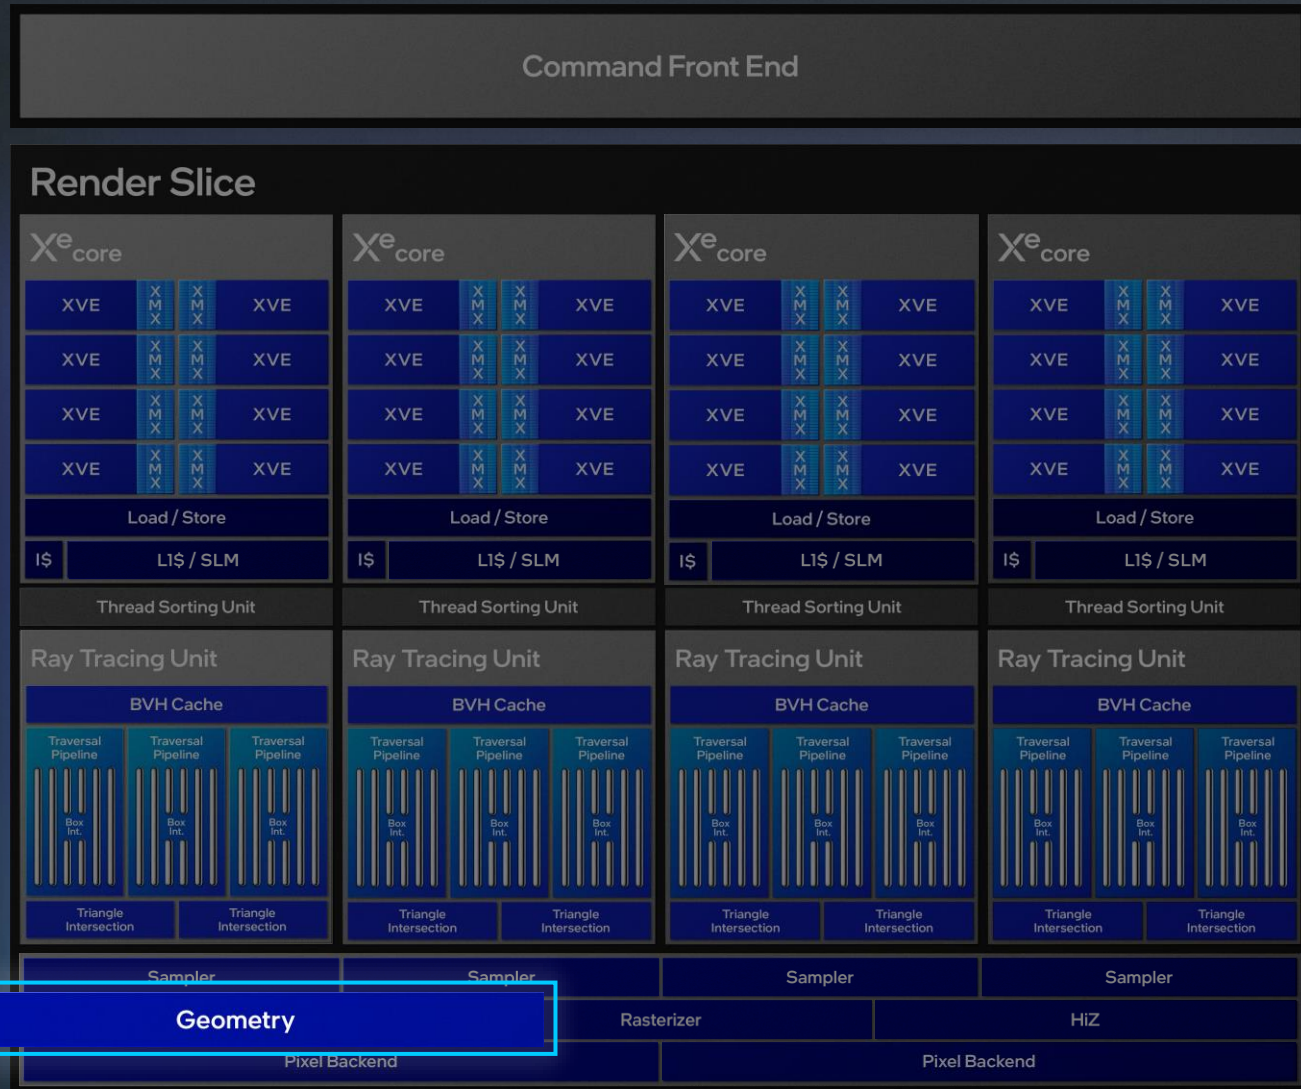
**Execute indirect**

Natively supported

# New Render Slice

3x vertex fetch  
throughput

3x mesh shading  
performance  
with vertex re-use



# New Render Slice

Out of order sampling  
with compressed textures

2x throughput  
for sampling without filtering

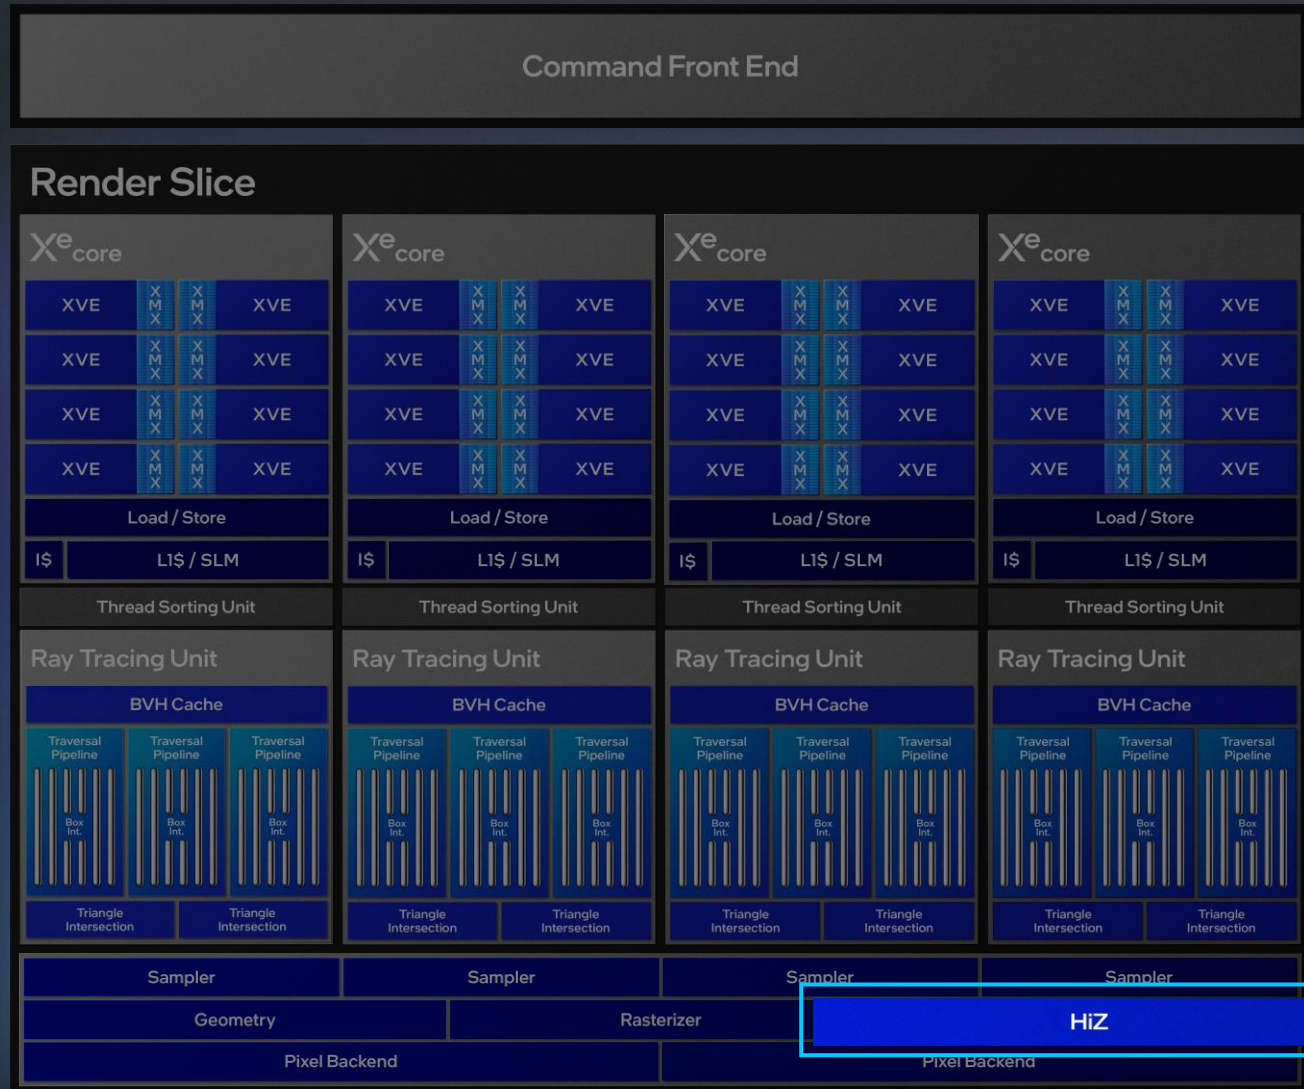
Programmable offsets



# New Render Slice

1.5x HiZ / Z / Stencil  
cache

Early HiZ culling  
of small primitives



# New Render Slice

**2x blending throughput**

for high granularity passes

**1.33x increase**

in pixel color cache

**Render target pre-fetch**

to L2\$



# New Render Slice

New 8:N compression

Fast clear  
for sub-resources





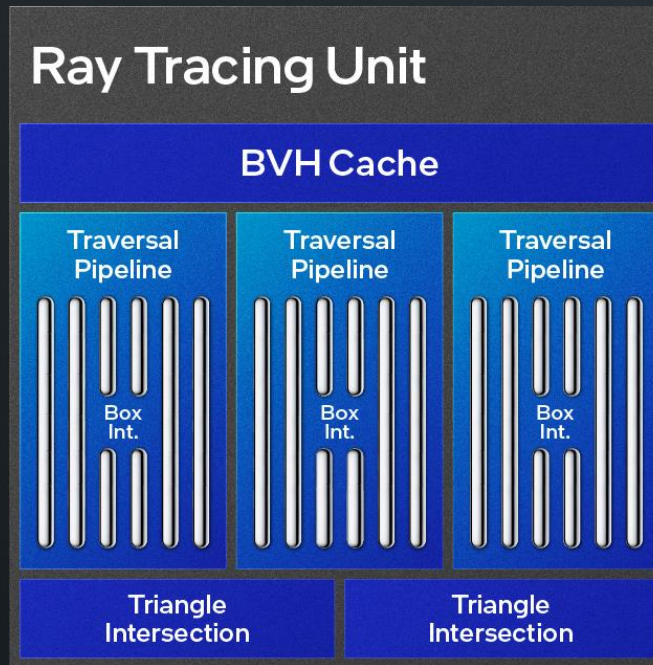


New  
**XeRTU**

3 Traversal pipelines

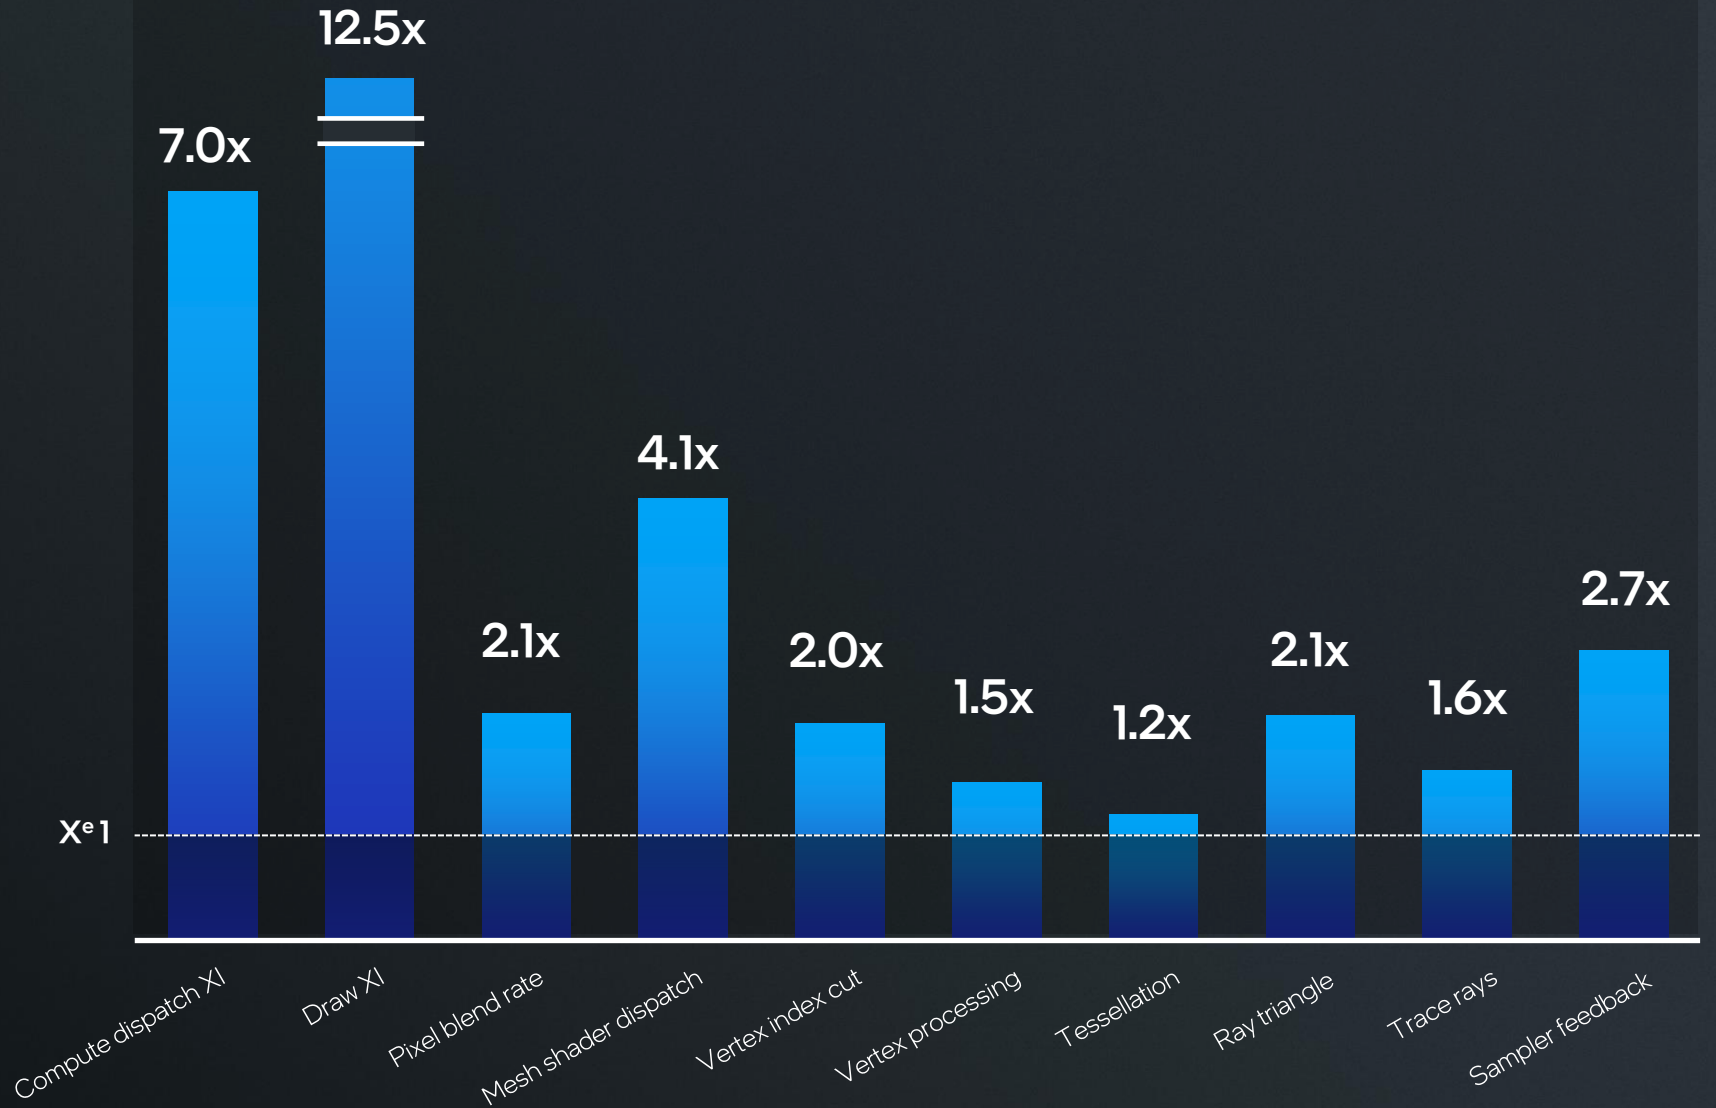
18 Box intersections

2 Triangle intersections



# Xe2

Improving IP  
Performance  
Efficiency





Relative IP performance (Higher is better)  
Normalized to configuration and clock frequency

See backup for workloads and configurations. Results may vary.


# Xe2

Intel GPU Architecture

## 2<sup>nd</sup> Gen Xe-core




Enhanced  
vector  
engines



Deeper  
caches

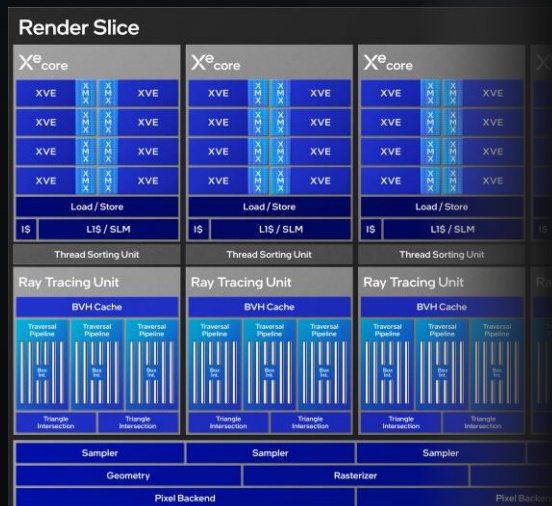
Larger  
ray tracing units




Native hardware support for  
execute indirect commands

Command Front End

Performance  
& efficiency  
optimized  
front to end



New  
XM  
engines



Implementing

# Xe2

Meteor Lake

Lunar Lake



1


Deliver on key mobile experiences

2

Step function in efficiency


3

Support for latest industry standards




# Lunar Lake

## Graphics overview



# Lunar Lake


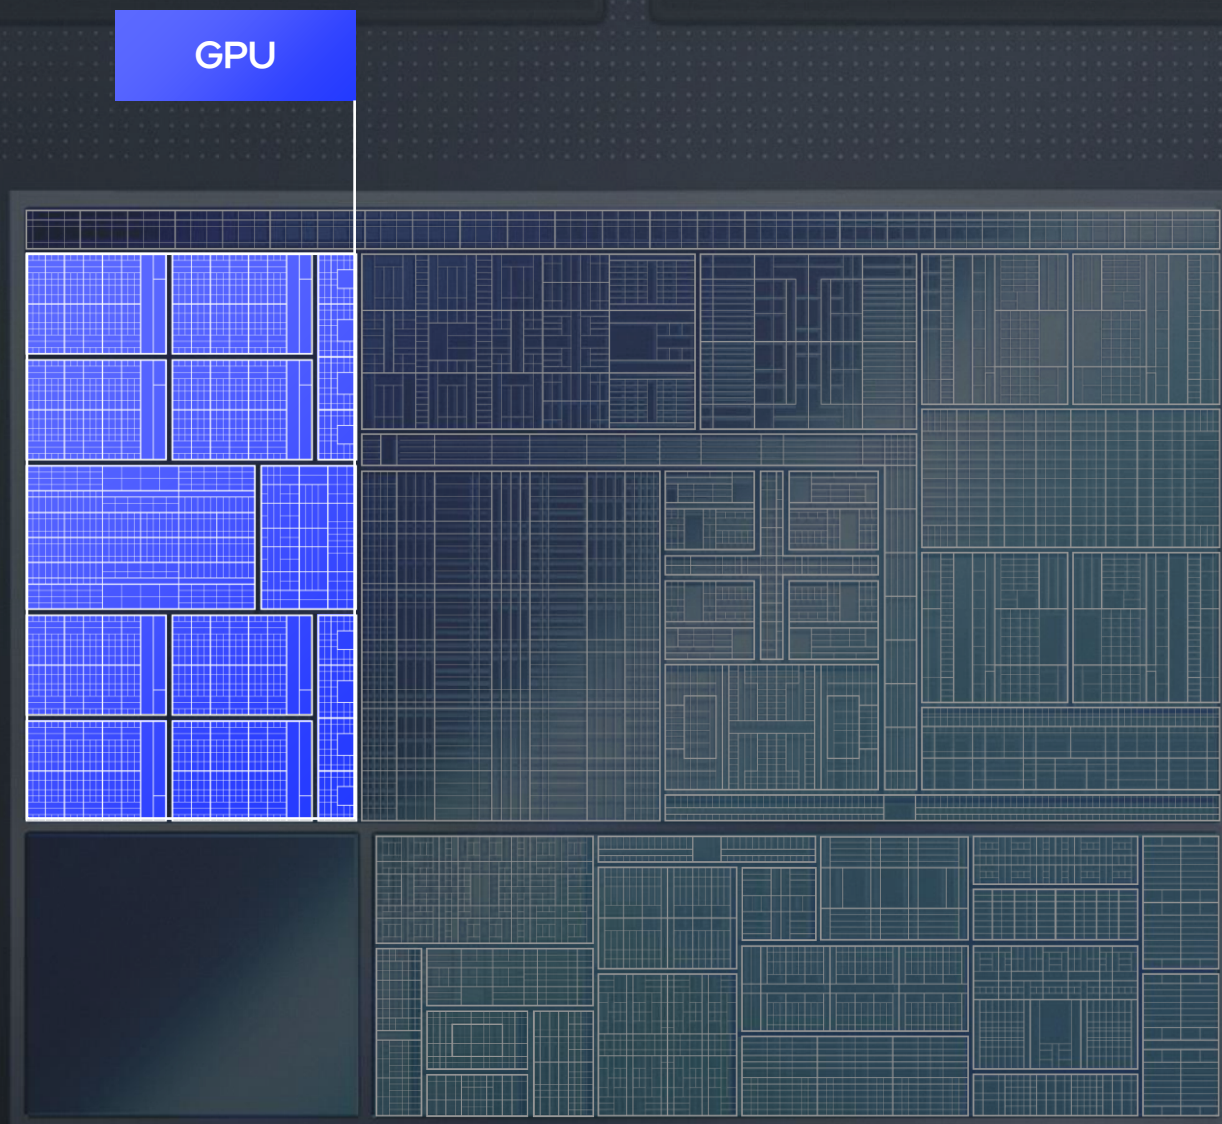
## Graphics overview



Lunar Lake

# Xe2 GPU

Optimized for  
performance efficiency



# Xe2

## Lunar Lake Configuration

8 Xe-cores

64 vector engines

2 geometry pipelines

8 samplers

4 pixel backends

8 ray-tracing units

8 MB L2\$





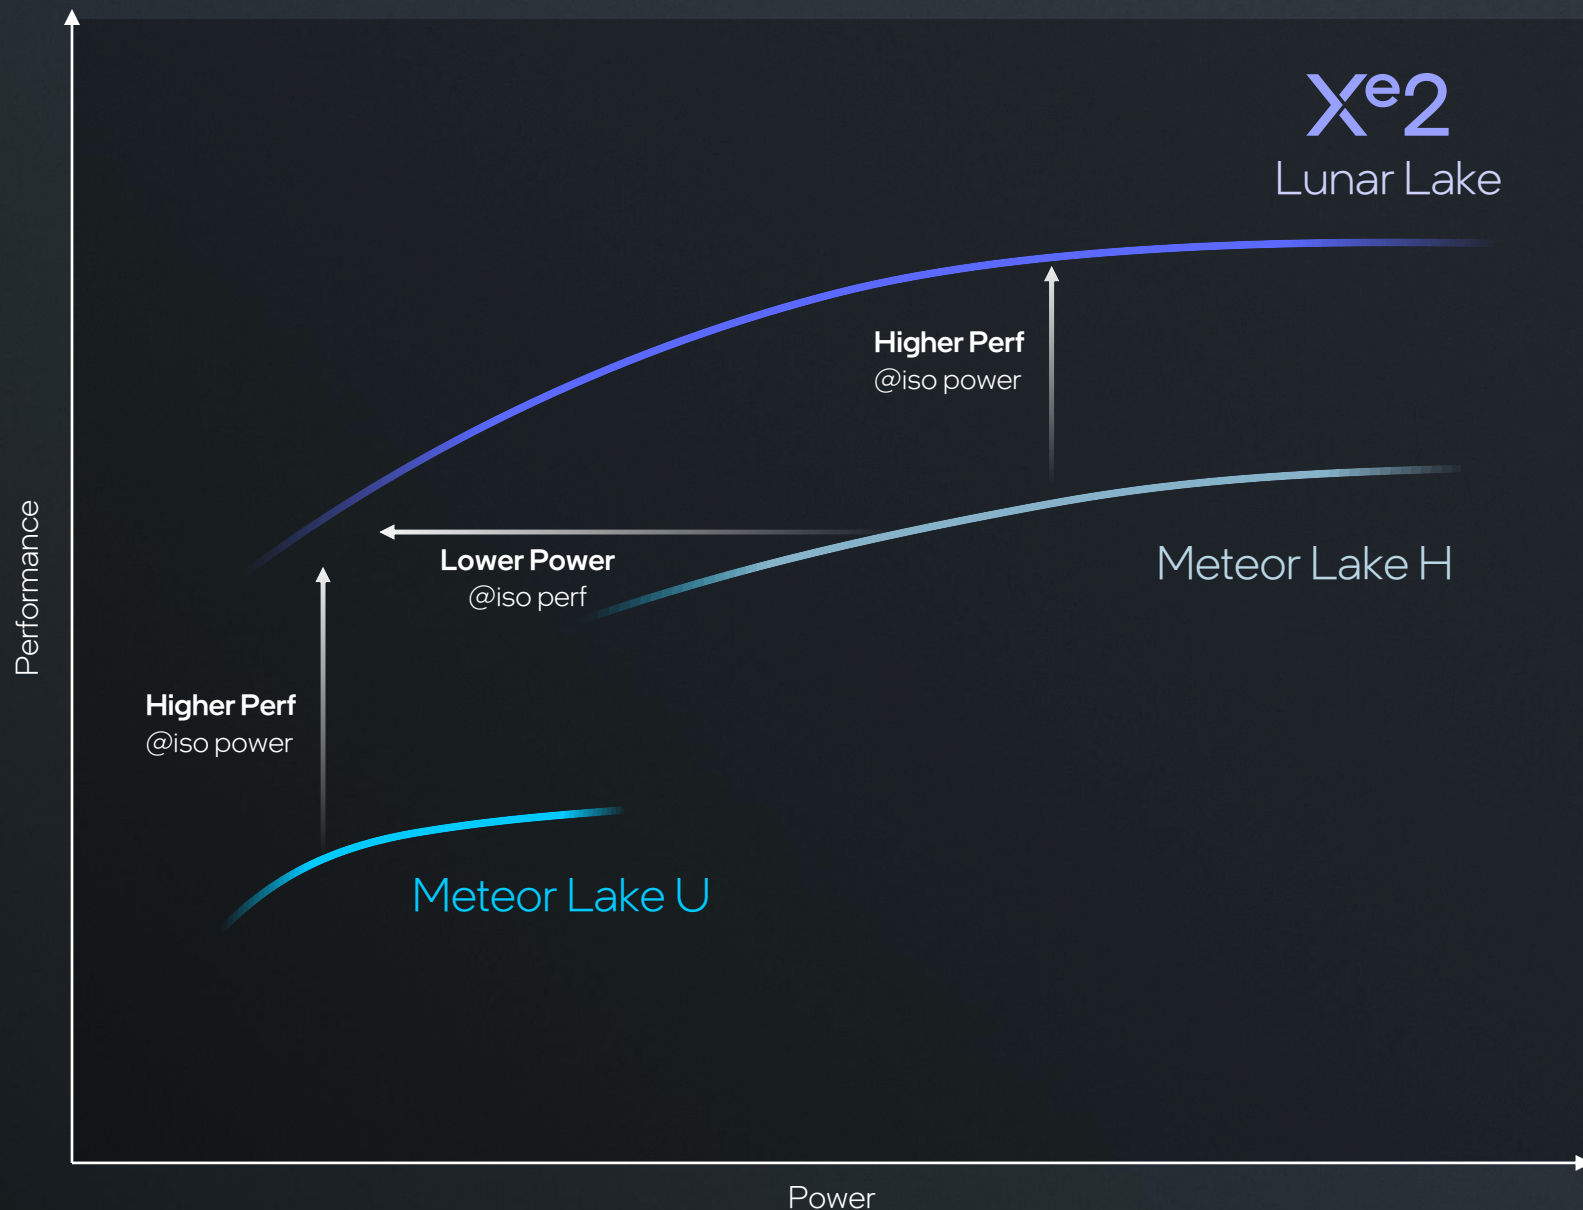
Lunar Lake

# Xe2 GPU

## Performance

~1.5x

vs. previous gen



Xe2  
Lunar Lake

Higher Perf  
@iso power

Lower Power  
@iso perf

Meteor Lake H

Higher Perf  
@iso power

Meteor Lake U

Power

Lunar Lake


# GPU AI Engine

**XMV**

X<sup>e</sup> Matrix  
Extensions



**67**

peak INT8  
TOPs



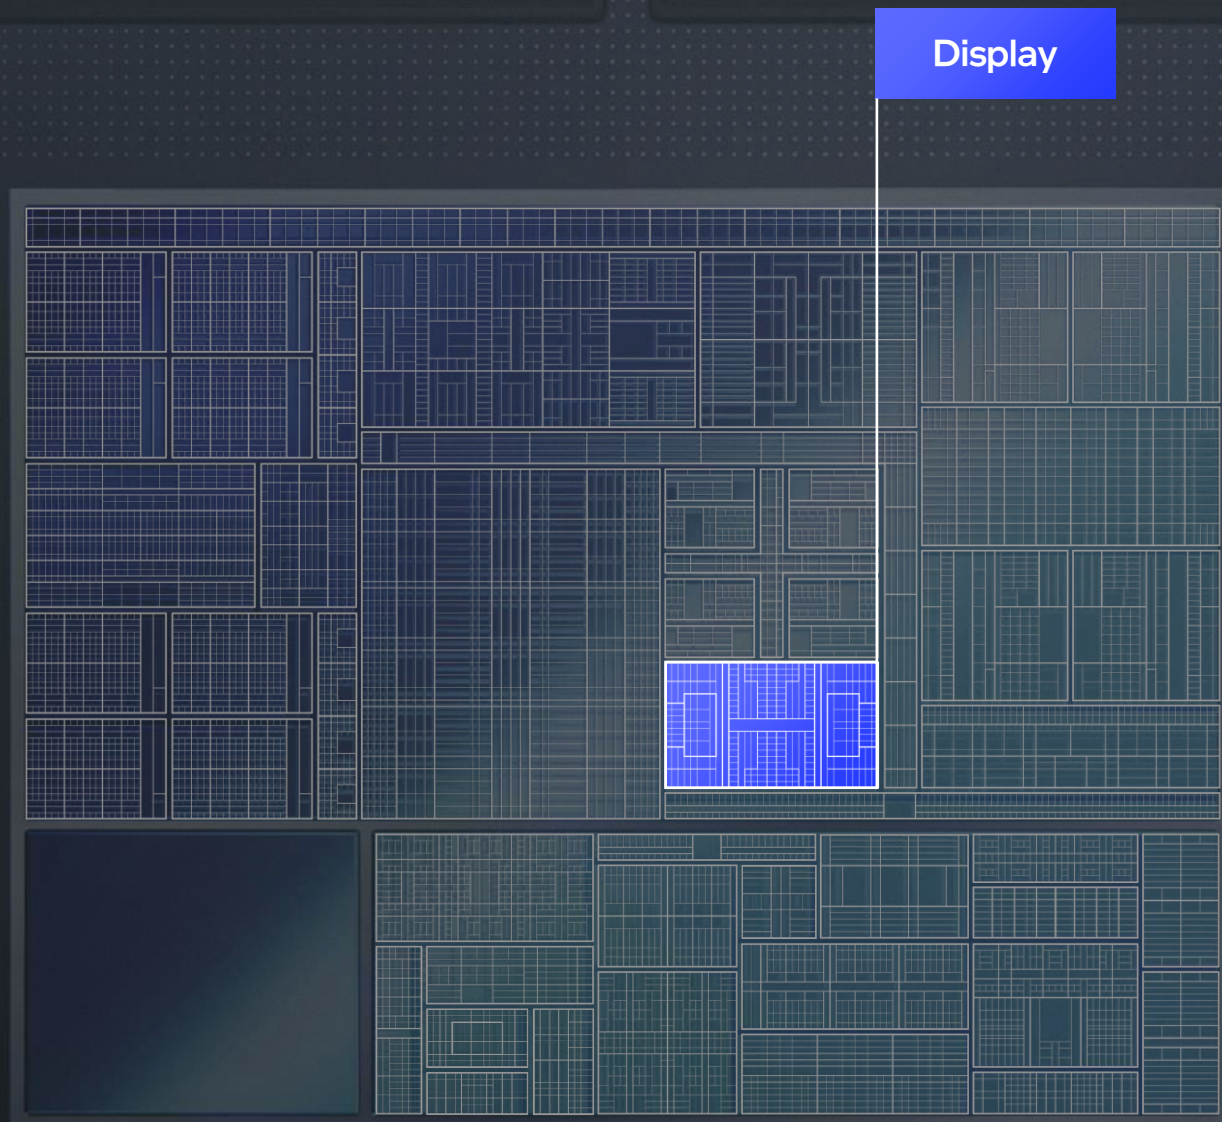
# Stable Diffusion

# Demo



Lunar Lake

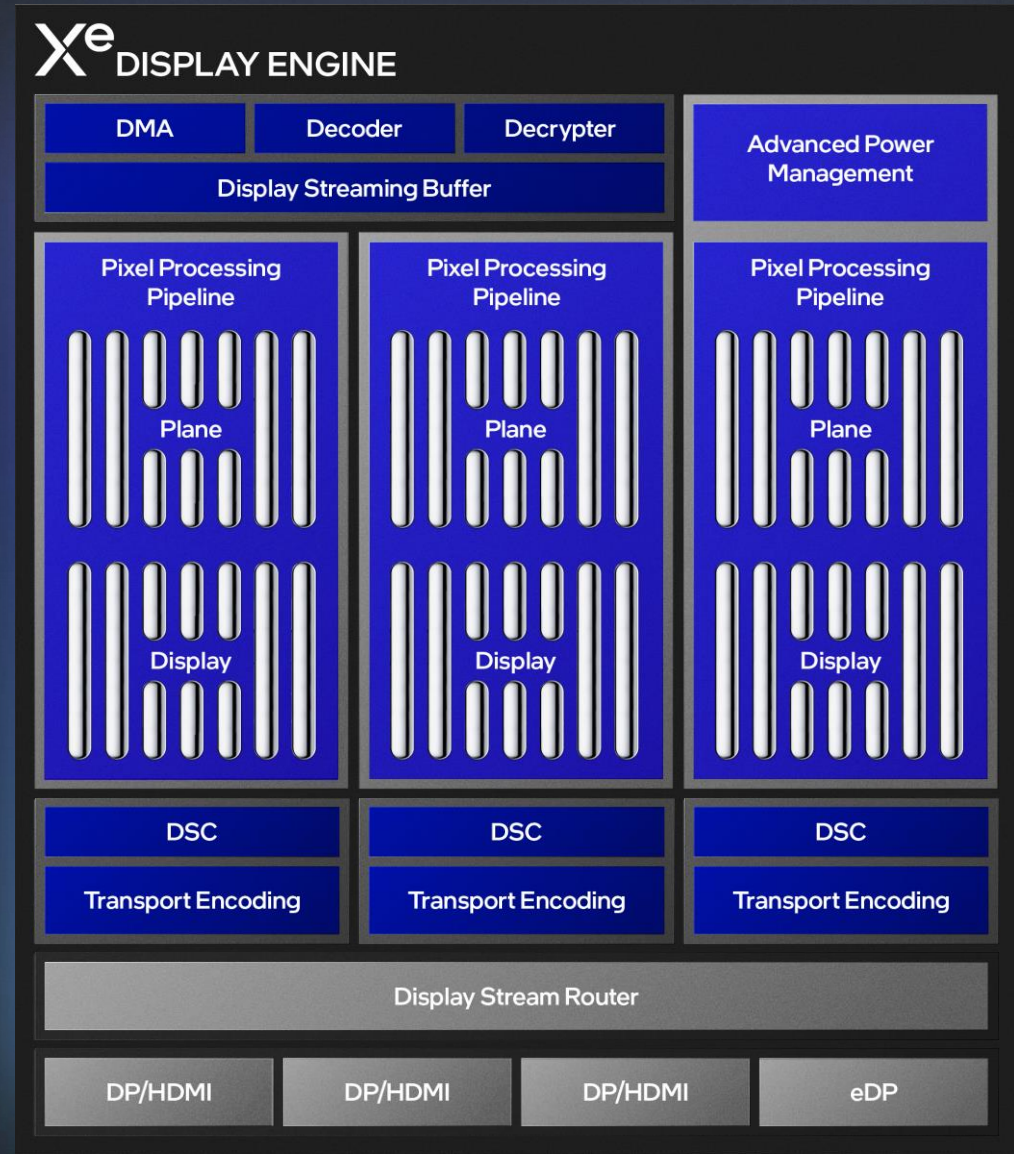
# Display Engine



Display

# Lunar Lake Display Engine

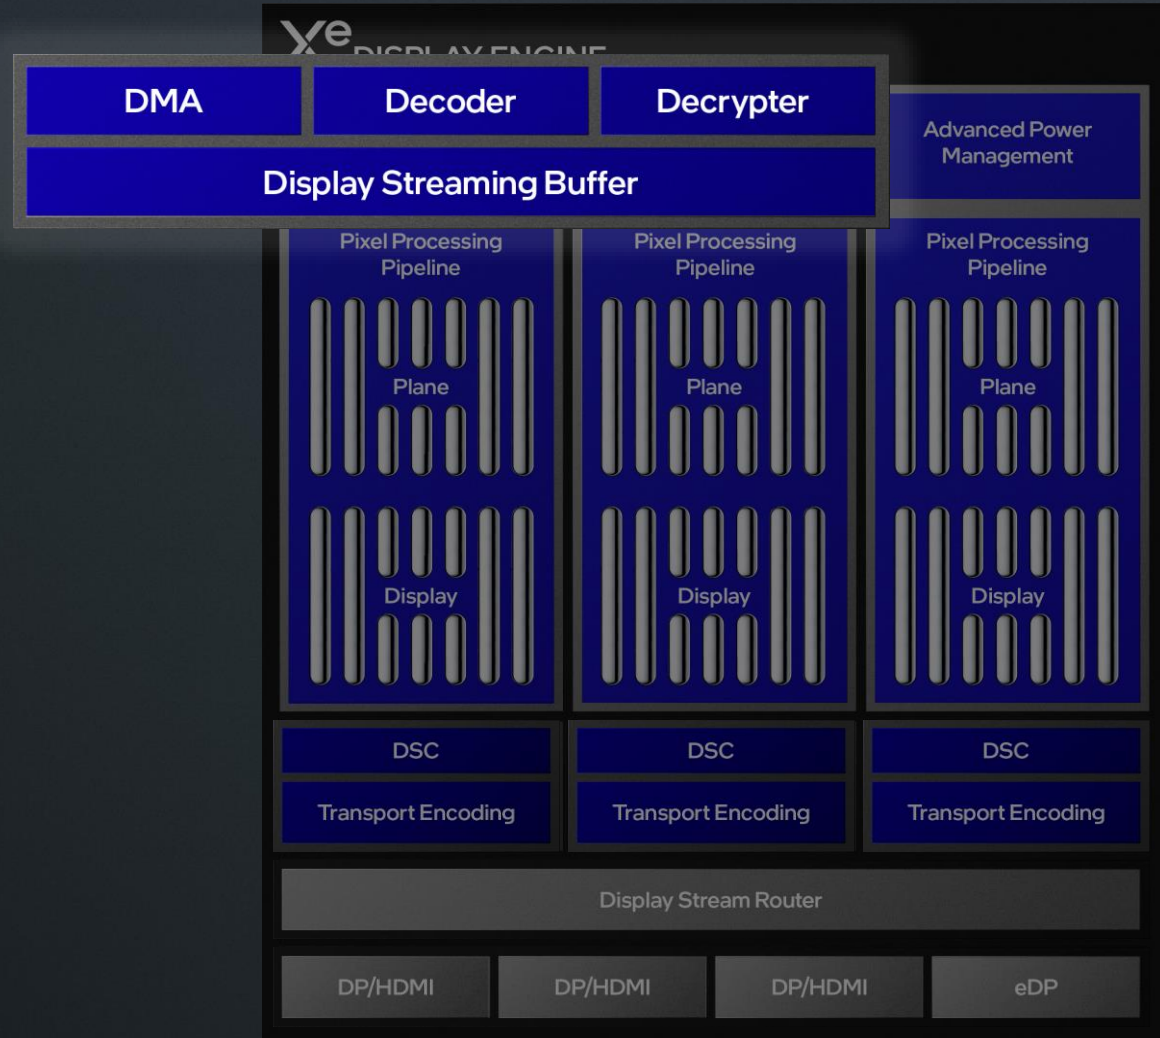
3 Display pipes	Up to <b>8k60 HDR</b>	Up to <b>3x 4k60 HDR</b>	Up to <b>1080p360 1440p360</b>
	HDMI 2.1	Display port 2.1	eDP 1.5 <span style="background-color: white; color: black; padding: 2px;">New</span>



# Display Engine Front End

Decode and decrypt

Streaming buffer



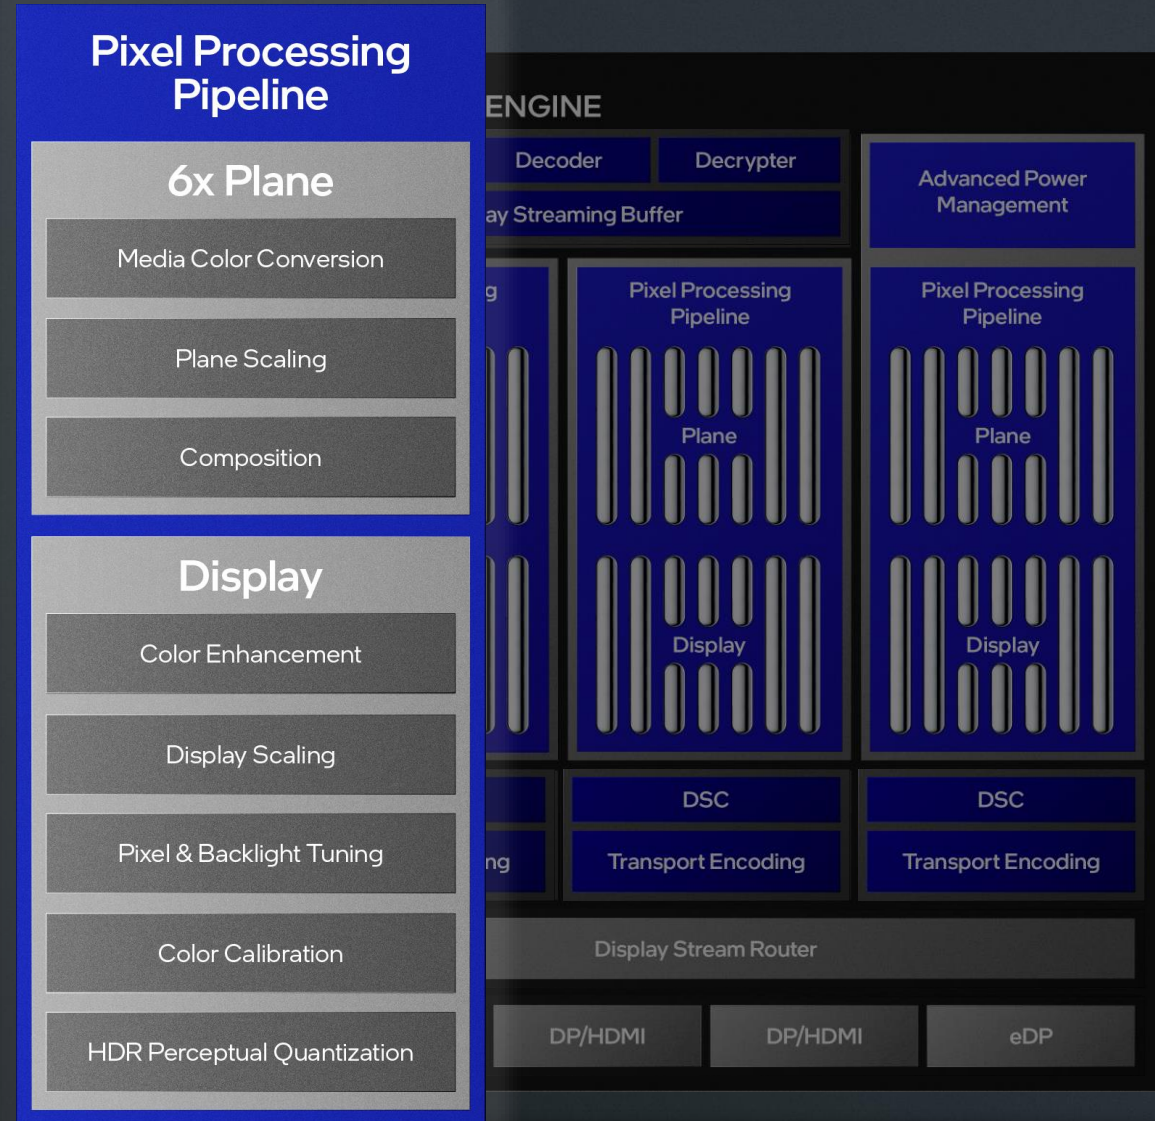
# Display Engine Pixel Processing Pipeline

## 6 planes per pipeline

Hardware support for color conversion and composition

## Flexible and power efficient

Designed to match any input format to any output format



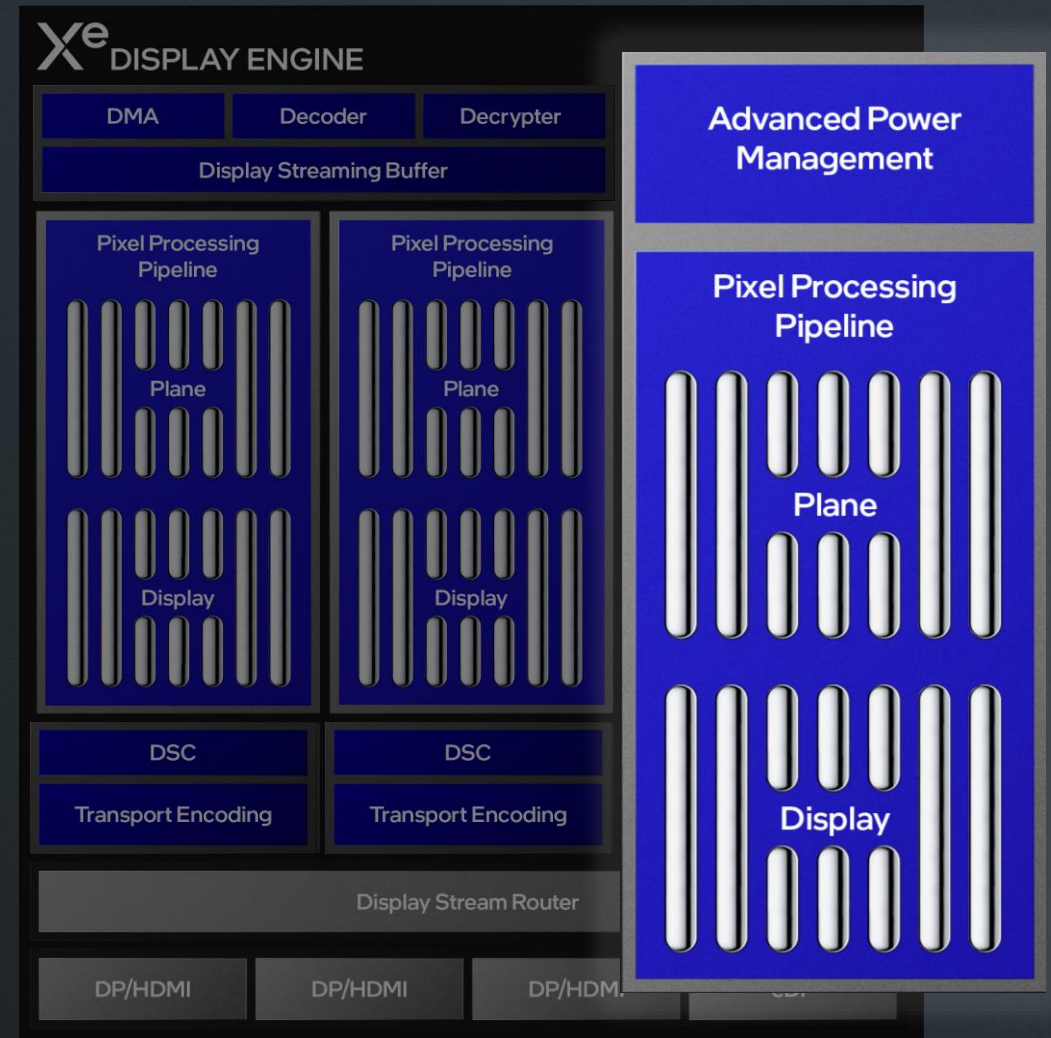
# Display Engine Low Power Optimized Pipeline

## Panel replay

Power gating during idle frames

## Brightness sensor with LACE

Local Adaptive Contrast Enhancement





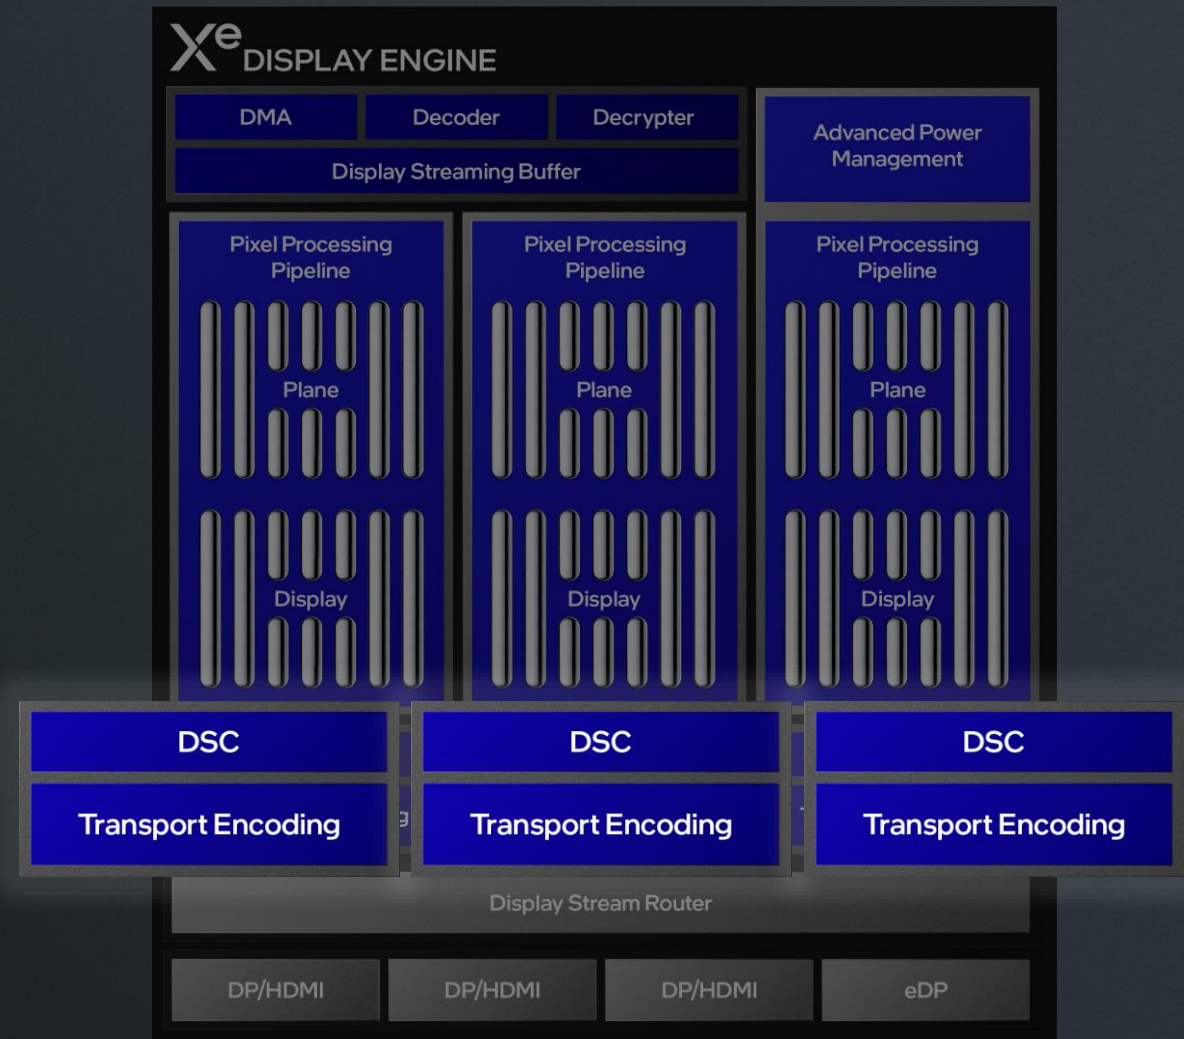
# Display Engine Compression & Encoding

## Display stream compression

3:1 visually lossless compression

## Transport encoding

Stream encode for HDMI and  
DisplayPort protocols



# Display Engine Router & Ports

## Stream assembly

Combine dual pixel pipeline streams and drive multi-stream transport

## Port Routing

Up to 4 ports are supported for added flexibility, including one eDP port





# eDisplayPort 1.5


## Panel Replay

Evolution of  
Panel Self  
Refresh

Selective  
update with  
early transport


Adaptive sync  
with panel  
replay

# Display Engine Content Matched Refresh Rate



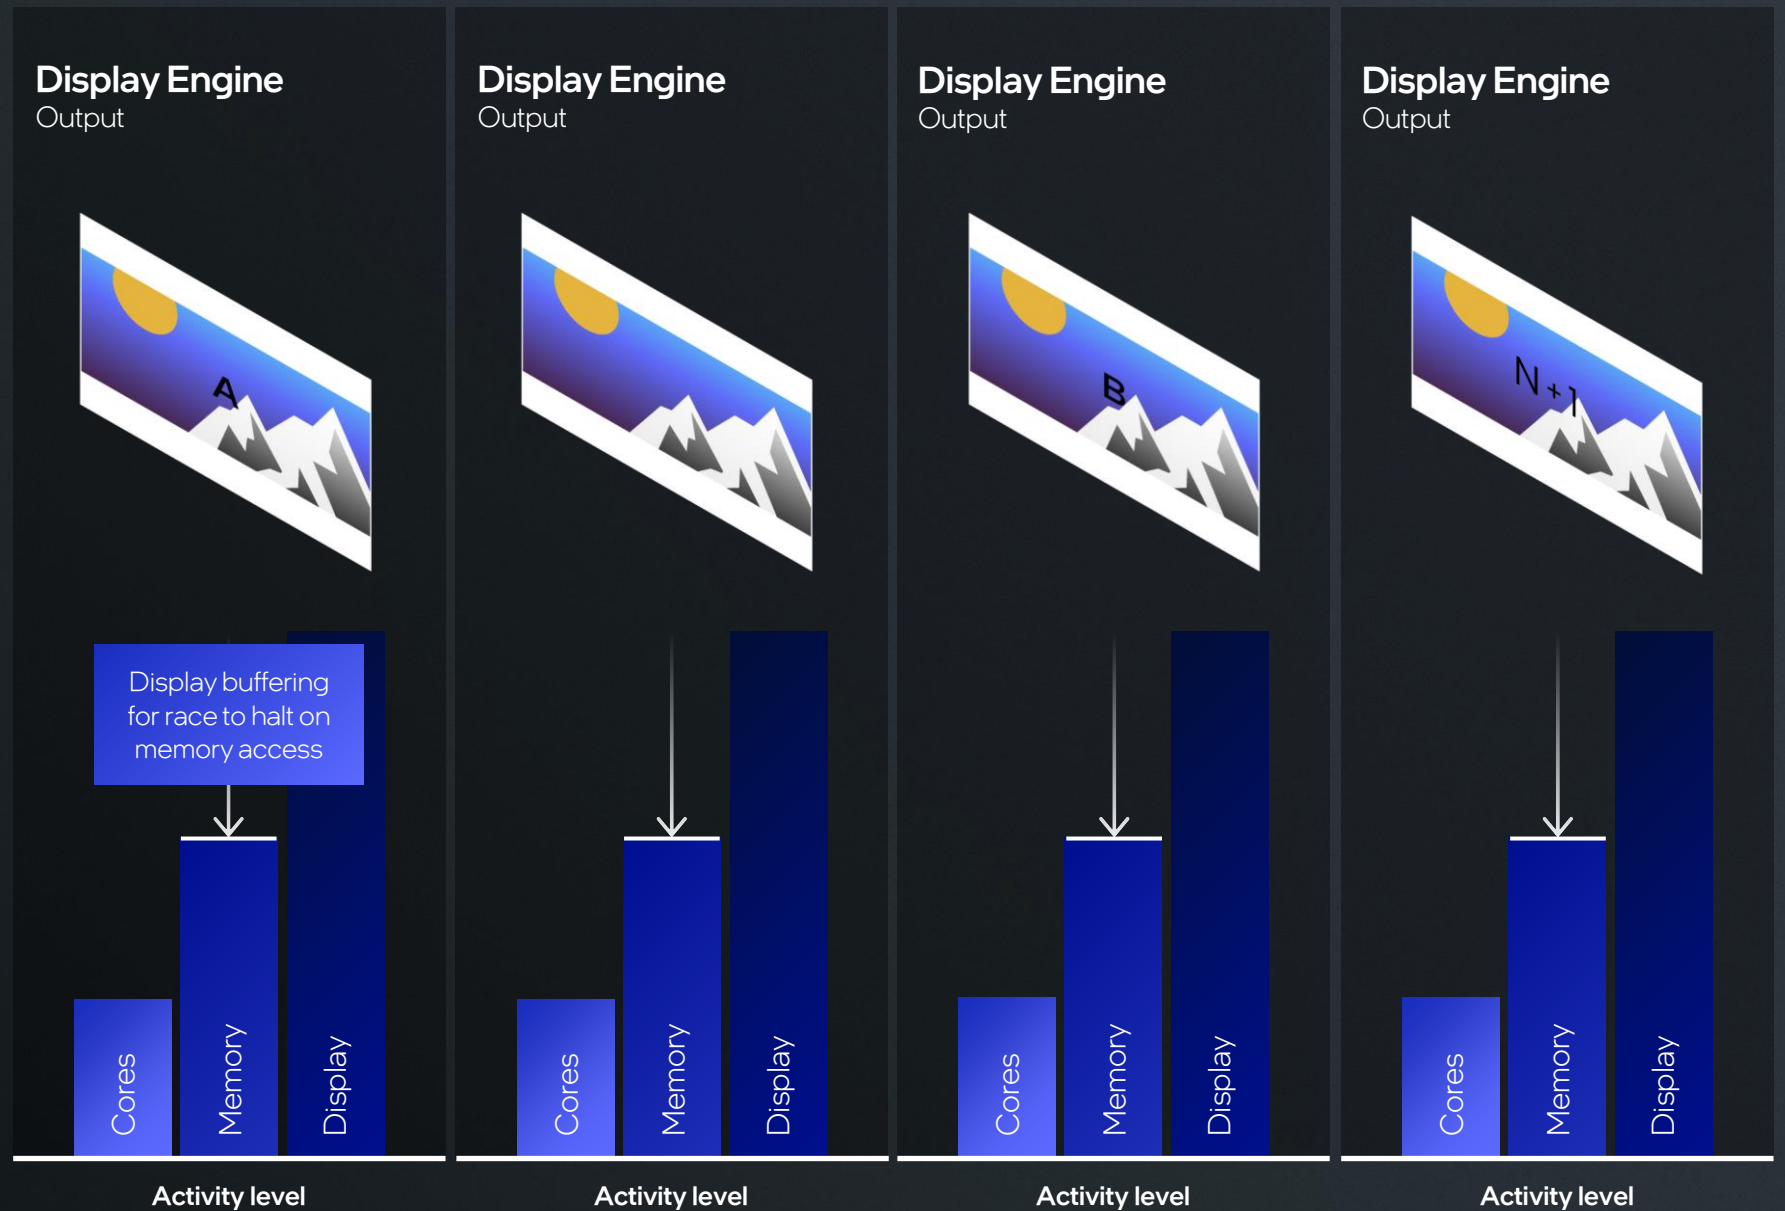
# Display Engine Power Optimization

Legacy



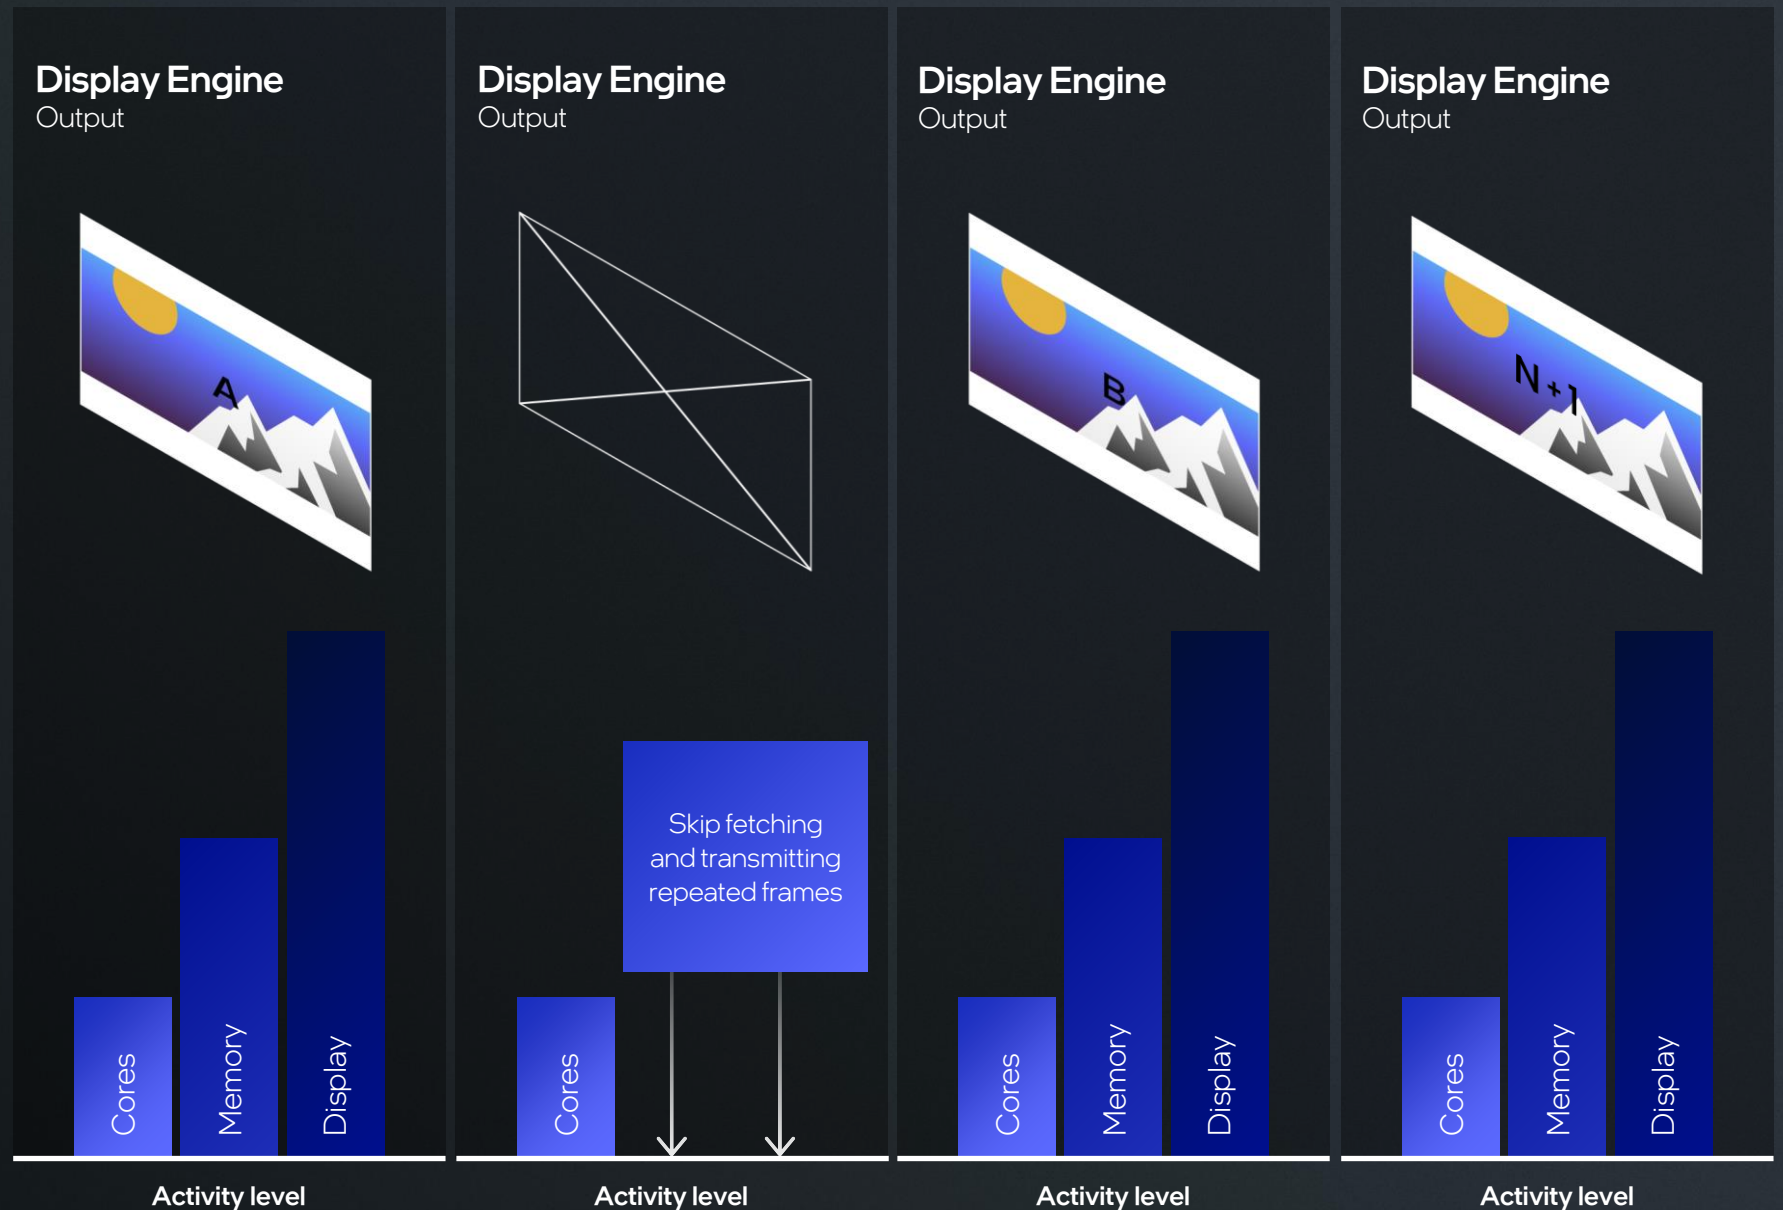
# Display Engine Power Optimization

Burst fill



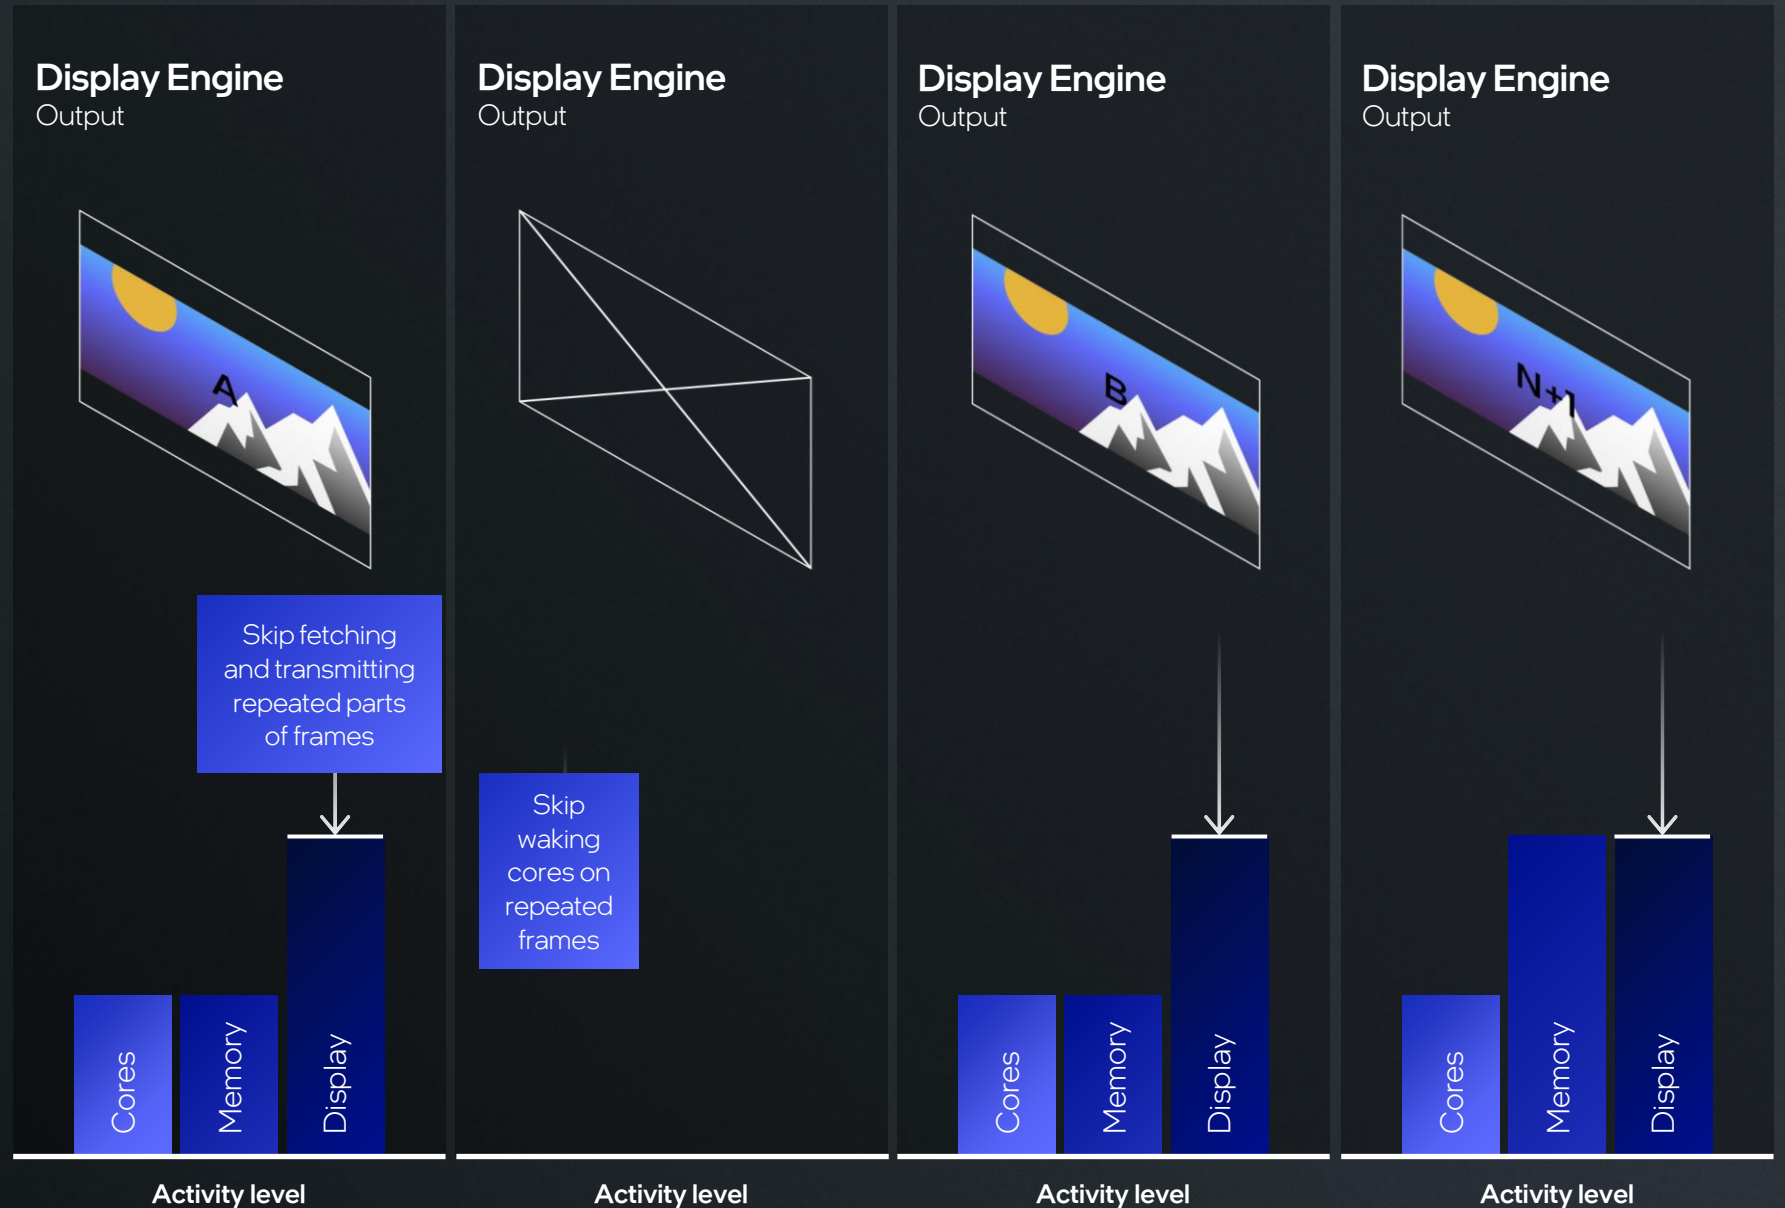
# Display Engine Power Optimization

Panel self refresh



# Display Engine Power Optimization

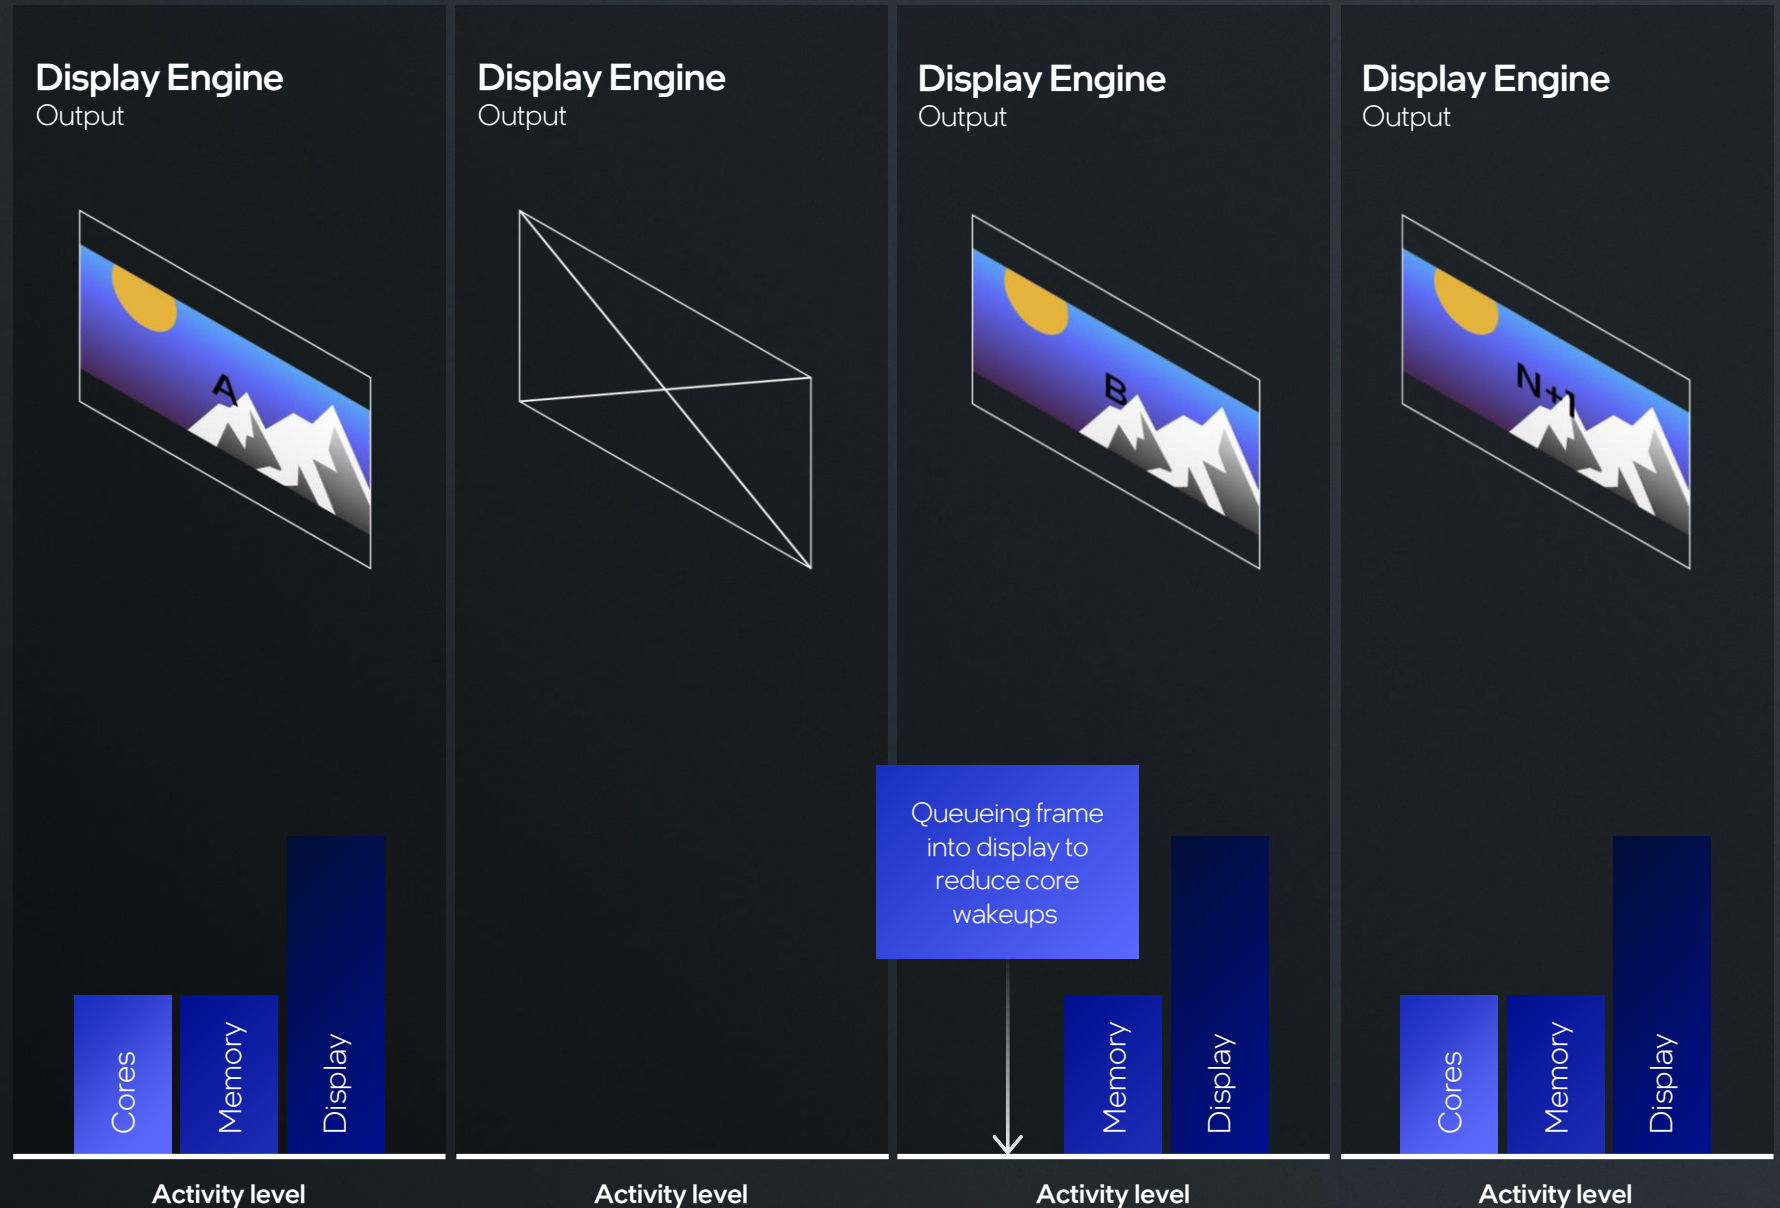
Selective update and optimized vertical blanking interrupts





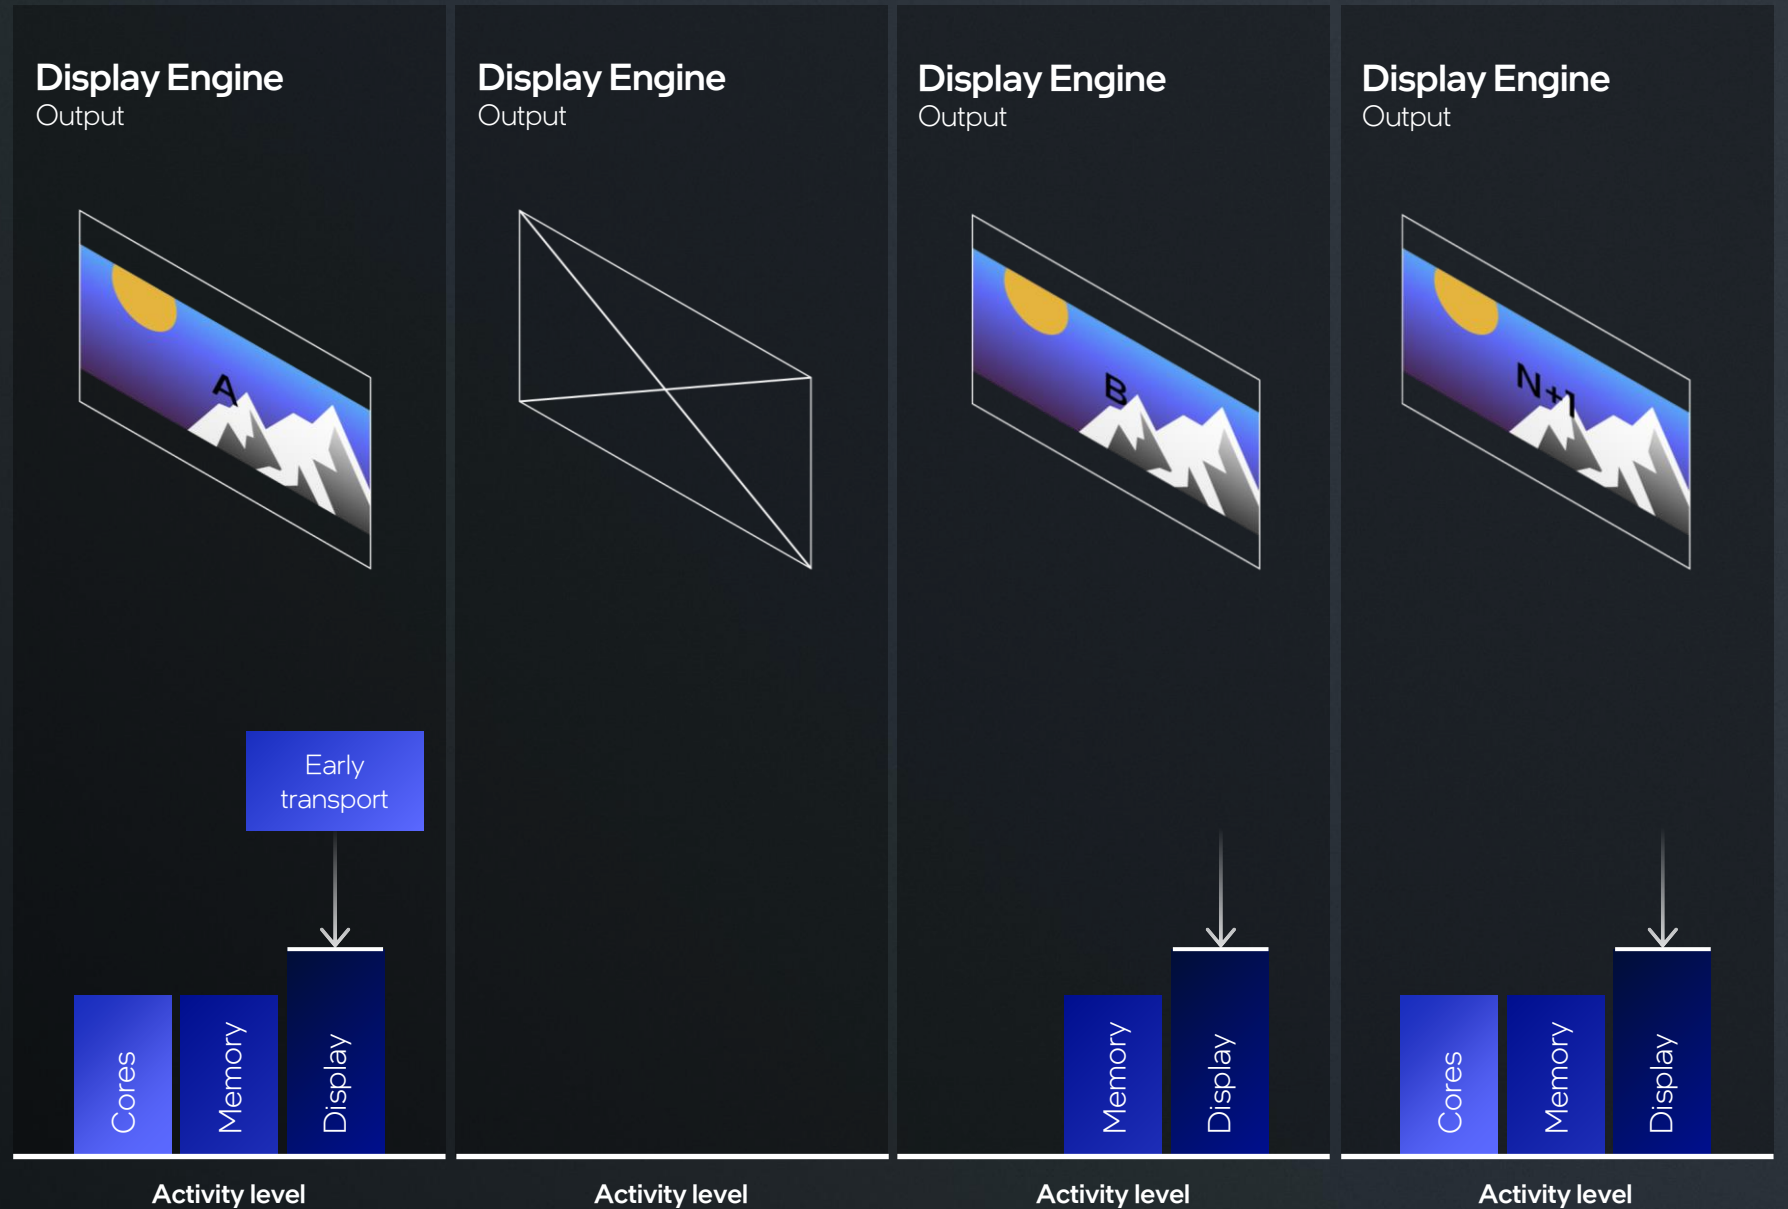
# Display Engine Power Optimization

Selective update and hardware queuing



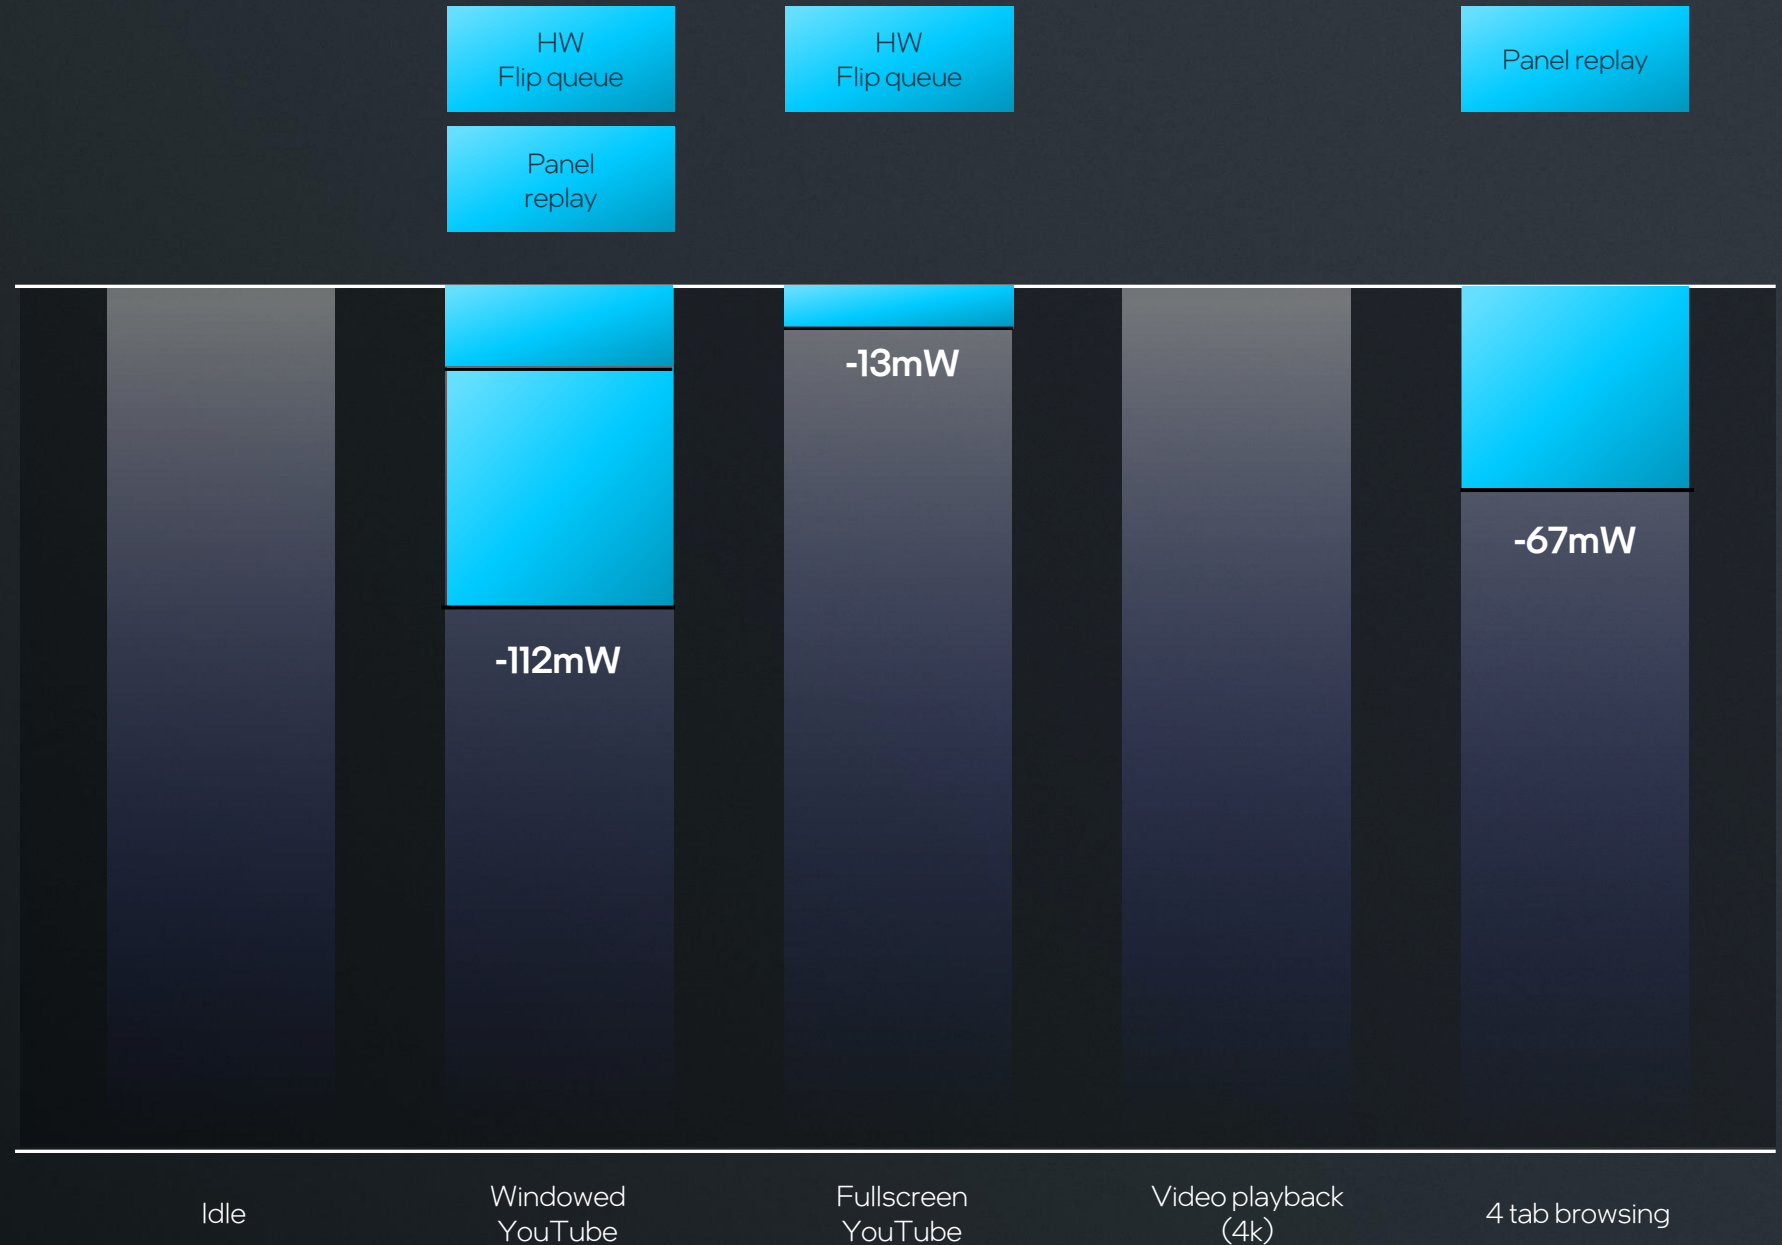
# Display Engine Power Optimization

+ early transport




# Display Engine Power Optimization

Efficiency gains




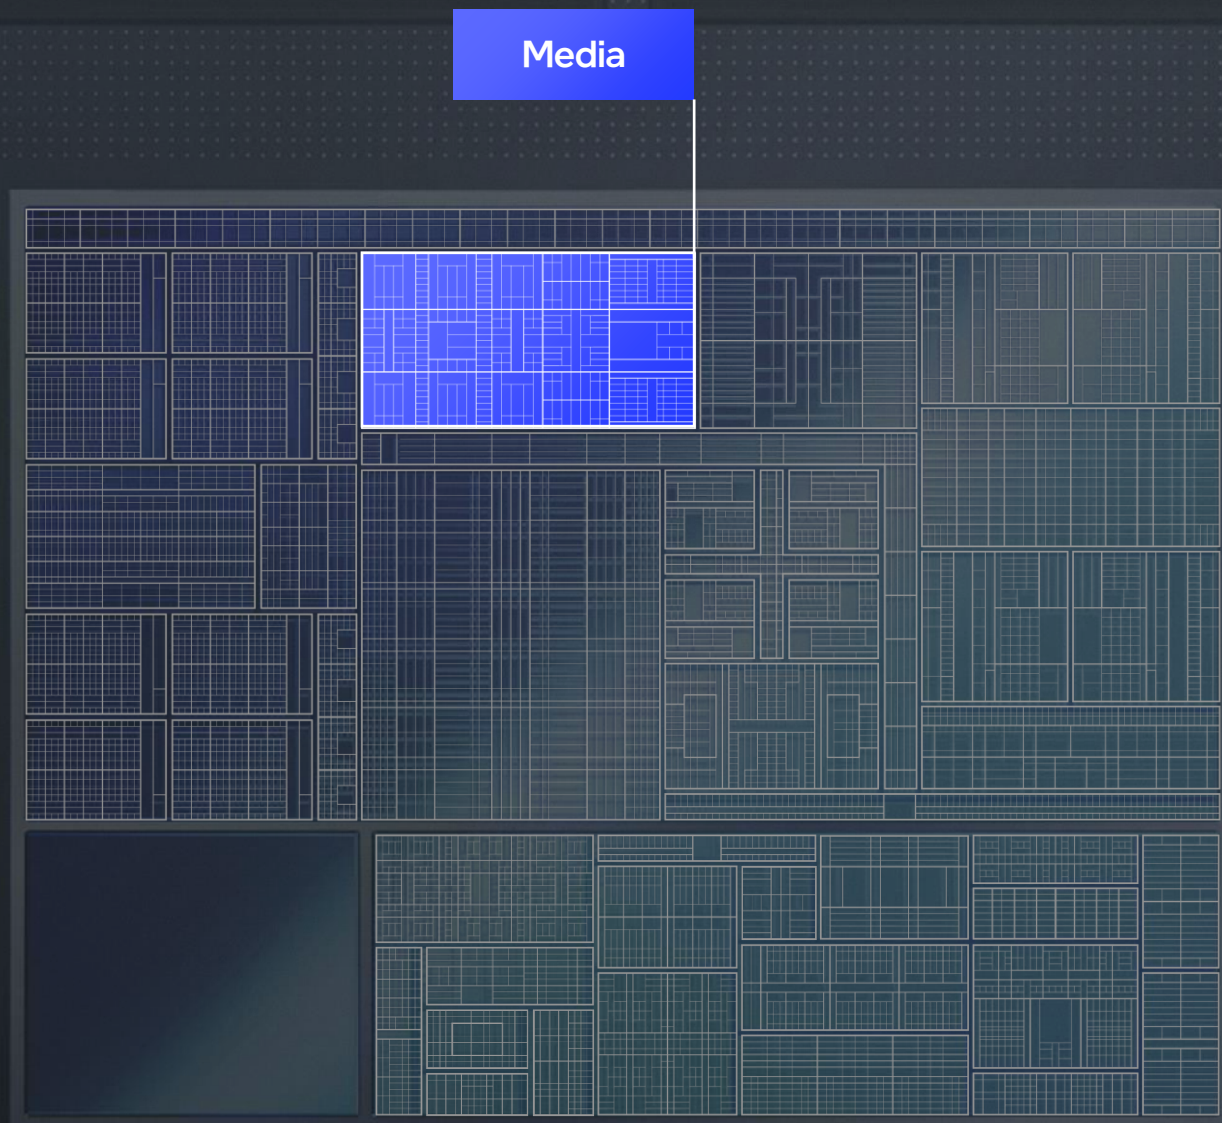
# Display Engine Power Optimization

Efficiency gains



Lunar Lake

# Media Engine




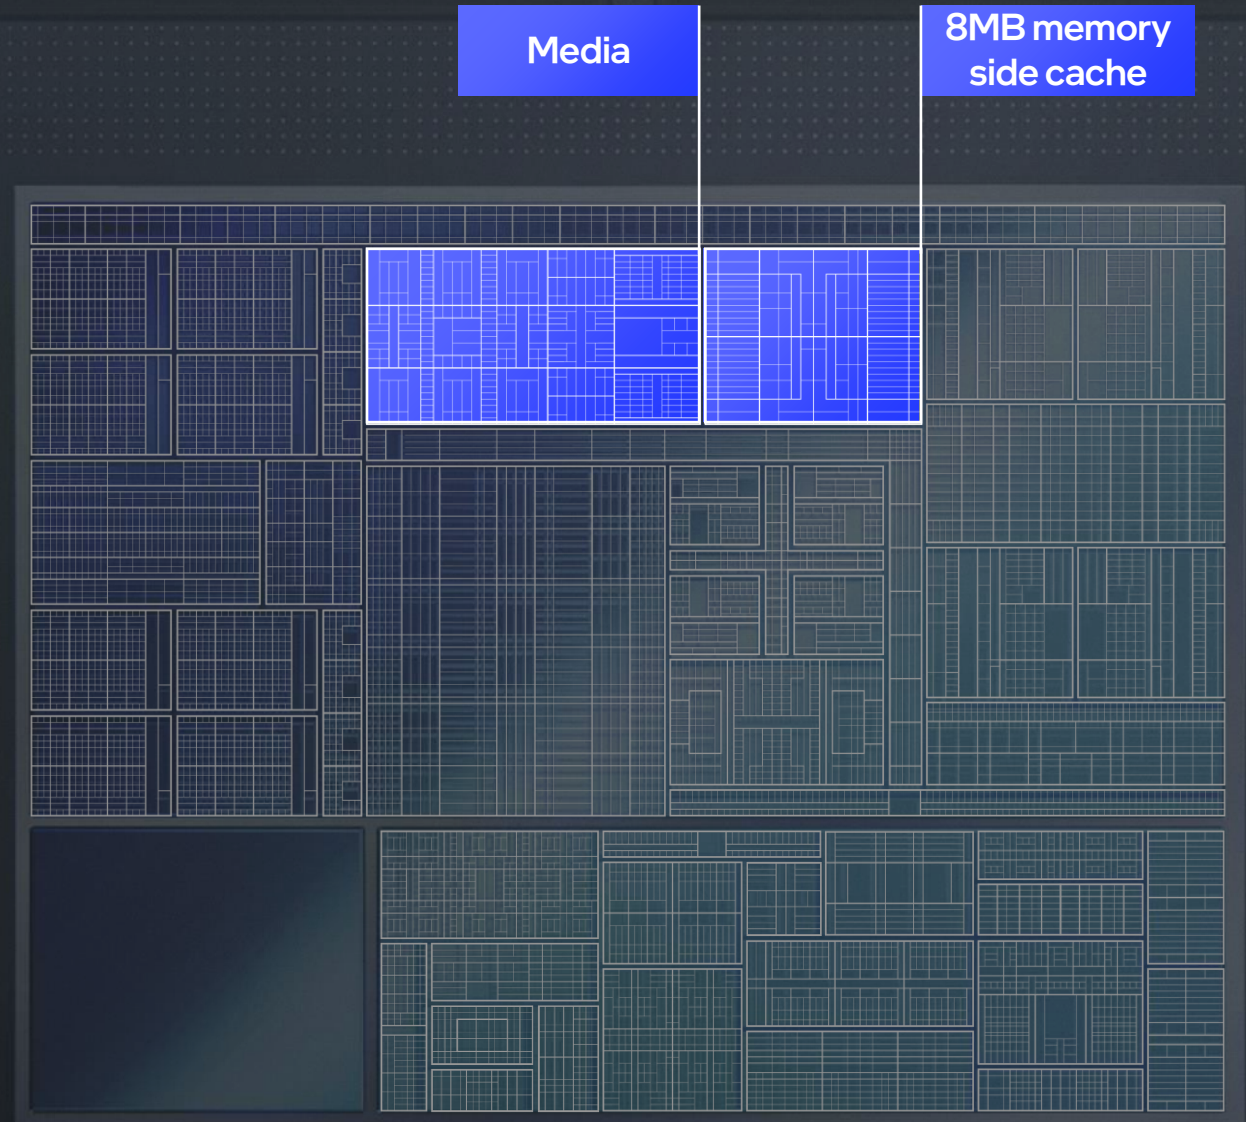
# New Memory Side Cache

## Bandwidth savings

Reduction in traffic to system memory  
across media workloads

## Power savings

Significant power reduction  
for encode workloads



# Lunar Lake Media Engine

Up to  
**8k60 10-bit**  
HDR decode

Up to  
**8k60 10-bit**  
HDR encode

AVC

VP9

**H.265**  
**HEVC**

**AV1**

**VVC**

New

**Xe** MEDIA ENGINE

**MFX**

Decoder

Encoder

Video Scaler

Color Space Converter

Video Enhancer

HDR Tone Mapper

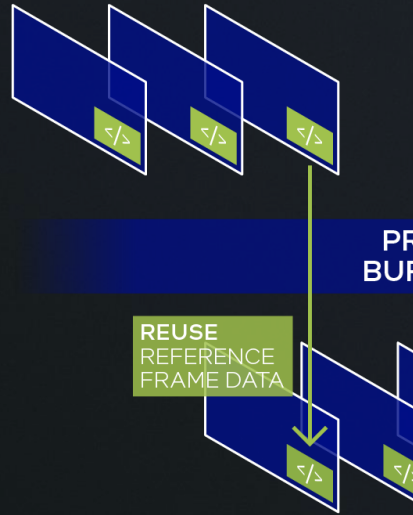
Bayer Processor




Significantly  
Reducing Bitrate at  
the Same Quality

10%

File size reduction  
over AV1



```
lude "XeSSRuntime" lude "XeSSRuntime"  
lude "XeSSJitter" lude "XeSSJitter"  
lude "Utility" lude "Utility"  
lude "Graphics" lude "Graphics"  
lude "ColorBuffer" lude "ColorBuffer"  
lude "DepthBuffer" lude "DepthBuffer"  
lude "CommandQueue" lude "CommandQueue"  
lude "Log.h" lude "Log.h"  
lude "Display" lude "Display"  
lude "xess/xess" lude "xess/xess"
```



Reduction  
in file size

Adaptive  
resolution  
streaming

Screen  
content  
coding

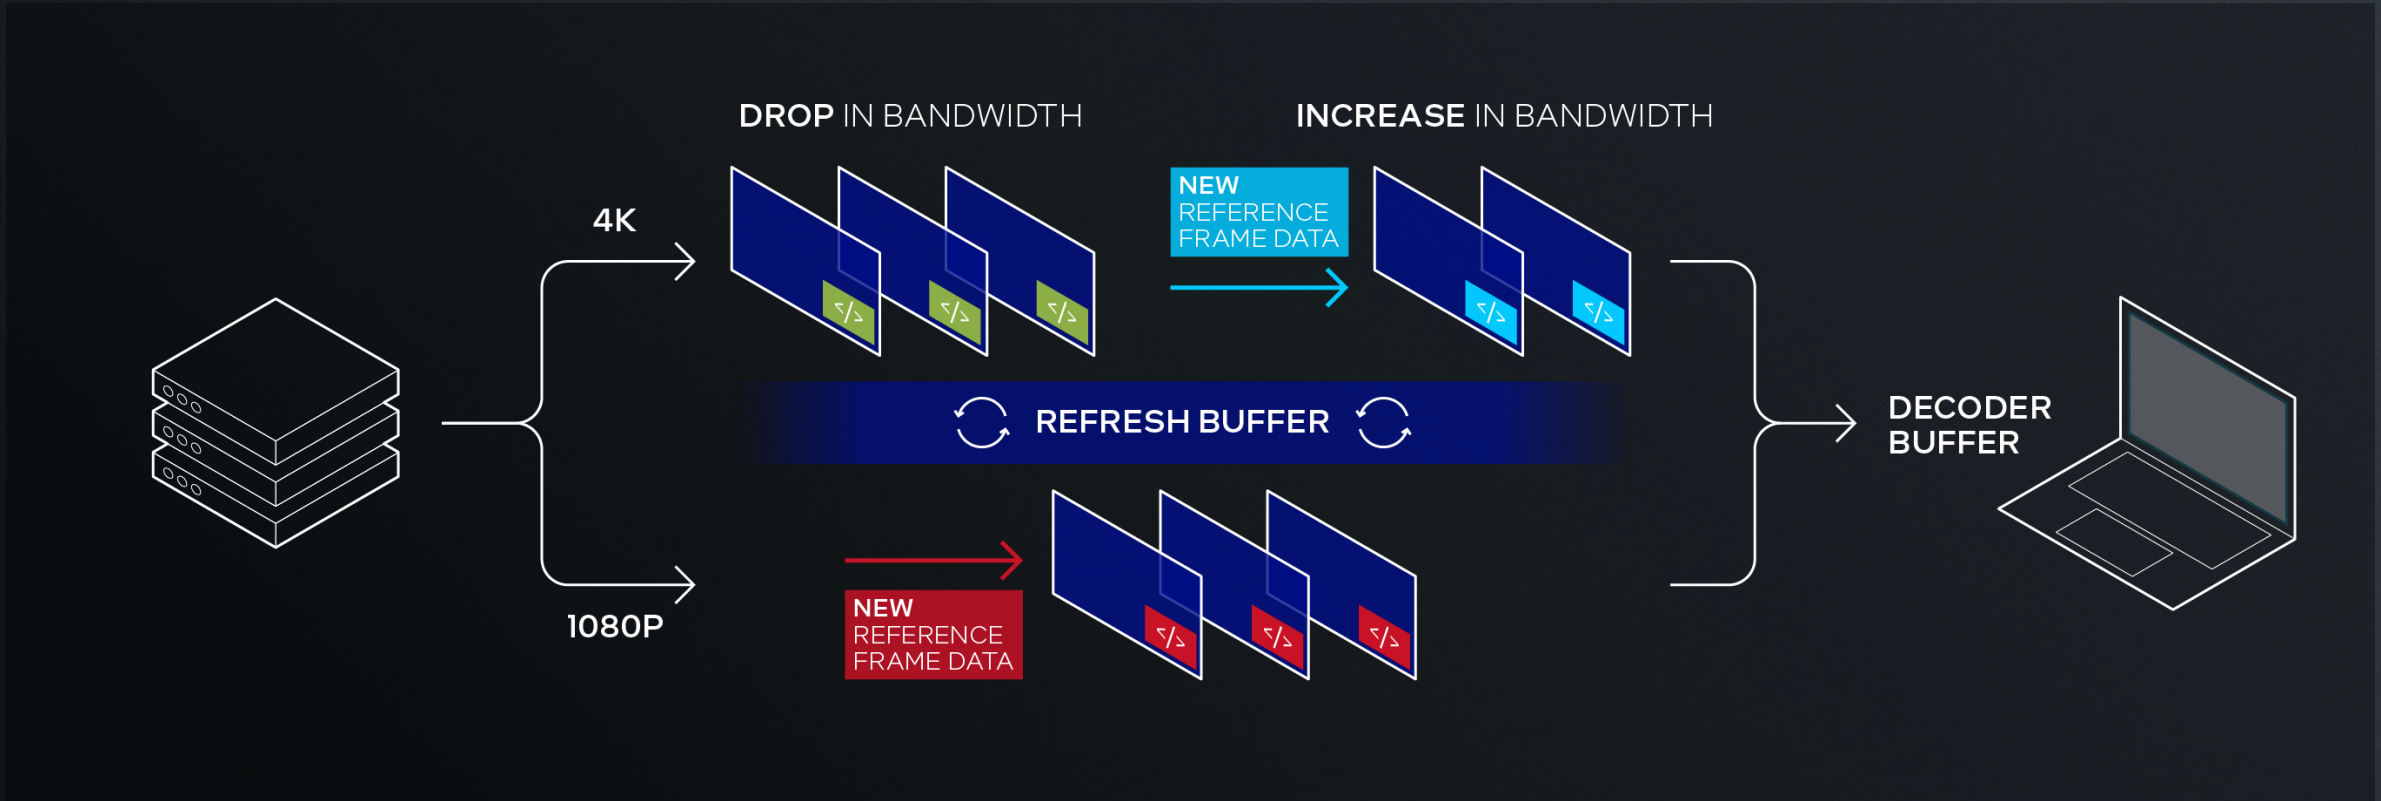
360-degree  
& panoramic



# Traditional Resolution Change

Send new reference data

Refresh decode buffer

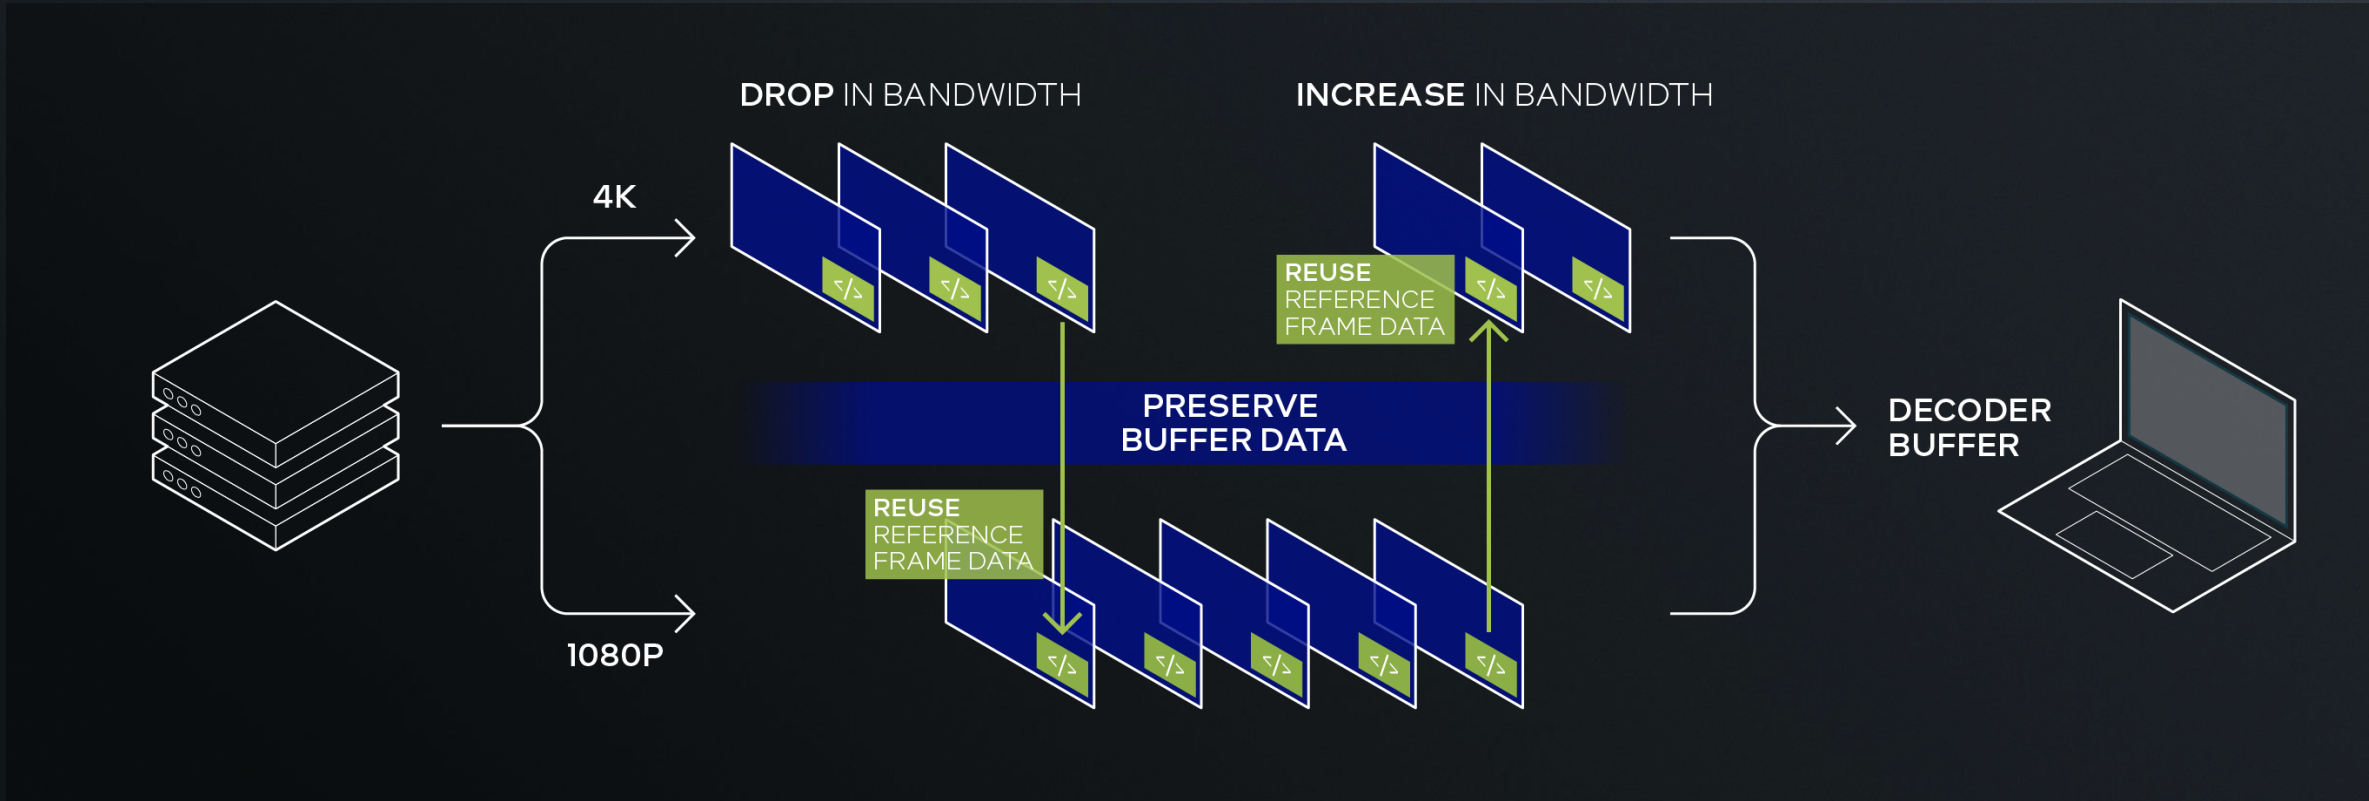




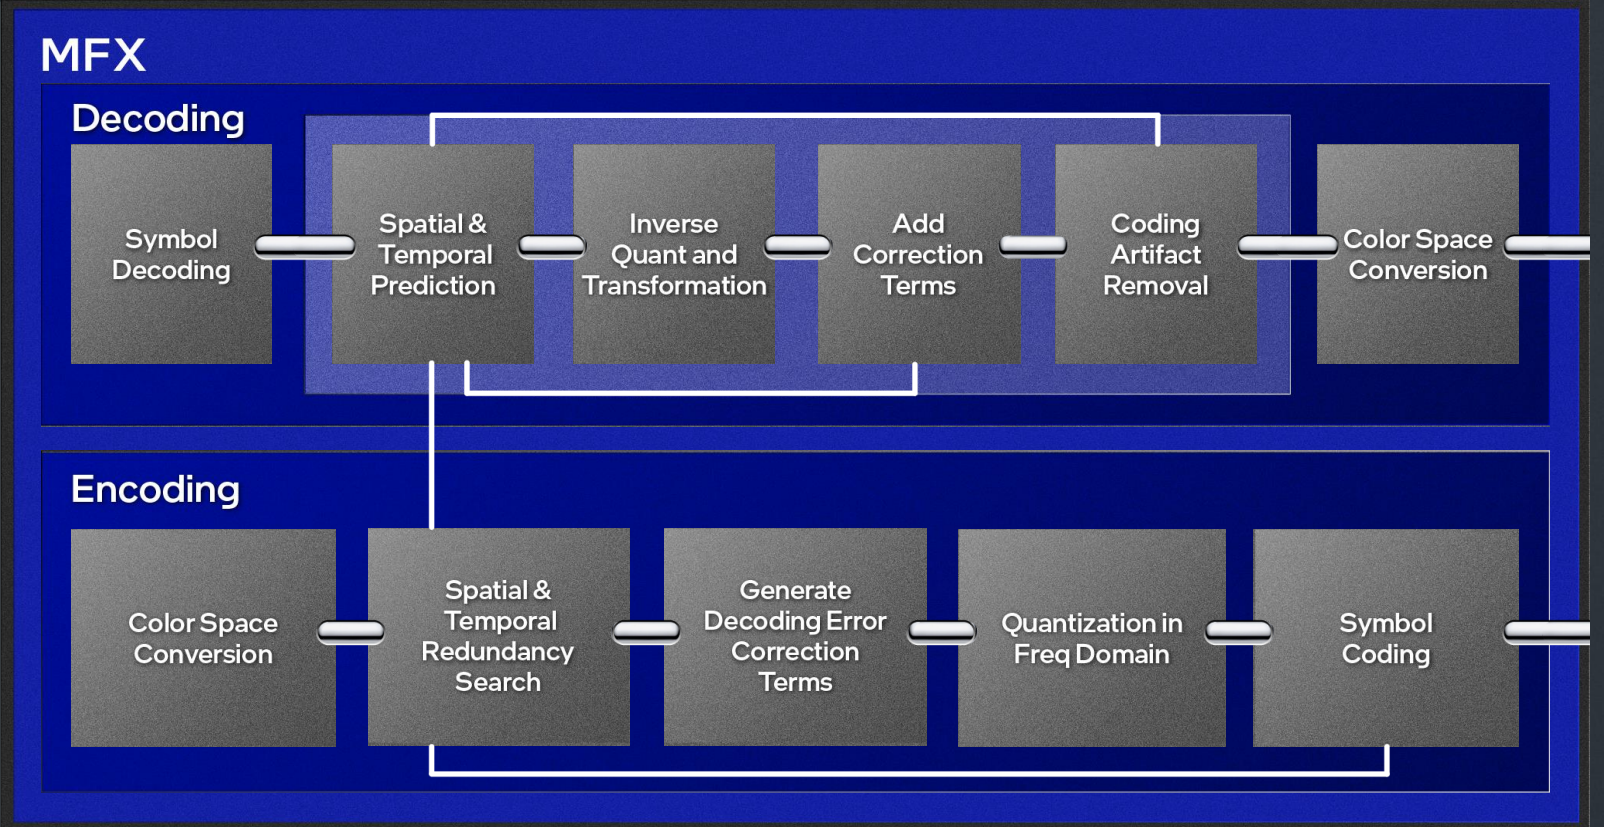


# Adaptive Resolution Streaming

Less data transfer

Less stream buffering



# Xe MEDIA ENGINE



H.265  
HEVC

AV1

VVC

# Screen Content Coding

Screen sharing

Remote desktop

Game streaming

```
#include "pch.h"

#include "XeSSRuntime.h"
#include "XeSSJitter.h"
#include "Utility.h"
#include "GraphicsCore.h"
#include "ColorBuffer.h"
#include "DepthBuffer.h"
#include "CommandContext.h"
#include "Log.h"
#include "Display.h"

#include "xess/xess_d3d12_debug.h"
```

AV1 with SCC


```
#include "pch.h"

#include "XeSSRuntime.h"
#include "XeSSJitter.h"
#include "Utility.h"
#include "GraphicsCore.h"
#include "ColorBuffer.h"
#include "DepthBuffer.h"
#include "CommandContext.h"
#include "Log.h"
#include "Display.h"

#include "xess/xess_d3d12_debug.h"
```

AV1 without SCC


# Evolution of Media Codecs



\*Same visual quality, Software Encoding | 1080p | Very Slow Preset.

Source: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11842/118420T/Towards-much-better-SVT-AV1-quality-cycles-tradeoffs-for-VOD/10.1117/12.2595598.full?tab=ArticleLinkFigureTable>

 **VVC**



Demo

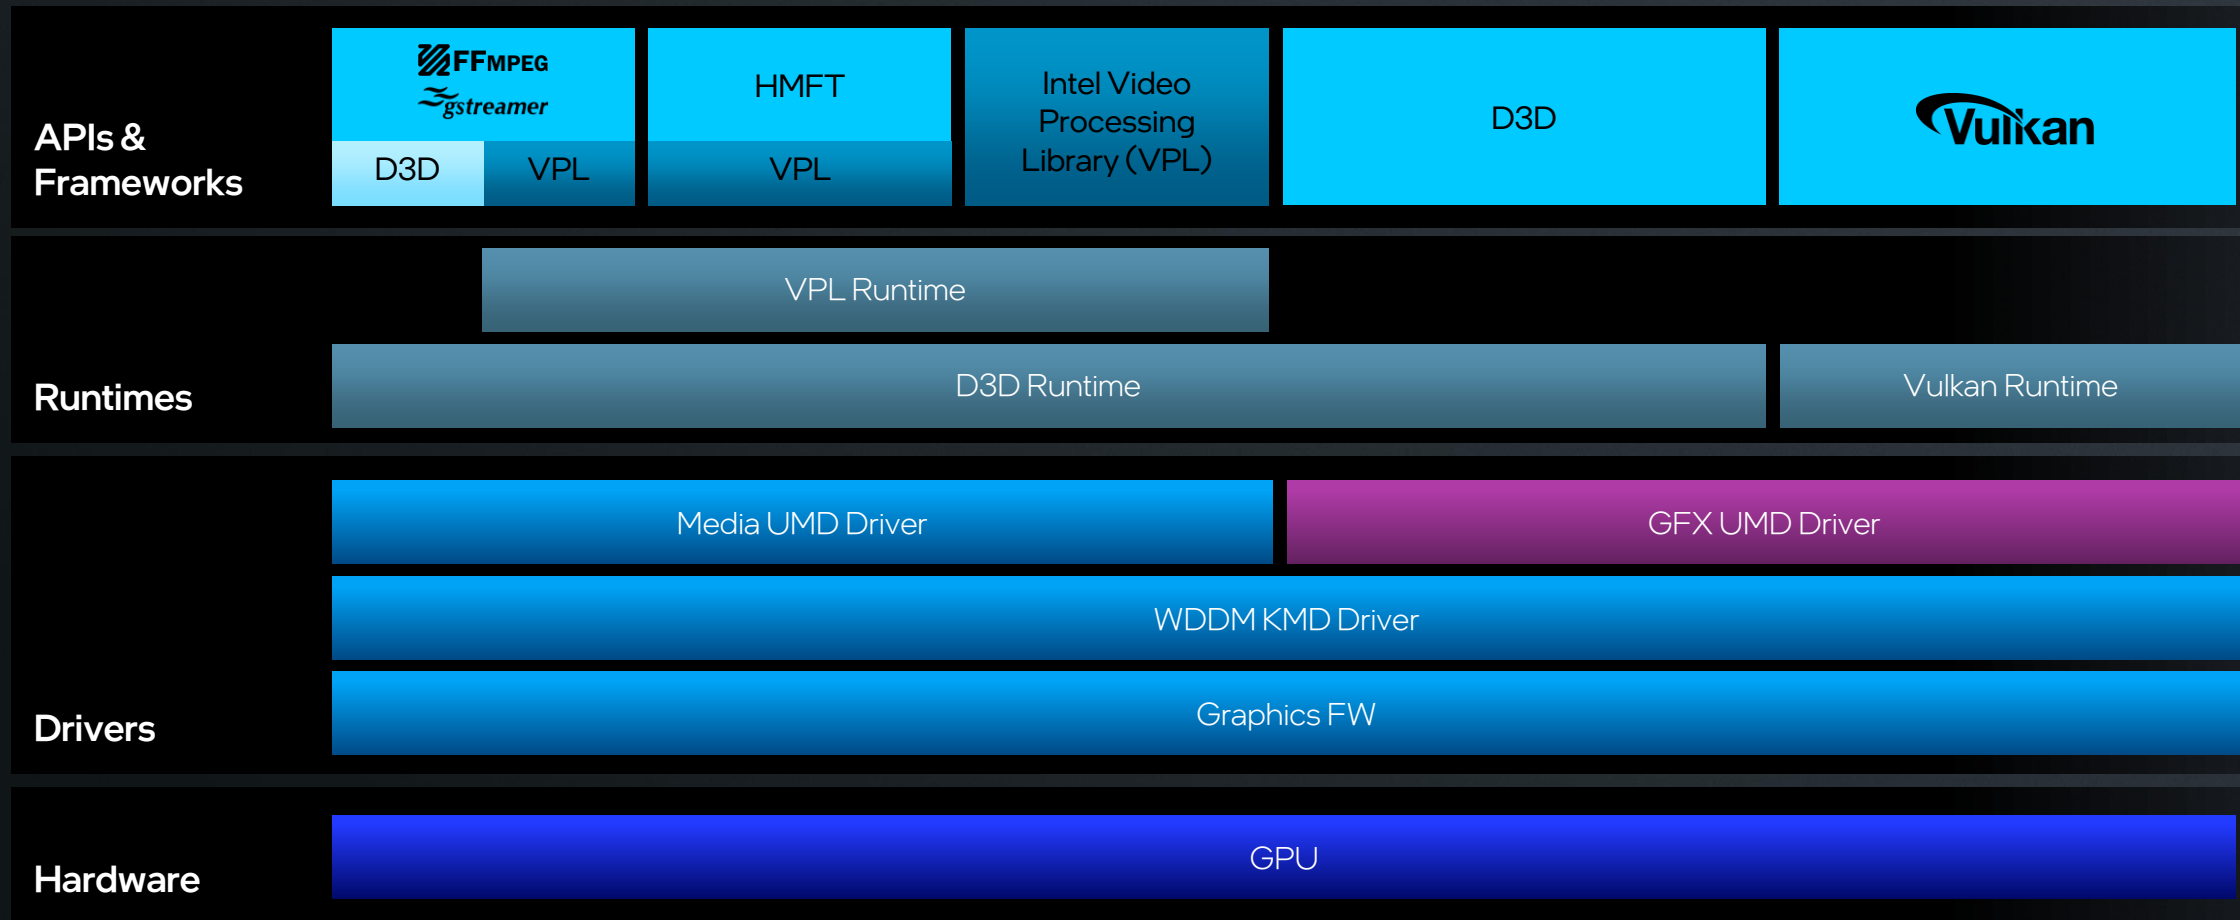
Lunar Lake

# GPU SW Stack



# Windows GPU Software Stack


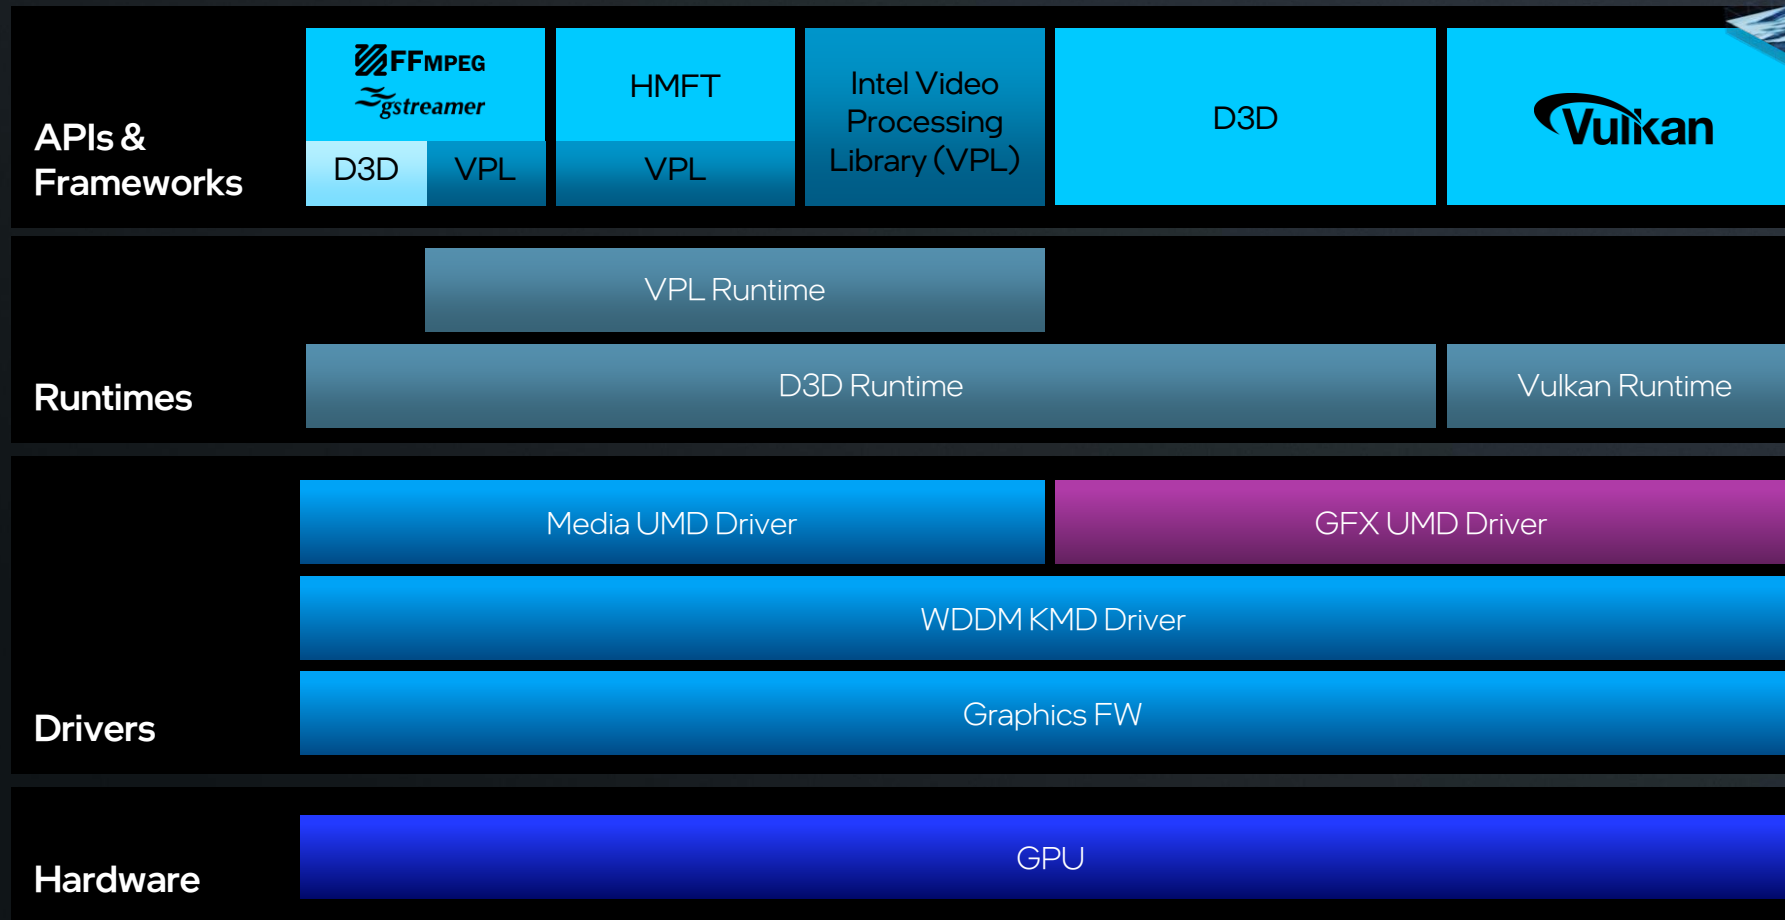
Ready for Xe2





# Windows GPU Software Stack

Ready for Xe<sup>e</sup>2



F1 '24 gaming



Demo

# Lunar Lake Graphics

Better, faster and more efficient on all fronts

up to  
**67 TOPS**

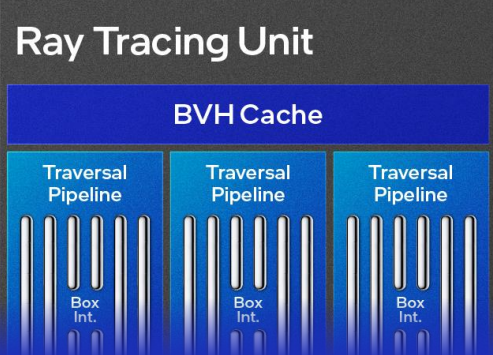
**1.5x**  
faster graphics performance  
vs. Meteor Lake GPU

**new**  
XM<sub>X</sub> engines




**8 enhanced**  
ray tracing units

Ray Tracing Unit



**8**  
**2<sup>nd</sup> gen**  
Xe<sup>e</sup>-cores

**enhanced**  
XeSS  
kernels



**D**  
eDP 1.5

**8MB**  
L2 cache

**efficiency optimized**  
media & display engines



**VVC**  
Decode

**AV1**  
Encode & decode

Up to  
**3x**  
4k60 HDR



The Intel logo is positioned in the top-left corner of the image. It consists of the word "intel." in a white, lowercase, sans-serif font, set against a solid blue rectangular background.

intel.

The logo for "TECH tour.TW" is located in the upper-left quadrant. It features the word "TECH" in a bold, white, uppercase font, with a small blue square containing a white dot to its right. Below "TECH" is the text "tour.TW" in a white, lowercase font. The entire logo is set against a blue rectangular background.

**TECH** .  
tour.TW

The words "Thank You" are centered in the lower half of the image. They are written in a large, white, sans-serif font. The text is overlaid on a dark blue, semi-transparent rectangular area that has a subtle gradient and a slight shadow effect.

Thank  
You

# Notices & Disclaimers

The preceding presentation contains product features that are currently under development. Information shown through the presentation is based on current expectations and subject to change without notice.

Results that are based on pre-production systems and components as well as results that have been estimated or simulated using an Intel Reference Platform (an internal example new system), internal Intel analysis or architecture simulation or modeling are provided to you for informational purposes only. Results may vary based on future changes to any systems, components, specifications or configurations.

Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](http://www.intel.com/PerformanceIndex).

AI features may require software purchase, subscription or enablement by a software or platform provider, or may have specific configuration or compatibility requirements. Details at [www.intel.com/AIPC](http://www.intel.com/AIPC).

No product or component can be absolutely secure. Intel technologies may require enabled hardware, software or service activation.

All product plans and roadmaps are subject to change without notice.

Performance hybrid architecture combines two core microarchitectures, Performance-cores (P-cores) and Efficient-cores (E-cores), on a single processor die first introduced on 12th Gen Intel® Core™ processors. Select 12th Gen and newer Intel® Core™ processors do not have performance hybrid architecture, only P-cores or E-cores, and may have the same cache size. See [ark.intel.com](http://ark.intel.com) for SKU details, including cache size and core frequency.

Built-in Intel® Arc™ GPU only available on select Intel® Core™ Ultra processor-powered systems; OEM enablement required.

Some images may have been altered or simulated and are for illustrative purposes only.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

# APPENDIX

Claim # & Statement	Slide # & Title/Details
	SLIDES 4 & 18:Improving IP performance efficiency
Xe2 IP performance per Xe-core is 1.2x to 12.5x higher than Xe1 IP across a set of various graphics functions.	Results are based on an internal suite of micro benchmarks and collected on a pre-release Xe2 engineering platform with pre-release GFX software. The comparison is a selected subset of micro benchmarks normalized for equal Xe-cores configuration and clock frequency.
	SLIDE 25: Lunar Lake Xe2 GPU Performance
1.5x graphics performance over Meteor Lake	Testing by Intel as of May2024. Data based on Lunar Lake reference validation platform measurement vs Meteor Lake reference validation platform as measured by 3DM Time Spy. 3DMark*
	SLIDES 43-44: Display Engine Power Optimization
Lunar Lake’s Display Engine benefits from a list of power savings optimization across a set of use cases	Testing conducted by Intel’s Display Engine engineering team to validate functionality of various power savings features on pre-release engineering platform with pre-release software.
	SLIDE 59: Lunar Lake Graphics
1.5x faster graphics performance vs. Meteor Lake GPU	Testing by Intel as of May2024. Data based on Lunar Lake reference validation platform measurement vs Meteor Lake reference validation platform as measured by 3DM Time Spy. 3DMark*

The Intel logo is centered on a dark blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®).

intel®