

Whitepaper



# Breaking Limit

True cross-platform ray tracing benchmark





## Introduction

*In a future where robots and AI technologies will be able to perform complex tasks and run factories, are we going to break the limit between artificial and real? Or are we driving at full speed towards a point of no return?*

The barrier between reality and simulation, in real-time graphics, is getting increasingly thinner with every year. With the help of AI techniques, it is likely that we will eventually achieve near-flawless realism. However, there is still some progress to be made before we reach that point.

Nevertheless, real-time ray tracing has been around for quite some time and was the first step toward achieving realism. Only a year and a half after Basemark released the first-ever ray tracing benchmark for mobile, the technology has already become essential in the handheld market. Many brands now offer devices that support ray tracing on mobile platforms. In desktop environments, real-time ray tracing established a firm foothold years ago, paving the way for the next technological advancements.

But what is the exact difference between desktops and mobiles when it comes to real-time graphics? What are the next steps the technology will follow? What will the next chapter be? These are some of the questions that GPUScore: Breaking Limit aims to answer.

# Breaking Limit Overview

Breaking Limit is a benchmark designed for desktop systems, laptops, and handheld devices. It features a true cross-platform workload (Breaking Limit) intended for premium smartphones, tablets, laptops, and lower-end desktop systems. Additionally, it includes a dedicated workload (Breaking Limit Ultra) for desktop systems and laptops with high-end GPUs.

What makes Breaking Limit truly unique is its cross-platform, allowing the exact same benchmark to run on all platforms and environments. As such, it is a very heavy benchmark for smartphones and lighter for powerful laptops and desktop systems. Therefore, we provide the Breaking Limit Ultra workload for high-end GPUs.

Both versions target ray tracing capable devices exclusively. For game developers, artists, and GPU device manufacturers, this benchmark offers valuable insights into the capabilities and content limitations that can be achieved. Additionally, as a true cross-platform benchmark, it allows precise comparisons of the performance of phones vs. laptops vs. desktops. By using the same workload without optimizations for different platforms, it provides accurate performance comparisons.

Since upscaling is often used in conjunction with ray tracing technologies to increase the frame rate, Breaking Limit Ultra includes two new upscaling tests.





## Technical Features

Both the cross-platform and Ultra workloads employ a hybrid ray tracing approach, using ray queries in conjunction with rasterization. As in all previous benchmarks, frame-based rendering is utilized in the Official benchmarks. The application uses offscreen rendering and shows a set of 9 rendered frames per presentation. This is done to prevent drivers from skipping frames and to provide more headroom for high-performing devices. The benchmarks consist of 2100 rendered frames.

### Breaking Limit (the cross-platform workload)

The cross-platform workload supports all common operating systems: Windows, Linux, iOS and Android. It has support for all of the modern graphics APIs – DirectX 12, Vulkan, and Metal. The rendering resolution used in the official benchmarks is Full HD (1920x1080).

As ray tracing effects, this workload uses Ray traced Reflections and Ray Traced Shadows. The shadows of the main directional light are ray traced. Other lights rely on shadow maps to preserve performance. The ray tracing of reflections is done at 75% of the main resolution to save performance while maintaining enough precision.

Other utilized high-end rendering features include Ground-truth Ambient Occlusion (GTAO), volumetric fog, motion blur, and Temporal Anti-aliasing (TAA). The cross-platform workload uses 11 shadow-casting light sources and the total GPU memory consumed is approximately 2.1GB.

To achieve high quality with a limited performance budget, the workload utilizes different optimization techniques, such as the usage of LOD meshes, similarly as games.

## Breaking Limit Ultra

Breaking Limit Ultra is a very intense workload, targeting laptops and desktop computers with high-end discrete GPUs. It is compatible with Windows and Linux and supports both DirectX 12 and Vulkan. The rendering resolution of the Official benchmark is 3840x2160.

Just like the cross-platform version, Breaking Limit Ultra uses Ray Traced Reflections and Ray Traced Shadows. In addition to those, it also uses Ray Traced Global Illumination. The resolution used in Ray Traced Reflections is higher than in the cross-platform workload and more lights are ray traced for shadows. The resolution used for reflections is the same as the main rendering resolution, ie. 3840x2160. A maximum of 8 light sources shadows are ray traced per frame.

Other utilized high-end rendering features include Ground-truth Ambient Occlusion (GTAO), volumetric fog, motion blur, and temporal anti-aliasing (TAA).

This workload has an average of 2.5 million triangles in the camera frustum, uses 49 shadow-casting light sources, having a total GPU memory consumption of approximately 6GB.

A distinctive feature of this workload is the inclusion of AI upscaling techniques in two distinct benchmarks. One employs NVidia's DLSS2 (Deep learning super sampling 2) and the other one employs AMD's FSR2 (FidelityFX™ Super Resolution 2).



# Included Benchmarks

GPUScore: Breaking Limit consists of several benchmarks, which offer information on the graphics performance from different perspectives.

## Official

The Official test runs the benchmark using the standard settings. Both the cross-platform and Ultra workloads have an Official benchmark.

**Breaking Limit** Official benchmark runs the benchmark at Full HD (1920x1080) resolution while **Breaking Limit Ultra** Official runs it at Ultra HD (3840x2160). Ultra workload also uses more demanding rendering settings to achieve higher quality.

This is the main benchmark, which provides an overall view of the gaming performance of a system. Results from these benchmarks can be uploaded to Powerboard, GPUScore's (formerly known as Basemark) online results platform.

## Official Native

This benchmark is available only on mobile platforms and the used workload is the cross-platform workload. The Official Native renders images at the native resolution of the device display. It uses the same settings as the Official mode, except for the resolution.

The Official Native benchmark offers more detailed and accurate information about the practical performance of the device. Using a fixed resolution benchmark might not be enough to evaluate the performance of graphical applications. For many devices, the fixed resolution of the Official mode might be different than the actual resolution of the device. Therefore, the device may under or over-perform, in a real-world scenario.

## Official DLSS2 and Official FSR2

Available only on the Breaking Limit Ultra workload. The Upscaled tests measure the performance of the devices using two different upscaling techniques, AMD's FSR2 (FidelityFX™ Super Resolution 2) and NVidia's DLSS2 (Deep learning super sampling), in conjunction with ray tracing. The two techniques are not comparable side by side, because, in addition to performance differences, they have differences in quality. Therefore, these exist as separate benchmarks. The results from these tests can be uploaded into two different categories on the Powerboard platform.

Both techniques use a balanced quality preset, which might introduce small differences in the initial rendering resolution. In addition, the FSR test renders one part of the content, while the DLSS test renders the other one, which makes the results incomparable.

The rendering resolution is upscaled to 3840x2160. The performance is calculated by running each involved section of the content in full resolution and then running it again by using a lower resolution which is upscaled to full resolution. The average difference between the two performance results is displayed as a percentage, which indicates the performance gain that the device gets from the technique compared to rendering natively at full resolution.

On the Powerboard platform, as well as in the application UI, both performance graphs and performance results are displayed (minimum/average/maximum framerate and the percentage gain), in addition to the improvement percentage.

## Custom

For both workloads, a Custom benchmark is available, allowing the end-user to configure the settings of the benchmark. This helps users to evaluate performance and quality impacts of different features, such as texture resolution or the number of bounces used in Ray traced Reflections.

Being a non-standard benchmark, with a lot of variations, this test does not upload the results to Powerboard, as they would be inconsistent and not comparable with each other. The Custom benchmarks, for both workloads, features the settings listed below:

**Benchmark mode: On / Off**

- The On / Off toggle allows to run the benchmark in a frame-based manner (On), which can be used for benchmarking, or in a time-based manner (Off).
- Rendering in a frame-based mode is the mode used in the Official tests. During the run “Offscreen image count” number of images are rendered off-screen per one displayed image. One of the images is shown as the main image and the others are shown in a small grid at the bottom of the screen.
- The time-based mode (Off) renders as many frames as possible while keeping the animation speed constant. This is how games render images, but due to varying set of rendered images, it’s not optimal for accurate benchmarking and thus it doesn’t produce a score.

**Render resolution**

- The main rendering resolution can be changed to different values: 720x480 (SD), 1280x720 (HD), 1920x1080 (FHD), 2560x1440 (QHD), 3840x2160 (UHD), 7680x4320\* (8k)

**Workload\***

- Whether to use the cross-platform or the Ultra workload.

**Vertical Sync\***

- Whether to try to use some form of vertical synchronization, depending on the device.

**Graphics API\***

- Which graphics API to use.

**Graphics Device\***

- Which GPU to use if multiple are present in the system.

**Full Screen\***

- Whether the benchmark should run in full screen or a window.



**Window resolution\***

- The resolution of the window when full screen is disabled. There are several options: 720x480, 1280x720, 1920x1080, 2560x1440, and 3840x2160. The Window resolution is different from the render resolution.

**Upscaling method\***

- The user can choose to optionally enable either FSR2 or DLSS2 upscaling technique.

**Upscaling quality\***

- When upscaling is enabled, this can be used to choose between different quality presets. These influence the quality of the upscaled output in relation to the performance. The available options are Performance, Balanced, and Quality.

**Loop count**

- The benchmark can be set to run infinitely in a loop or for a specific number of times. This is very useful for example to monitor the stability of performance on mobile devices.

**Off-screen Image Count**

- This option is only effective when benchmark mode is enabled. This option allows the user to change the number of frames rendered per displayed image. The first image is displayed as a larger image and the rest are shown as a grid below the main image. The benchmark performs frame-based rendering but presents results to the user every Nth frame depending on this value.

**Time Step**

- The animation time step between frames can be customized from the default value. The larger this value is, the fewer frames will be rendered.

**Initial Frame**

- It sets which frame is the first one to be rendered, from the whole benchmark run. For example, if the time step is 0.1 seconds and you set the initial frame to 10, the rendering starts at 1 second of the animation.

**Frame Count**

- It sets how many frames should be rendered, starting from the initial frame.

**Use Half Resolution Textures**

- Since textures require a big part of the memory allocation, using them at half resolution decreases memory usage. This may improve performance or allow one to run a benchmark on a device that does not have enough memory to support full texture resolution.

### **Ray Tracing**

- This parameter can be used to disable the usage of all ray tracing effects for testing purposes. The effects are not replaced by anything and thus the rendering will not look complete.

### **Ray Traced Reflections Resolution**

- This option modifies the resolution used for rendering the reflections. It can be 50%, 60%, 75%, or 100% of the main rendering resolution.

### **Max Bounce Count for Ray Traced Reflections\***

- This option controls how many reflection bounces are allowed before terminating the reflection path. The Official benchmarks of the cross-platform workload use a maximum of 1 bounce, while in Ultra maximum of 2 is used. The use of 2 instead of 1 allows reflections of reflections but decreases performance.

### **Ray Traced Reflections Shader Type**

- This option defines whether the ray queries are instantiated from a compute or a fragment shader

### **Reflections Workgroup Size**

- It defines the compute workgroup size (WGS) used by the Ray Traced Reflections effect when the selected “Rayquery Shader Type” is “compute”. Supported values are 8x4, 8x8, or automatic. When automatic is selected, both 8x4 and 8x8 are briefly tested before the benchmark starts and the one with better performance is selected.

### **Warmup Frame Count**

- To achieve consistency in the benchmarked frames, the user can define a specific number of frames that are rendered before the benchmarking starts. Due to device driver behavior, the first few rendered frames are often heavier than the following frames, regardless of their content, and thus the user might want to smoothen the result by rendering some warmup frames. In the Official benchmarks, we use 10 warmup frames before starting the benchmark.

*\*Available only on desktop applications*

## Experience

- The experience mode allows the user to run the benchmark content using time-based rendering. Although frame-based rendering is the best approach when benchmarking, Breaking Limit also features this time-based mode to visualize the content in a game-like smooth way. By using time-based rendering as many frames per second are rendered as possible, but the animation speed is kept constant. The time-based approach offers a more visually pleasant experience, and it is what actual games use. The disadvantage of this approach is that it cannot produce precise and comparable performance results. Therefore, the Experience mode does not produce a score and thus does not upload to Powerboard.

## Technical Requirements

Breaking Limit is a graphics benchmark that utilizes ray tracing. Thus, a ray tracing-capable GPU is required.

For the **cross-platform workload, Breaking Limit**, the minimum requirements are:

Android devices:

- Android 14+
- 4GB+ of shared memory
- Ray tracing-capable GPU and driver

iOS devices:

- iOS 17+
- 6GB+ of shared memory

Windows devices:

- x64 CPU
- Windows 10 or 11
- 8GB+ of RAM
- Ray tracing-capable GPU with 4GB+ of VRAM



Linux devices:

- x64 CPU
- Ubuntu 20.04 or 22.04 (other distributions may be supported, but are not tested)
- 8GB+ of RAM
- Ray tracing-capable GPU with 4GB+ of VRAM

For **Breaking Limit Ultra**, the minimum requirements are:

Windows devices:

- x64 CPU
- Windows 10 or 11
- 8GB of RAM
- Ray tracing-capable GPU with 8GB+ of VRAM

Linux devices:

- x64 CPU
- Ubuntu 20.04 or 22.04 (other distributions may be supported, but are not tested)
- 8GB+ of RAM
- Ray tracing-capable GPU with 8GB+ of VRAM

## Final Words

Breaking Limit opens a new perspective on graphics benchmarking, by extending what can be benchmarked with a sufficient true cross-platform workload. It is a daring move to extend this approach to real-time ray tracing. Although mobile devices are expected to be less performant than laptop or desktop hardware, face-to-face comparisons still provide useful insight into the differences between the platforms. By having the same singular reference, these comparisons are now possible.

Also, by including the Breaking Limit Ultra workload, for high-end desktops, GPUScore: Breaking Limit becomes a new reference for high-end real-time graphics benchmarking. It enables the user to explore and evaluate the latest and most innovative hardware capabilities of the latest high-end graphics hardware.

