



disguise Technologies
Head Office
88-89 Blackfriars Road
Southwark, London, SE1 8HA
United Kingdom
info@disguise.one
www.disguise.one

Introduction and Disclaimer

This document outlines the method taken to discover and monitor the video playback performance of each disguise machine. The configuration used to do so is extremely specific, with no additional pieces of product functionality active, making it closer to laboratory conditions than a real production. These data, therefore, are theoretical maximums and **are not** a guarantee of performance in any given scenario.

Goals

The goal of this test is to discover the amount of video layers of a single codec, resolution, and frame rate that can be played while maintaining a target output frame rate of 60 frames per second. The test and the content used is identical for each machine in the disguise catalog to allow for direct comparison. Since different codecs stress the video playback pipeline in different ways, these tests make it possible to determine potential strengths and weaknesses of each machine.



disguise Technologies
Head Office
88-89 Blackfriars Road
Southwark, London, SE1 8HA
United Kingdom
info@disguise.one
www.disguise.one

Project Setup

The test is designed to give the disguise software the best possible environment to play back the most demanding content possible. A set of rules is in place when setting up this test to ensure this is the case, and to keep the test as repeatable as possible.

The test project abides to these rules, which can be divided into categories of stage, timeline, feed, and hardware:

Stage:

- The stage of the project will have no meshes aside from the screens used to map and output the videos.
- One LED screen per resolution tested (Full-HD and DCI 4k).

Timeline:

- One codec and frame rate tested per track.
- One unique video file per layer.
- Video layers in each track arranged in a stair-step pattern, adding one layer per section break.
- Section breaks are placed every 30 seconds along the track.
- Cue tags and notes at the start of each section to track how many video layers are active in that section at-a-glance, and to control the playhead during the test.

Feed:

- All screens must be assigned to every output of the test machine.

Hardware:

- The test machine must be run as an actor with its GUI disabled.
- If the test machine cannot be run as an actor, a minimum of GUI elements should be present, and measurements are to be taken while viewing the stage.



disguise Technologies
Head Office
88-89 Blackfriars Road
Southwark, London, SE1 8HA
United Kingdom
info@disguise.one
www.disguise.one

- The test machine must never also be running the automated performance script (outlined below).
- The session director machine may run the automated performance script or the script may be run on a third-party on the network.
- All test machine outputs must be synchronized if supported.
- The GUI output of the test machine must be connected to a real display.
- All other outputs on the test machine must be at least emulated.
- If the test machine features a RAID array, it must be in its as-sold configuration.
- All machines must be using the unaltered latest version of their respective OS images.



disguise Technologies
Head Office
88-89 Blackfriars Road
Southwark, London, SE1 8HA
United Kingdom
info@disguise.one
www.disguise.one

Videos

The following codecs are tested in Full-HD and DCI 4k, at 30fps and 60fps:

- HAP
- HAP-Q
- Quicktime Animation
- TIFF image sequence
- NotchLC

10-bit capable systems are also tested with the following more demanding image sequence formats in 4K DCI at 30fps and 60fps:

- DPX, 8-bit, RGB
- DPX, 10-bit, RGB
- TGA, 8-bit, RGB
- TGA, 8-bit, RGBA
- TIFF, 8-bit, RGB
- TIFF, 8-bit, RGBA

Other supported codecs were not included for the following reasons:

DXV: Similar, if sometimes slightly lighter, performance impact to HAP. Less popular than HAP due to lack of support for version 3.

HAP-Alpha: Identical performance impact and data rate to HAP-Q, since both make use of a fourth channel.

Photo-JPEG: Deprecated.

The same videos are used for every test. There is one video per codec and resolution that is duplicated to fulfill the needs of the test. Video frame rate is controlled during the test using the disguise software. This is done for the sake of practicality, as the full array of test content already outstrips the media drive capacity of everything but the flagship system.

The content is chosen to fit the following requirements:

- Able to enforce a continuous bitrate for variable bitrate codecs.
- Plenty of movement in the frame for manual analysis scenarios.
- No embedded audio.
- At least 1 minute in length.
- Maximum or otherwise “final” quality render – eg: spatial quality at 100 for HAP.
- Full-HD content rendered at 1920*1080@30 and DCI content rendered at 4096*2160@30.



Above: A frame from the chosen content. The sweeping motions of the footage make it easy to spot frame drops during manual investigations. A small amount of animated colour noise has been added to the video to enforce a constant bitrate.



disguise Technologies
Head Office
88-89 Blackfriars Road
Southwark, London, SE1 8HA
United Kingdom
info@disguise.one
www.disguise.one

How is the performance data gathered?

Data is gathered using the internal automated testing library.

The performance script, as it is called, is designed to interface with a companion d3 project that matches the specification described above. It is able to cycle through a list of test cases (one case is one codec at a specific resolution and frame rate) which are divided into separate tracks in the d3 project. The script executes the following routine for every case/track:

1. Begin at section 1, which has a single active layer, and commence playback.
2. After 20 seconds, obtain the buffered performance data for the actor from the machine monitoring system. This data is exported from the director machine, so the action of downloading it does not disturb the actor.
3. Check against four different data points (outlined below) to determine if the actor was stable during the previous 20 seconds of recorded data.
4. If the actor was stable, move the playhead forward to the next section (adding another video layer in the process) and begin the cycle again.
5. As the layer count increases, the actor will begin to reach its performance threshold.
6. If a frame drop is detected on any one of the four data points, the actor is flagged as unstable and the script re-runs the test at that section, waiting a further 20 seconds to buffer a new set of performance data for analysis.
7. One of two things can happen here:
 - a. If the actor is found to be stable during this “double check” routine, another layer is added like normal and the test case continues.
 - b. If the second set of data also shows evidence of a frame drop, the current layer count minus one is recorded and the test moves onto section 1 of the next track/case in the list.

The script will run through the list of tracks/cases multiple times. Published figures will be obtained from the average of no fewer than 10 such iterations.

Multiple actors can be tested at once, but it is recommended that this only be done for multiples of the same machine type.

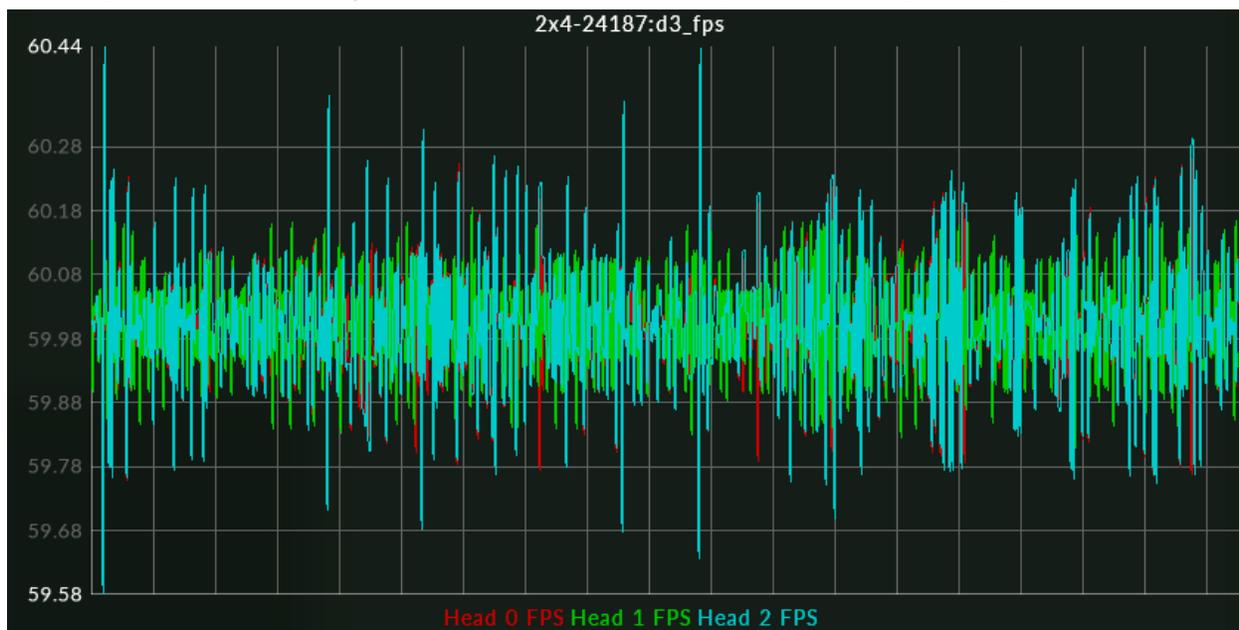
Four data points are used to assess the playback performance of each case as different codecs run into different bottlenecks on the system. GPU-bound codecs will cause the “d3_fps” monitor



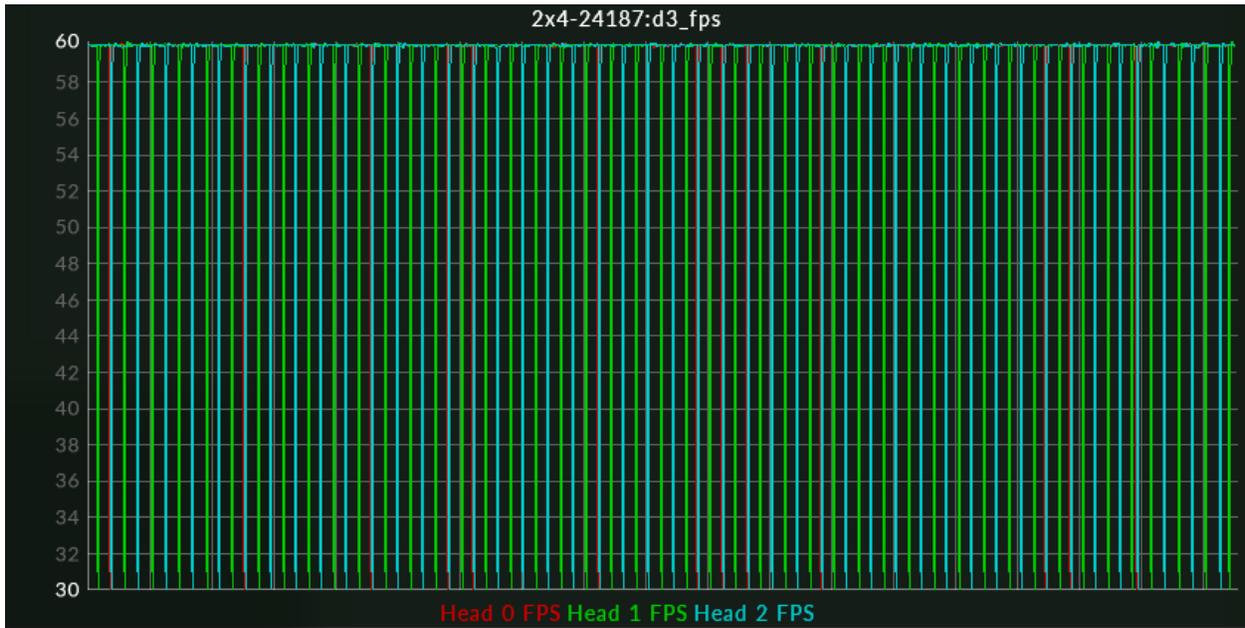
disguise Technologies
Head Office
88-89 Blackfriars Road
Southwark, London, SE1 8HA
United Kingdom
info@disguise.one
www.disguise.one

to drop, while CPU-bound and harddrive-bound codecs will show drops on the prefetcher monitors.

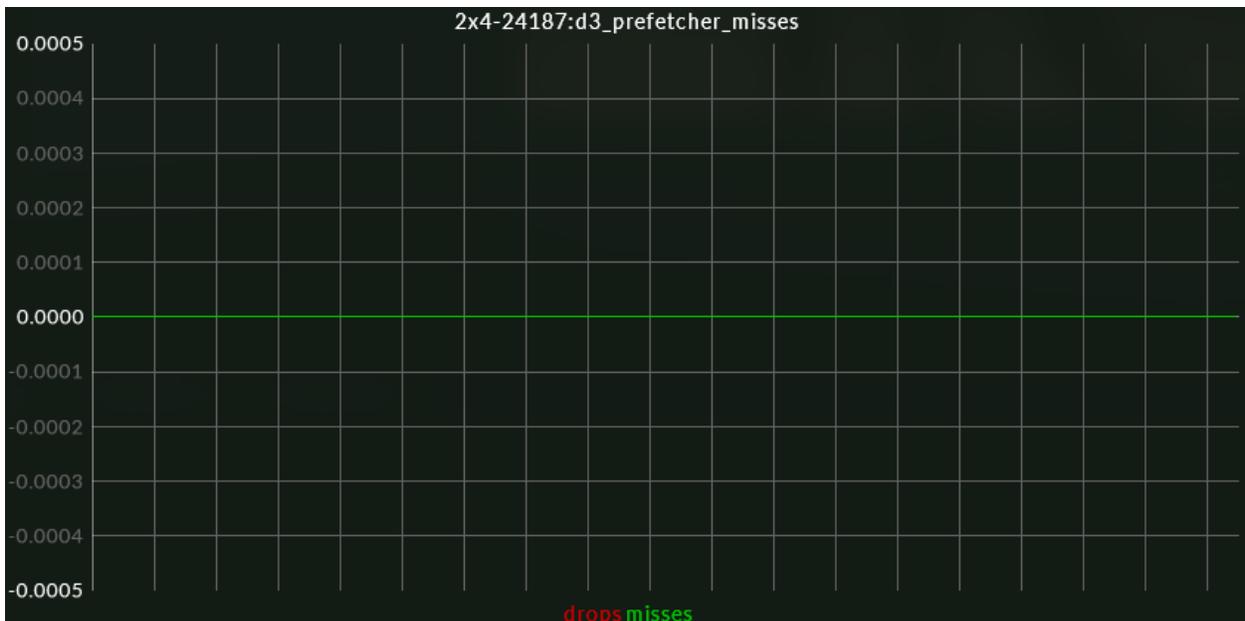
The script measures average FPS, instantaneous FPS, prefetcher drops and prefetcher misses.



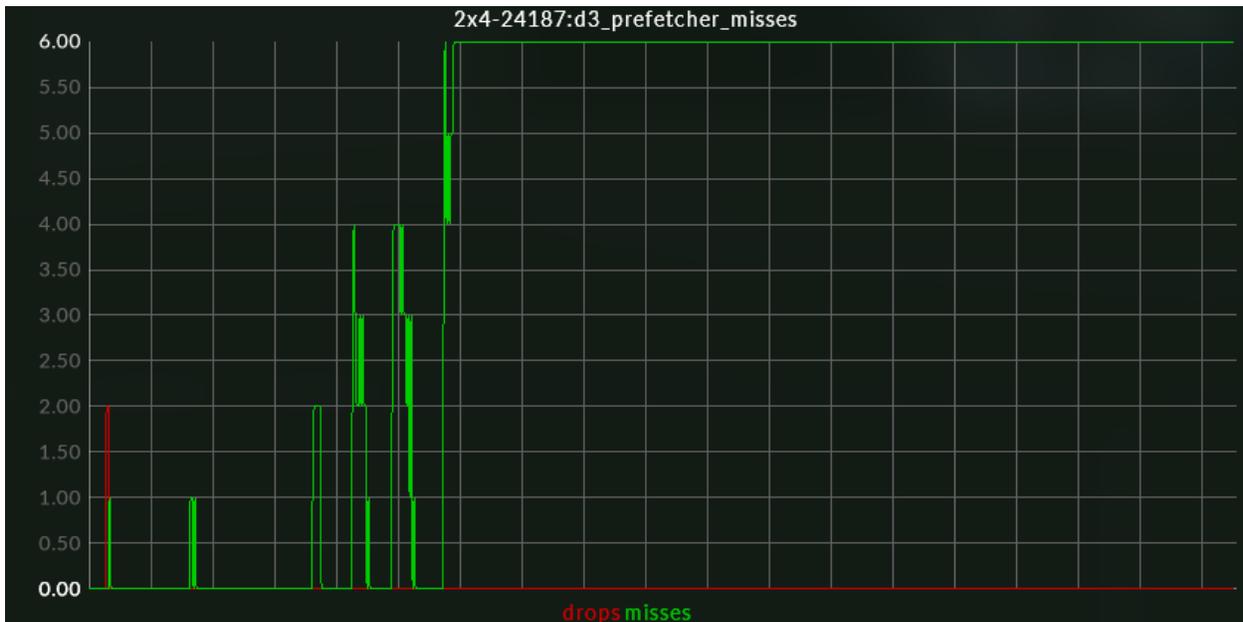
Good FPS: Y-axis shows all three heads are at values within a frame of 60fps. Some variation within this range is expected and attributed to minute differences in timing between the renderer and the metric gathering system. Head 0 is the GUI head, and is ignored by the script.



Bad FPS: Consistent spikes down to 30fps on head 1 and 2. In this case the average and instantaneous FPS figures would flag the machine as unstable.



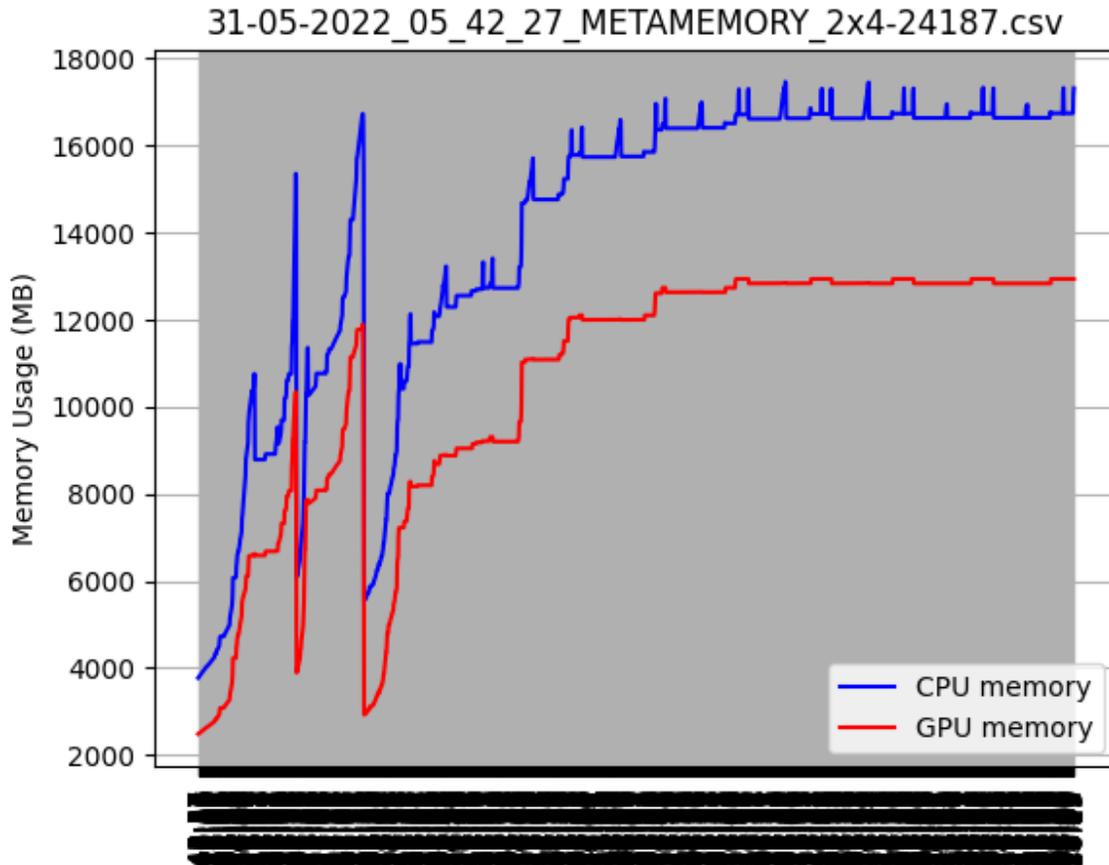
Good prefetcher drops and misses: A flat line at 0. Every frame prefetched is getting rendered.



Bad prefetcher drops and misses: If the value for either statistic ever goes above 0, the test is marked as a failure.

CPU, GPU, memory usage, and disk usage are also logged for posterity.

The script will also chart the CPU and GPU memory usage for the entire test run (across all iterations) for internal review.



An example of the “metamemory” chart.



disguise Technologies
Head Office
88-89 Blackfriars Road
Southwark, London, SE1 8HA
United Kingdom
info@disguise.one
www.disguise.one

Observations from testing

- The Solo and previously 2x2plus have extremely predictable performance despite the complication of not being able to run as actors. No matter the codec tested, the bottleneck is always at the media drive. As long as the data rate of the video load is below 660MB/s (the bandwidth of a SATA3 SSD), the machine will be able to run it.
- Older dual-CPU systems can have an advantage over newer single-CPU systems for formats like image sequences, but the true, practical, advantage comes from PCI bandwidth gains on newer platforms.
- Machines with excellent Notchmark scores might perform worse overall than machines with lower Notchmark scores. The video playback pipeline has a lot of moving parts aside from the GPU, and Notchmark is a largely GPU-based score.
- Older machines would exhibit performance degradation from the use of Quad-type VFC cards. This is not the case with newer machines, such as those in the vx range.