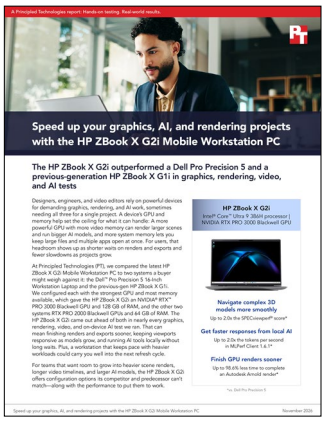

The science behind the report:

Speed up your graphics, AI, and rendering projects with the HP ZBook X G2i Mobile Workstation PC

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report Speed up your graphics, AI, and rendering projects with the HP ZBook X G2i Mobile Workstation PC.

We concluded our hands-on testing on September 22, 2026. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on July 19, 2026 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to https://facts.pt/calculating-and-highlighting-wins. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: Results of our testing.

Performance & productivity tests	HP ZBook X G2i	HP ZBook X G1i	Dell™ Pro Precision 5
Geekbench 7 Pro v7.0.0			
GPU OpenCL score	106,026	88,896	93,187
Background Blur	80,005	66,564	68,414
Face Tracking	103,698	82,723	81,734
Super Resolution	138,250	117,323	115,505
Horizon Detection	135,875	91,097	100,988
Photo Filter	122,138	97,230	107,858
Video Filter	107,159	87,784	91,412
RAW	99,052	92,650	105,086
Feature Matching	117,668	84,725	95,223
Path Tracer	103,610	96,687	106,100
Particle Physics	93,057	90,786	95,729
Fluid Simulation	83,044	79,159	70,403

Performance & productivity tests	HP ZBook X G2i	HP ZBook X G1i	Dell™ Pro Precision 5
3DMark® v2.32.8874			
Steel Nomad overall score	2,902	1,922	2,291
FPS	29.02	19.23	22.92
Speed Way overall score	3,277	2,261	2,549
Graphics (FPS)	32.77	22.61	25.49
Port Royal overall score	8,109	5,673	6,247
FPS	37.55	26.26	28.92
Procyon AI Computer Vision 2.0 Benchmark - GPU (Procyon v2.12.2641, Procyon AI Computer Vision 2.0 v1.0.567)	Windows ML - NVIDIA RTX PRO 3000 GPU, float16	Windows ML - NVIDIA RTX PRO 2000 GPU, float16	Windows ML - NVIDIA RTX PRO 2000 GPU, float16
Overall score	2,625	2,121	2,262
ConvNeXt throughput (images/s)	1,583.740	1,164.790	1,257.240
BLIP - Decoder Inference speed (it/s)	235.668	260.977	335.911
BLIP - Encoder Inference speed (it/s)	170.785	143.347	157.466
SAM 2.1 - Decoder Inference speed (it/s)	502.228	625.431	785.772
SAM 2.1 - Encoder Inference speed (it/s)	57.201	43.708	45.769
DETR - Inference speed (it/s)	104.729	87.710	88.945
ESRGAN - Inference speed (it/s)	11.711	8.814	9.053
Procyon AI Image Generation Benchmark - Stable Diffusion 1.5 (FP16) (benchmark v1.1.265)	NVIDIA TensorRT	NVIDIA TensorRT	NVIDIA TensorRT
Overall score	1,729	1,157	1,243
Overall duration (sec)	57.824	86.421	80.403
Overall image generation speed (sec/image)	3.614	5.401	5.025
Procyon AI Image Generation Benchmark - Stable Diffusion 1.5 (INT8)	NVIDIA TensorRT	NVIDIA TensorRT	NVIDIA TensorRT
Overall score	22,009	14,666	16,425
Overall duration (sec)	11.359	17.046	15.220
Overall image generation speed (sec/image)	1.420	2.131	1.903
Average UNET speed (it/s)	36.697	24.452	27.301
Procyon AI Image Generation Benchmark - Stable Diffusion XL (FP16)	NVIDIA TensorRT	NVIDIA TensorRT	NVIDIA TensorRT
Overall score	1,481	904	927
Overall duration (sec)	404.867	663.448	646.695
Overall image generation speed (sec/image)	25.304	41.466	40.418

Performance & productivity tests	**HP ZBook X G2i**	**HP ZBook X G1i**	**Dell™ Pro Precision 5**
Geekbench AI GPU v1.7.0 (ONNX framework)	DirectML - NVIDIA RTX PRO 3000 Blackwell	DirectML - NVIDIA RTX PRO 2000 Blackwell	DirectML - NVIDIA RTX PRO 2000 Blackwell
Single precision score	18,334	15,727	14,975
Half precision score	33,066	32,199	26,668
Quantized score	14,531	12,634	11,596
Image Classification (SP)	6,229	6,915	5,055
Image Classification (HP)	12,074	13,014	8,170
Image Classification (Q)	6,046	6,545	4,839
Image Segmentation (SP)	14,843	11,859	11,875
Image Segmentation (HP)	23,185	19,570	18,311
Image Segmentation (Q)	13,706	10,971	10,827
Pose Estimation (SP)	93,235	69,708	74,025
Pose Estimation (HP)	298,561	226,779	221,991
Pose Estimation (Q)	89,954	64,120	67,524
Object Detection (SP)	9,225	8,055	7,061
Object Detection (HP)	13,806	16,768	11,675
Object Detection (Q)	7,944	7,739	6,382
Face Detection (SP)	25,073	20,689	20,871
Face Detection (HP)	38,757	39,459	33,563
Face Detection (Q)	22,556	19,173	18,586
Depth Estimation (SP)	44,086	34,150	37,113
Depth Estimation (HP)	92,162	88,925	74,958
Depth Estimation (Q)	37,772	29,258	31,170
Style Transfer (SP)	156,689	105,752	114,190
Style Transfer (HP)	447,542	368,622	373,837
Style Transfer (Q)	133,361	92,671	101,313
Image Super-Resolution (SP)	22,751	19,165	18,536
Image Super-Resolution (HP)	42,866	46,546	34,768
Image Super-Resolution (Q)	20,190	16,057	15,561
Text Classification (SP)	3,169	2,866	2,756
Text Classification (HP)	3,819	3,755	3,297
Text Classification (Q)	1,728	1,849	1,377
Machine Translation (SP)	4,522	4,818	3,835
Machine Translation (HP)	4,827	5,290	4,364
Machine Translation (Q)	1,828	1,898	1,634

Performance & productivity tests	HP ZBook X G2i	HP ZBook X G1i	Dell™ Pro Precision 5
Amuse v3.7.3			
Time to generate image (sec) (lower is better)	92.70	135.15	123.42
Time to generate video (sec) (lower is better)	493.10	665.70	623.23
MLPerf Client v1.6.1 (Llama 3.1 8B Instruct)	NVIDIA Windows ML NvTensorRtRtx GPU	NVIDIA Windows ML NvTensorRtRtx GPU	NVIDIA Windows ML NvTensorRtRtx GPU
Overall time to first token (TTFT) (sec) (lower is better)	0.43	1.07	1.40
Overall tokens per second (TPS)	56.4	29.1	27.1
Topaz Video v1.7.0			
3,840 x 2,160 (4K) Artemis 1X (FPS)	2.64	1.63	2.23
3,840 x 2,160 (4K) Artemis 2X (FPS)	2.00	1.20	1.56
3,840 x 2,160 (4K) Artemis 4X (FPS)	0.75	0.45	0.53
3,840 x 2,160 (4K) Iris 1X (FPS)	3.33	2.01	2.82
3,840 x 2,160 (4K) Iris 2X (FPS)	1.94	1.12	1.66
3,840 x 2,160 (4K) Iris 4X (FPS)	0.65	0.33	0.39
3,840 x 2,160 (4K) Proteus 1X (FPS)	2.76	1.61	2.37
3,840 x 2,160 (4K) Proteus 2X (FPS)	2.03	1.21	1.97
3,840 x 2,160 (4K) Proteus 4X (FPS)	0.81	0.46	0.59
3,840 x 2,160 (4K) Proteus Natural 2X (FPS)	0.89	0.48	0.74
3,840 x 2,160 (4K) Gaia 1X (FPS)	0.98	0.53	0.75
3,840 x 2,160 (4K) Gaia 2X (FPS)	0.68	0.37	0.50
3,840 x 2,160 (4K) Gaia 4X (FPS)	0.49	0.27	0.38
3,840 x 2,160 (4K) Gaia Animation 2X (FPS)	2.87	1.76	2.13
3,840 x 2,160 (4K) Nyx 1X (FPS)	1.06	0.57	0.84
3,840 x 2,160 (4K) Nyx 2X (FPS)	1.43	0.79	1.11
3,840 x 2,160 (4K) Nyx Fast 1X (FPS)	2.31	1.26	1.77
3,840 x 2,160 (4K) Nyx XL 1X (FPS)	0.23	0.14	0.18
3,840 x 2,160 (4K) Rhea 4X (FPS)	0.48	0.27	0.38
3,840 x 2,160 (4K) RXL 4X (FPS)	0.49	0.27	0.37
3,840 x 2,160 (4K) Hyperion HDR 1X (FPS)	12.33	9.57	10.31
3,840 x 2,160 (4K) 4X Slowmo Apollo (FPS)	4.18	2.29	3.39
3,840 x 2,160 (4K) 4X Slowmo APFast (FPS)	15.48	9.03	10.73
3,840 x 2,160 (4K) 4X Slowmo Chronos (FPS)	2.04	1.21	1.69
3,840 x 2,160 (4K) 4X Slowmo CHFast (FPS)	4.08	2.29	3.21
3,840 x 2,160 (4K) 16X Slowmo Aion (FPS)	2.42	1.63	1.88

Performance & productivity tests	HP ZBook X G2i	HP ZBook X G1i	Dell™ Pro Precision 5
Blender Benchmark using GPU (Blender Benchmark Launcher v3.3.0, Blender v5.2.0)			
Blender Open Data - Monster	1,716.67	1,165.91	1,227.53
Blender Open Data - Junkshop	1,361.16	942.50	1,024.84
Blender Open Data - Classroom	1,087.73	771.61	837.16
Maxon Redshift Benchmark (Maxon One App v2026.5, Redshift v2026.8.1)			
Time to complete (m:ss) (lower is better)	02:31	03:23	03:07
Puget Bench for Creators: Adobe® Premiere® Pro Extended (Puget Bench v2.0.1, Adobe Premiere Pro v26.0)			
Overall score (Standard)	91,914	75,191	78,064
Overall score (Extended)	73,202	59,438	61,769
LongGOP score (Standard)	86	75	76
Intraframe score (Standard)	98	87	94
RAW score (Standard)	149	128	126
GPU Effects score (Standard)	57	38	42
LongGOP score (Extended)	64	58	55
Intraframe score (Extended)	53	45	50
RAW score (Extended)	129	108	107
GPU Effects score (Extended)	65.70	44.40	49.60
Encoding - H.264 50Mbps 8-bit UHD (FPS)	158.80	137.54	151.31
Encoding - H.264 10Mbps 8-bit HD (FPS)	336.90	462.83	510.18
Encoding - DNxHR LB UHD (FPS)	60.45	56.10	63.12
Encoding - DNxHR SQ UHD (FPS)	60.34	57.29	63.78
Encoding - DNxHR HQX UHD (FPS)	38.81	35.43	39.59
Encoding - ProRes 422 Proxy UHD (FPS)	222.21	172.42	188.58
Encoding - ProRes 422HQ UHD (FPS)	142.71	112.60	139.48
Encoding - ProRes 4444 UHD (FPS)	88.61	81.35	88.31
Processing - 4K H.264 150Mbps 4:2:0 8-bit (FPS)	135.51	152.52	135.18
Processing - 4K H.264 150Mbps 4:2:0 8-bit 4x MultiStream (FPS)	57.88	45.20	41.29
Processing - 4K HEVC 100Mbps 4:2:2 10-bit (FPS)	101.76	93.15	93.63
Processing - 4K HEVC 100Mbps 4:2:2 10-bit 4x MultiStream (FPS)	32.38	34.00	30.70
Processing - 8K HEVC 100Mbps 4:2:0 8-bit (FPS)	57.78	50.32	47.88
Processing - 4K ProRes 422 Proxy (FPS)	102.13	103.20	109.69
Processing - 4K ProRes 422 (FPS)	99.62	94.60	104.10
Processing - 4K ProRes 422 4x MultiStream (FPS)	32.44	27.28	24.90
Processing - 4K DNxHR LB (FPS)	210.72	189.32	208.83
Processing - 4K DNxHR SQ (FPS)	217.05	197.56	192.29
Processing - Canon Cinema RAW Light (FPS)	77.52	68.48	65.33

Performance & productivity tests	HP ZBook X G2i	HP ZBook X G1i	Dell™ Pro Precision 5
Processing - ARRIRAW (FPS)	176.74	132.76	121.46
Processing - 4K Nikon N-RAW (FPS)	85.00	66.59	68.62
Processing - Sony X-OCN ST (FPS)	19.70	18.99	21.03
Processing - RED R3D (FPS)	63.73	53.31	53.84
Processing - RED R3D 4x MultiStream (FPS)	15.44	12.25	11.76
Real World - YouTube Gamer (FPS)	56.78	43.12	43.93
Real World - Behind the Scenes (FPS)	24.46	18.75	15.90
Real World - Behind the Scenes (ProRes) (FPS)	21.57	16.66	18.83
Processing - 8K RED (FPS)	50.45	41.42	39.55
GPU Effects - VR Digital Glitch x20 (FPS)	14.72	10.01	11.67
GPU Effects - VR De-Noise x4 (FPS)	6.79	4.75	5.34
GPU Effects - Echo Glow x3 (FPS)	9.03	6.21	6.90
GPU Effects - Bokeh Blur x2 (FPS)	5.24	3.31	3.53
GPU Effects - Edge Glow x3 (FPS)	9.10	6.15	6.75
GPU Effects - Wonder Glow x3 (FPS)	9.66	6.65	7.50
Autodesk® Arnold v7.5.0.0			
GPU render time - hr:min:sec (lower is better)	0:06:10	3:06:58	7:53:05
Chaos V-Ray GPU Benchmark v6.00.02 - GPU + CUDA			
Overall score	2,945	2,094	2,304
Chaos V-Ray GPU Benchmark - GPU + RTX			
Overall score	4,543	3,206	3,542
OctaneBench® 2025 (benchmark v2025.2.1)			
Overall score	459.01	313.59	342.90
SPECviewperf® 15.0.1 - 1920x1080 (FHD)			
3ds Max® composite score	59.08	31.04	28.66
Blender composite score	72.36	49.21	55.06
Catia composite score	71.62	52.52	57.31
Creo composite score	168.46	132.62	146.90
Energy composite score	87.05	58.32	63.87
Enscape composite score	39.78	28.37	31.73
Maya® composite score	176.26	126.12	151.36
Medical composite score	142.04	97.10	114.12
SNX composite score	91.15	99.84	104.84
Solidworks composite score	89.02	61.44	73.46
Unreal Engine composite score	69.61	55.39	63.18

System configuration information

Table 2: Detailed information on the systems we tested.

System	HP ZBook X G2i 16"	Dell Pro Precision 5 16"	HP ZBook X G1i 16"
Processor			
Vendor	Intel®	Intel	Intel
Model number	Core™ Ultra 9 386H	Core Ultra 9 386H	Core Ultra 9 285H
Core frequency (GHz)	1.6 – 4.9	1.6 – 4.9	1.0 – 5.4
Number of cores	(16) 4 x Performance Cores 8 x Efficient Cores 4 x Low Power Efficient Cores	(16) 4 x Performance Cores 8 x Efficient Cores 4 x Low Power Efficient Cores	(16) 6 x Performance Cores 8 x Efficient Cores 2 x Low Power Efficient Cores
Number of threads	16	16	16
L3 cache (MB)	18	18	24
AI engine capabilities NPU TOPS	50	50	13
NPU driver	Intel 32.0.100.4778	Intel 32.0.100.4778	Intel 32.0.100.4724
Memory			
Amount (GB)	128	64	64
Type	DDR5	LPDDR5X	DDR5
Speed (MT/s)	6,400	8,533	5,600
Graphics			
Vendor	NVIDIA	NVIDIA	NVIDIA
Model number	RTX PRO 3000 Blackwell	RTX PRO 2000 Blackwell	RTX PRO 2000 Blackwell
Driver	NVIDIA 32.0.15.9608	NVIDIA 32.0.15.9671	NVIDIA 32.0.15.9595
Storage			
Model	Samsung® MZVL81T0HFLB-00BH1	SK Hynix® PVC10	Samsung MZVL21T0HCLR-00BH1
Amount (TB)	1	1	1
Type	PCIe Gen4 NVMe M.2 SSD	PCIe Gen4 NVMe M.2 SSD	PCIe Gen4 NVMe M.2 SSD
Connectivity/expansion			
WLAN	Intel Wi-Fi 7 BE211	Intel Wi-Fi 7 BE211	Intel Wi-Fi 6E AX211
Bluetooth	6.0	6.0	5.3
USB	2 x Thunderbolt 4 USB Type-C 40Gbps 2 x USB Type-A 5Gbps	2 x Thunderbolt 4 USB Type-C 40Gbps 2 x USB Type-A 5Gbps	2 x Thunderbolt 4 USB Type-C 40Gbps 2 x USB Type-A 5Gbps
Video	HDMI 2.1	HDMI 2.1	HDMI 2.1
Battery/AC adapter			
Type	Lithium-ion	Lithium-ion	Lithium-ion
Rated capacity (Wh)	96	96	83

System	HP ZBook X G2i 16”	Dell Pro Precision 5 16”	HP ZBook X G1i 16”
Display			
Type	WUXGA 400 nit	WUXGA 400 nit	WUXGA 400 nit
Size (inches)	16	16	16
Resolution	1,920 x 1,200	1,920 x 1,200	1,920 x 1,200
Touchscreen	No	No	No
Operating system			
Vendor	Microsoft	Microsoft	Microsoft
Name	Windows 11 Pro	Windows 11 Pro	Windows 11 Pro
Build number or version	25H2 Build 26200.9168	25H2 Build 26200.9168	25H2 Build 26200.9168
BIOS			
BIOS name and version	HP Y97 Ver.01.02.01 (07/07/2026)	Dell 1.5.2 (07/22/2026)	HP X97 Ver.01.06.05 (07/22/2026)
Dimensions			
Height (inches)	0.98	0.54 – 0.72	0.9
Width (inches)	14.15	14.09	14.15
Depth (inches)	9.88	10.08	9.88
Weight (pounds)	5.02	5.10	4.86

How we tested

Setting up the systems

Setting up and updating the OEM image

1. Boot the system.
2. Follow the on-screen instructions to complete installation, using the default selections when appropriate.
3. Set the Windows (plugged in) Power Mode to Best Performance.
4. Set Screen and Sleep options to Never:
    a. Right-click the desktop, and select Display settings.
    b. From the left column, select System.
    c. Click Power & Battery.
    d. For all power options listed under Screen and Sleep, select Never.
5. Disable User Account Control notifications:
    a. Select Windows Start, type UAC, and press the Enter key.
    b. Move the slider control to Never notify, and click OK.
6. Run Windows Update, and install all available updates.
7. Run the Support Assistant utilities, and install all recommended BIOS and driver updates available.
8. Verify the date and time are correct, and synchronize the system clock with the time server.
9. Pause Automatic Windows Updates:
    a. Click the Windows Start button.
    b. Type Windows Update settings and press Enter.
    c. From the Pause updates drop-down menu, select Pause for 5 weeks.

Capturing an image

1. Connect an external HDD to the system.
2. Click the Windows Menu button, and type `Control Panel` in the search bar. Click Control Panel → System and Security → Backup and Restore (Windows 7) → Create a system image.
3. Verify that the external HDD is selected as the save drive, and click Next.
4. Verify that all drives are selected to back up, and click Next.
5. Click Start backup.
6. When you see the prompt to create a system repair disc, select No, and close the dialogs.

Restoring an image

1. Connect an external HDD to the system.
2. Press and hold the Shift key while restarting the system.
3. Select Troubleshoot.
4. Select Advanced options.
5. Select See more recovery options.
6. Select System image recovery.
7. Select the User account.
8. Enter the system password, and click Continue.
9. At the Restore system files and settings screen, select Next.
10. Verify that the external HDD is selected, and click Next.
11. Once the recovery has completed, click Finish.

Testing with 3DMark

Setting up the test

1. Purchase and download 3DMark from https://benchmarks.ul.com/3dmark.
2. To install 3DMark with the default options, double-click the 3DMark installer.exe file.
3. To launch 3DMark, double-click the 3DMark desktop icon.
4. Enter the registration code, and click Register.
5. Exit 3DMark.

Running the test

1. To launch the benchmark, double-click the 3DMark desktop icon.
2. At the 3DMark Home screen, click More Tests.
3. Select the benchmark you want to run.
4. Move the slider button to turn off the "Include Demo" feature.
5. Click Run.
6. When the benchmark run completes, record the results.
7. Perform steps 1 through 6 two more times for each benchmark, and report the median of the three runs.

Testing with Amuse Generate AI

Setting up the test

1. Download the Amuse application from https://github.com/saddam213/AmuseAI/releases/tag/v3.7.3.
2. Using all the defaults, run the installer, and install the application.
3. Launch the Amuse application.
4. Under the Image tab, click Load, and download the Anima model.
5. Click the Video tab, and download the AnyFlow 14B model.

Running the test

1. Launch Amuse.
2. Click Load to load the downloaded Anima model.
3. Once the model has loaded, select the Image tab, and click Load to load the downloaded model.
4. In the prompt type A cute raccoon playing guitar on the beach.
5. Under the Options tab verify the following settings:
   - Resolution = 768 x 768 Square.
   - Seed = 8.
   - Scheduler = Euler.
   - Steps = 30.
   - Guidance = 7.5.
6. Under the Automation tab, verify the following settings:
   - Automation = Seed Batch
   - Image Count = 4.
7. Click Automate to generate Images.
8. After the image generation has completed, record the elapsed time.
9. Wait 5 minutes and repeat step 7 two more times and capture the median of 3 runs.
10. Select the Video tab.
11. Click Load to load the AnyFlow 14B model.
12. Once the model has downloaded, select the Video tab, and click Load to load the downloaded model.
13. Click Create to create the Python Virtual Environment.
14. In the prompt type A cute raccoon playing guitar on the beach.

15. Under the Options tab verify the following settings:
    - Resolution = 720 x 480 Landscape.
    - Seed= 8.
    - Scheduler= FlowMatchEuler.
    - Frames= 33.
    - Frame Rate= 16 fps.
16. Under the Automation tab verify the following settings:
    - Automation = Seed Batch
    - Image Count = 4.
17. Click Automate to generate videos.
18. After the video generation has completed, record the elapsed time.
19. Wait 5 minutes and repeat step 17 two more times and capture the median of 3 runs.

Testing with Arnold Renderer

Setting up the test

1. Purchase and download Arnold from https://www.autodesk.com/.
2. Download and install 7-zip from https://www.7-zip.org/.
3. Go to https://dpel.aswf.io/alab/ and download the two test scenes to the Downloads directory:
    - Asset Structure - v.2.3.0 (aka Alab-2.3.0.zip)
    - Techvar Assets - v.2.2.0 (aka alab-techvars.v2.2.0.zip)
4. Launch PowerShell in Administrator mode, and enable long path support by typing:

```
New-ItemProperty -Path "HKLM:\SYSTEM\CurrentControlSet\Control\FileSystem" -Name "LongPathsEnabled"
-Value 1 -PropertyType DWORD -Force
```

5. Reboot the system.
6. Using 7-zip, extract the two test scenes using the following PowerShell script:

```
$ErrorActionPreference = 'Stop'
$7z    = "C:\Program Files\7-Zip\7z.exe"
$root  = "C:\alab"          # repo root
$usd   = "C:\alab\ALab"     # actual USD root
cd "$HOME\Downloads"

# Guard: Move-Item nests instead of renaming if the target already exists
if (Test-Path $root) { throw "$root already exists. Delete or rename it first." }

# 1. Asset Structure -> C:\alab
& $7z x ALab-2.3.0.zip -oC:\stage_alab
if ($LASTEXITCODE -gt 1) { throw "7-Zip failed on Asset Structure (exit $LASTEXITCODE)" }
Move-Item C:\stage_alab\ALab-2.3.0 $root
if (-not (Test-Path "$usd\entry.usda")) { throw "entry.usda not found at $usd" }
Remove-Item C:\stage_alab -Recurse -Force

# 2. Techvars -> staging folder (zip has a 'techvar_assets' wrapper)
& $7z x alab-techvars.v2.2.0.zip -oC:\stage_tv
if ($LASTEXITCODE -gt 1) { throw "7-Zip failed on Techvars (exit $LASTEXITCODE)" }
if (-not (Test-Path C:\stage_tv\techvar_assets\fragment)) { throw "techvar_assets\
fragment not found" }

# 3. Merge techvar fragment onto the structure fragment, overwriting
robocopy C:\stage_tv\techvar_assets\fragment "$usd\fragment" /E /IS /IT /MOVE /R:1 /W:1 /NFL /NDL
if ($LASTEXITCODE -ge 8) { throw "robocopy failed (exit $LASTEXITCODE)" }
Remove-Item C:\stage_tv -Recurse -Force

# 4. Verify
Test-Path "$usd\entry.usda"
"{0:N2} GB" -f ((Get-ChildItem "$usd\fragment" -Recurse -File | Measure-Object Length
-Sum).Sum / 1GB)
```

7. Install Arnold.
8. Extract the Arnold zip file.
9. Inside the Bin directory, click the ArnoldLicenseManager.exe.
10. Install a single-user license.
11. Install Autodesk Licensing software when prompted.
12. Log in with your Autodesk account.
13. Click the Windows Start button, and type Edit the system environment variables.
14. Under the Advanced tab, click Environment Variables.
15. Under the System variables section, scroll down and select Path, and click Edit.
16. Click New, browse to where the Arnold\bin directory is located, and click OK.

Running the GPU render test

1. Open PowerShell in Administrator mode, browse to the Arnold\bin path, and type:

```
kick C:\alab\ALab\entry.usda -frame 1010 -dw -dp -device gpu -set options.camera /root/camera01/
GEO/renderCam_hrc/renderCam_buffer/renderCam_srt/renderCam -set options.xres 1920 -set options.yres
1080 -set options.AA_samples 6 -set options.GI_diffuse_depth 2 -set options.GI_specular_depth 2 -set
options.GI_transmission_depth 2 -set options.GI_total_depth 10 -ii -o C:\alab\out\gpu_warmup.exr -v
5 -logfile C:\alab\out\gpu_warmup.txt
```

2. Record the render time reported in the log file.
3. Repeat steps 1 and 2 two more times.
4. Record the median result.

Testing with the Blender Benchmark

Setting up the test

1. Download the Blender Benchmark from https://opendata.blender.org/.

Running the test

1. Launch the Blender Benchmark.
2. At the Welcome screen, click Next.
3. Select Blender version 5.2.0, and click Next.
4. At the Benchmark Scenes screen, click Next.
5. At the Benchmark Device screen, select either the CPU or GPU option, and click Start Benchmark.
6. Record the results.
7. Wait 15 minutes before performing the next run.
8. Repeat steps 1 through 7 two more times, and record the median result.

Testing with Chaos V-Ray benchmark

Setting up the V-Ray benchmark

1. Download the V-Ray benchmark from https://www.chaos.com/benchmark-download.

Running the V-Ray benchmark

We captured the following metrics for this testing.

- V-Ray – GPU – RTX
- V-Ray – GPU – CUDA

1. Run V-Ray benchmark.
2. Using V-Ray for GPU select RTX, select only the GPU, and run v-Ray for 5 minutes. Record the results.
3. Wait 5 minutes, then repeat step 2 twice.
4. Using V-Ray for GPU select CUDA, select only the GPU, and run the test for 5 minutes. Record the results.
5. Wait 5 minutes, then repeat step 4 twice.

Testing with Geekbench 7 Pro GPU OpenCL

Setting up the test

1. Purchase a Pro license and download and install Geekbench 7 Pro from https://www.geekbench.com/download/.

Running the test

1. Launch Geekbench.
2. Select OpenCL and click Run Compute Benchmark.
3. Record the result.
4. Wait 5 minutes before re-running.
5. Repeat steps 1 through 4 two more times and report the median result.

Testing with the Geekbench AI benchmark

Setting up the test

1. Purchase and download a Geekbench AI Pro license from https://www.geekbench.com/ai/download/.
2. Using all the defaults, run the installer, and install the benchmark.

Running the test

1. Launch Geekbench AI.
2. Enter the license key.
3. For GPU testing, select:
   - AI Framework: ONNX
   - AI Backend: DirectML
   - AI Device: graphics card
4. Click Run AI Benchmark.
5. Wait 5 minutes and repeat steps 3 and 4 two more times and capture the median result.

Testing with MLPerf Client

Setting up the test

1. Download the latest MLPerf Client GUI (we used version 1.6.1) from https://github.com/mlcommons/mlperf_client/releases
2. Extract the MLPerf Client zip file to the desktop.
3. Install the VC++ Redistributable executable found inside the extracted MLPerf Client file.

Running the test

1. Inside the extracted ML Perf Client directory, to launch the benchmark, click mlperf-windows-gui.exe.
2. Accept the license agreement, and click Continue.
3. Select the benchmark (Llama 3.1 8B Instruct on NVIDIA Windows ML NvTensorRtRtx GPU), and click Run Benchmark Test.
4. To download the necessary benchmark files, click Yes.
5. Record the results.
6. Wait 15 minutes before rerunning.
7. Repeat steps 1 through 6 two more times, and report the median result.

Testing with Maxon Redshift

Setting up the test

1. Purchase, download, and install a 1-month license for Maxon Redshift from https://www.maxon.net/en/buy#monthly.
2. Launch the Maxon app, and install the Redshift license onto the system under test.

Running the test

1. To open a command prompt, click the Windows icon, type cmd and press Enter.
2. Inside the command prompt, type cd C:\Program Files\Maxon Redshift 2026\bin, and press Enter.
3. To start the benchmark, type RunBenchmark.bat, and press Enter.
4. Record the results.
5. Wait 5 minutes before rerunning the test twice, and report the median result.

Testing with OctaneBench 2025 Benchmark

Setting up the OctaneBench 2025 Benchmark

1. Download and extract the OctaneBench Benchmark from https://render.otoy.com/octanebench/.

Running the OctaneBench benchmark

1. Click the Octane.exe to launch the benchmark.
2. Click Run. Record the results.
3. Wait 5 minutes, then repeat step 2 twice, and report the median result.

Testing with Procyon AI Computer Vision 2.0 Benchmark

Setting up the test

1. Purchase and download the Procyon AI Computer Vision 2.0 benchmark from https://benchmarks.ul.com/procyon.
2. Double-click the installer.
3. Click Next.
4. To agree to the EULA, click Agree, and click Next.
5. Click Next.
6. Launch Procyon.
7. Select Settings, and input the license key.
8. Close Procyon.

*Note: It is necessary to have Windows Update enabled for the first run so that the benchmark can check online for the latest Windows ML version.

Running the test

1. Launch Procyon.
2. Select the Computer Vision 2.0 Benchmark test.
3. For GPU testing, select the GPU for Inference device type, and leave the default Inference Engine set to Windows ML.
4. To begin the test, click Run.
5. When the test completes, record the results.
6. Wait 15 minutes before rerunning.
7. Repeat steps 1 through 6 two more times, and report the median result.

Testing with Puget Bench for Creators: Premiere Pro

We used the following applications:

- Adobe Premiere Pro
- Puget Bench for Creators

Setting up the test

1. Launch Adobe Premiere Pro.
2. Click through the Tutorial pop-up tips.
3. Close Adobe Premiere Pro.
4. Purchase a Puget Bench for Creators license from https://www.pugetsystems.com/pugetbench/creators/.
5. Click the Download Puget Bench for Creators for Windows or Mac.
6. After the download completes, double-click the installation file to install Puget Bench.
7. Enter the license key in the license field. Click Activate.
8. Click Download Assets.

Running the test

1. Boot the system.
2. Open Puget Bench for Creators.
3. Select the Premiere Pro test on the left side of the app.
4. Under the Benchmark Preset drop-down menu, select Extended.
5. Click Start Test.
6. When the benchmark finishes, record the overall score.
7. Close Puget Bench for Creators, and restart the system under test.
8. Wait 30 minutes before performing the next run.
9. Repeat steps 1 through 8 two more times, and record the median results.

Testing with SPECviewperf 15.0.1

Setting up the test

1. Navigate to https://gwpg.spec.org/benchmarks/benchmark/specviewperf-15_0_1/, and purchase and download the vendor license of the benchmark.
2. Download and unzip SPECviewperf 15.0.1.
3. Navigate to the extracted SPECviewperf 15.0.1 directory, and click the .exe file to install.

Running the test

1. Launch SPECviewperf.
2. Click Run.
3. When the test is complete, record the results.
4. Repeat steps 1 through 3 two more times, and report the median result.

Testing with Topaz Video

Setting up the test

1. Purchase, download and install Topaz Labs Video from https://www.topazlabs.com/downloads.

Running the test in 3,840 x 2,160 (4K)

1. Launch Topaz Labs Video.
2. Enter license information, and click OK.
3. Select Process → Benchmark.
4. From the Input Resolution drop-down menu, select 3,840 x 2,160 (4K).
5. Click Benchmark.
6. When the test completes, record the results.
7. Wait 15 minutes before retesting.
8. Repeat steps 3 through 7 two more times.

Testing with the Procyon AI Image Generation Benchmark

Setting up the test

1. Purchase and download the Procyon AI Image Generation benchmark from https://benchmarks.ul.com/procyon.
2. To the Procyon benchmark, double-click the installer.
3. Click Next.
4. Click to agree to the EULA, and click Next.
5. Click Next.
6. Launch Procyon.
7. Select Settings, and input the license key.
8. Close Procyon.

Running the test

1. Launch Procyon.
2. Select the Image Generation Benchmark test.
3. Under the Stable Diffusion 1.5 (FP16) test option, select the appropriate AI Inference Engine and select the graphics device name to be used.
4. To begin the test, click Run.
5. Complete 3 runs, and report the median result.
6. Under the Stable Diffusion 1.5 (INT8) test option, select the appropriate AI Inference Engine and select the graphics device name to be used.
7. To begin the test, click Run.
8. Complete and record 3 runs.
9. Under the Stable Diffusion XL (FP16) test option, select the appropriate AI inference Engine and select the graphics device name to be used.
10. To begin the test, click Run.
11. Complete and record 3 runs and report the median.

This project was commissioned by HP

Read the report ▶

Primary contributors

Tech: Jon D.

Writing: Elyse K.

Design: Jared W.

PM: Peter H.

How we created this report

A PT team, which includes the contributors we've listed and others, created this report and performed the technical work behind it. We used AI to help outline, draft, and proofread this report, as well as to aid in chart creation, image editing, and content layout.

Principled Technologies®

Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the science behind this report.

DISCLAIMER OF WARRANTIES; LIMITATION OF LIABILITY:

Principled Technologies, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing, however, Principled Technologies, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree that Principled Technologies, Inc., its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result.

In no event shall Principled Technologies, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Principled Technologies, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Principled Technologies, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein.