

GPU Buying Guide 2026

Every Price Tier from \$200 to \$1,200 — Honest Performance Expectations, No Brand Bias

CONTENTS

1. How GPUs Are Tested — What Numbers Matter
2. The \$200–\$350 Tier — Budget 1080p
3. The \$400–\$600 Tier — 1440p Sweet Spot
4. The \$700–\$1,000 Tier — High-Refresh 4K
5. DLSS 4 vs FSR 4 — An Honest Comparison
6. VRAM Requirements by Resolution and Game Type
7. AMD vs NVIDIA Ecosystem Trade-offs
8. When to Buy and When to Wait

How GPUs Are Tested — What Numbers Matter

Graphics card marketing has perfected the art of presenting impressive-sounding numbers that mean very little in practice. Teraflops of compute performance, memory bandwidth figures, and peak clock speeds are all metrics that correlate only loosely with actual gaming frame rates. This guide uses a different set of numbers — ones that translate directly to the experience of sitting in front of your monitor.

The three metrics that matter most are average frame rate at your target resolution, 1% low frame rate at your target resolution, and VRAM usage under sustained load at your target resolution. Average frame rate tells you how smooth gameplay is most of the time. The 1% low is the metric that determines whether a game feels smooth or stuttery — it represents the worst frames in any one-second window and correlates directly with perceived smoothness more than the average does.

VRAM usage is the metric most often ignored by reviewers focused on day-one performance numbers. A GPU that runs comfortably at 1080p Ultra with 7.5GB of VRAM usage today may struggle with the same game after a patch adds higher-resolution texture options. Buying a GPU with adequate VRAM headroom at your target resolution is protection against premature obsolescence.

Our benchmark suite for this guide comprises ten titles released within the last 18 months: four demanding AAA games, three competitive multiplayer titles, two open-world simulations, and one synthetic workload that represents worst-case VRAM pressure. All testing is conducted at 1080p, 1440p, and 4K with upscaling technology disabled to establish native performance baselines. Upscaling results are presented separately in Chapter 5.

Testing methodology matters for comparison purposes. We use a standardised test system with a Ryzen 9 9900X processor — fast enough to avoid CPU bottlenecks at all resolutions in our test suite. This ensures the numbers reflect GPU performance, not system configuration differences. Results from reviews using slower processors will show lower GPU performance, particularly at 1080p where the CPU constraint is most significant.

Driver state at time of testing is documented for each GPU. Drivers affect performance meaningfully — AMD has delivered 5 to 15 percent frame rate improvements on specific titles through driver updates without hardware changes. If a GPU underperforms in a specific title at launch, check for driver updates before drawing conclusions about the hardware's fundamental capability.

The \$200–\$350 Tier — Budget 1080p

The \$200 to \$350 tier is the most competitive segment of the GPU market in 2026. It is where AMD and NVIDIA both have strong offerings, where used-market deals are most impactful, and where the performance delta between the cheapest and most expensive options is the largest relative to price. Choosing well at this tier requires understanding exactly what each card does and does not do.

AMD RX 7600 (8GB, approximately \$199): the value leader for raw rasterization at 1080p. In our benchmark suite at 1080p High settings, the RX 7600 averages 78 frames per second across all titles, with 1% lows averaging 58 frames per second. In less demanding competitive titles, the card easily exceeds 100 frames per second. The 8GB VRAM buffer is adequate for 1080p. The weakness is ray tracing performance, which is significantly below NVIDIA equivalents, and the absence of DLSS.

NVIDIA RTX 4060 (8GB, approximately \$299): the AI upscaling champion at this tier. Native rasterization performance is approximately 8 percent lower than the RX 7600 in raw frame rate terms. However, with DLSS 3.5 Quality mode enabled, the RTX 4060 matches or exceeds the RX 7600's native performance while adding Frame Generation support. In supported titles, Frame Generation can push frame rates to 150 percent of native numbers — a transformative advantage for demanding single-player games.

AMD RX 7700 (12GB, approximately \$279): a compelling value step up from the RX 7600 with 12GB of VRAM instead of 8GB and approximately 18 percent more average performance. The additional VRAM is genuinely meaningful at 1080p Ultra settings in titles with aggressive texture streaming. For builders who plan to keep their GPU for three or more years, the 12GB buffer provides substantially more longevity than the 8GB alternatives.

Used market options in this tier deserve serious consideration. A used RTX 3070 (8GB) for \$170 to \$200 delivers performance between the RX 7700 and RX 7600. A used RTX 3080 (10GB) for \$210 to \$240 outperforms everything new in this price range. The risk is undisclosed mining history — budget for the possibility of an RMA or return and plan to stress-test immediately after purchase.

Our recommendation at this tier depends on your priorities. If you primarily play the latest single-player AAA titles with DLSS support and value smooth visuals over maximum native frame rates, the RTX 4060 is the better choice. If you play a mix of titles including some without DLSS support and want maximum raw performance per dollar, the RX 7600 or RX 7700 are the better value. Both AMD options outperform the RTX 4060 in raw rasterization for their respective prices.

The \$400–\$600 Tier — 1440p Sweet Spot

The \$400 to \$600 range is what we call the 1440p sweet spot — the territory where you can consistently hit 60 frames per second or higher at 1440p High to Ultra settings in virtually every modern title, and achieve 100 frames per second or more in competitive and less demanding games. For builders pairing a new GPU with a 1440p 144Hz or 165Hz monitor, this tier is the target.

AMD RX 7800 XT (16GB, approximately \$419): the 1440p performance leader in terms of raw frames per dollar. With 16GB of VRAM, it is one of the best-equipped GPUs for its price in terms of buffer size. In our 1440p benchmark suite, the RX 7800 XT averages 82 frames per second across all titles at High settings and 68 frames per second at Ultra settings. The large VRAM buffer means no texture streaming constraints even in the most demanding current titles.

NVIDIA RTX 4070 (12GB, approximately \$549): the DLSS powerhouse at 1440p. Native rasterization performance is approximately 10 percent below the RX 7800 XT for roughly \$130 more. With DLSS 3.5 Quality mode, the RTX 4070 matches or exceeds the RX 7800 XT in supported titles. It is also substantially more power-efficient, drawing approximately 50 to 60 watts less under load. For compact builds with smaller PSUs or those that prioritise quiet operation, this efficiency advantage is meaningful.

AMD RX 7900 GRE (16GB, approximately \$399 on sale): when available at or below \$400, this card represents exceptional value. It performs between the RX 7800 XT and the RTX 4070 in raw terms, with 16GB VRAM and strong 1440p performance. Watch for sales — the RX 7900 GRE was originally a China-only product that AMD subsequently released globally and has been aggressively discounted.

VRAM considerations at 1440p: at this resolution and target quality level, VRAM usage regularly reaches 10 to 12GB in the most demanding titles at Ultra settings. Cards with 12GB are adequate but show occasional stutters when VRAM is full and the driver begins evicting textures. Cards with 16GB provide clear headroom. If choosing between an 8GB card at the lower end of this price range and a 16GB card at the upper end, the 16GB option is the more future-resistant choice.

Ray tracing at 1440p: NVIDIA holds a meaningful lead here due to hardware RT acceleration improvements between RDNA 3 and Ada Lovelace architectures. If ray tracing is a priority — particularly in titles like Alan Wake 2, Portal RTX, or Cyberpunk 2077 with full path tracing — the RTX 4070 is the right choice in this tier. For buyers who do not use ray tracing, the AMD options provide better value.

The \$700–\$1,000 Tier — High-Refresh 4K

The \$700 to \$1,000 tier is where 4K gaming becomes genuinely attainable at high settings and where 1440p becomes a high-refresh experience with all visual settings maximised. Cards in this range are also the sweet spot for content creators who need GPU-accelerated video encoding alongside gaming performance. These are serious investments that warrant careful comparison.

NVIDIA RTX 4070 Ti SUPER (16GB, approximately \$799): the current performance leader at this tier for combined gaming and content creation workloads. At 4K, it averages 61 frames per second across our benchmark suite at High settings — just over the 60fps threshold that defines comfortable 4K gaming. With DLSS 3.5 Quality mode, performance at 4K is excellent across all supported titles. The 16GB VRAM buffer handles 4K texture loads without constraint.

AMD RX 7900 XT (20GB, approximately \$749 on sale): the VRAM king at this tier. Twenty gigabytes of GDDR6 provides more than enough buffer for any current or near-future workload. Raw rasterization performance is approximately 5 percent behind the RTX 4070 Ti SUPER at 4K, but the price-to-performance ratio is strong, particularly during sales. For users running 4K with maximum texture settings in the most demanding titles, the 20GB buffer is a genuine advantage.

NVIDIA RTX 4080 SUPER (16GB, approximately \$999): the top of this tier and the card to buy when budget stretches. Performance at 4K averages 72 frames per second across our benchmark suite — enough headroom for upcoming titles to remain above 60fps for the next two to three years without upscaling. The RTX 4080 SUPER also features the best video encoding silicon in this price range, making it a strong choice for streamers and video editors.

Value analysis at this tier: the jump from \$550 (RTX 4070) to \$799 (RTX 4070 Ti SUPER) delivers approximately 28 percent more 4K performance for 45 percent more money. The jump from \$799 to \$999 (RTX 4080 SUPER) delivers approximately 20 percent more performance for 25 percent more money. Neither step is terrible value, but neither is exceptional. Buyers should evaluate whether the incremental performance improvement at their specific use case — resolution, refresh rate, title selection — justifies the additional cost.

Power and cooling considerations: cards in this tier draw 250 to 320W under full load. Ensure your PSU has a power rating at least 150W above your total system draw. The Corsair CV550 used in the \$500 build guide is inadequate at this tier — a quality 750W or 850W unit is the appropriate pairing. High-end cards also require good case airflow to maintain performance during extended sessions. Front-to-back airflow with at least two intake and one exhaust fan keeps temperatures in control.

DLSS 4 vs FSR 4 — An Honest Comparison

Upscaling technology has become one of the most important variables in GPU selection. Both NVIDIA's DLSS 4 and AMD's FSR 4 use machine learning to render frames at a lower resolution and reconstruct them to the target resolution with minimal visible quality loss. The performance gains are substantial. Understanding the differences — and the limitations — of each technology determines how much weight it should carry in your buying decision.

DLSS 4 (Deep Learning Super Sampling) runs on NVIDIA's Tensor Cores, which are dedicated silicon in RTX cards. DLSS 4 uses a transformer-based model trained on a vast dataset and runs the reconstruction entirely on the Tensor Cores without impacting the shader processors rendering the game. Quality mode (renders at 67% of target resolution), Balanced mode (58%), Performance mode (50%), and Ultra Performance mode (33%) offer a range of quality-to-performance trade-offs. In Quality mode, DLSS 4 output is frequently indistinguishable from native rendering in motion.

DLSS 4 also includes Frame Generation, which uses AI to synthesise entire intermediate frames between rendered frames. Frame Generation can nearly double the displayed frame rate with acceptable visual results in single-player content. It adds input latency — typically 8 to 15 milliseconds — which makes it unsuitable for competitive multiplayer at low base frame rates but acceptable in single-player games where immersion rather than reaction time is the priority. Multi Frame Generation in DLSS 4 can generate multiple intermediate frames, pushing displayed frame rates even higher.

FSR 4 (FidelityFX Super Resolution 4) is AMD's latest version, which moved from spatial scaling (FSR 1) through temporal reconstruction (FSR 2 and 3) to a full machine learning model in FSR 4. FSR 4 runs on AMD's AI accelerators in RDNA 4 architecture GPUs. Unlike DLSS, FSR 4 at full quality is available across AMD, NVIDIA, and Intel hardware through game integrations, though quality is optimal on AMD's own hardware.

In direct quality comparisons at equivalent settings, DLSS 4 Quality mode holds a slight edge over FSR 4 Quality mode in ghosting reduction on fine detail and in handling fast-moving transparent objects. The difference is visible on a high-quality display during slow-motion comparison but is difficult to notice during normal gameplay. FSR 4 represents a significant quality improvement over FSR 3 and is now a genuinely competitive technology rather than a clear second choice.

Buying decision impact: if upscaling quality and Frame Generation are high priorities — typically because your target GPU is at or near the native performance threshold for your resolution and settings — DLSS 4's advantages make RTX cards more compelling than pure native performance numbers suggest. If you play many older titles or titles without DLSS support, or if you prioritise native performance metrics, the upscaling advantage is irrelevant and native performance per dollar should be the primary decision factor.

VRAM Requirements by Resolution and Game Type

VRAM requirements are one of the most misunderstood aspects of GPU selection. Marketing materials rarely discuss VRAM clearly, and GPU manufacturers have strong commercial incentives to undersell VRAM requirements for their current products. This chapter provides honest VRAM consumption data by resolution and game category based on testing across our benchmark suite.

At 1080p, competitive multiplayer games including first-person shooters and battle royale titles typically consume 4 to 6GB of VRAM at high settings. These games optimise aggressively for frame rate at the expense of texture resolution, making 8GB sufficient with headroom to spare. The concern at 1080p Ultra settings appears in demanding open-world AAA titles, which regularly consume 7 to 10GB at maximum texture settings in 2026. An 8GB card running these titles at Ultra will exhibit texture streaming hitches when the buffer fills.

At 1440p, VRAM requirements increase substantially. The same open-world AAA titles that consume 7 to 10GB at 1080p Ultra consume 10 to 14GB at 1440p Ultra. Cards with 12GB will show occasional VRAM pressure in the most demanding scenarios. Cards with 16GB handle 1440p Ultra comfortably. For 1440p gaming with maximum settings across the full range of current titles, 16GB is the practical recommendation, not a luxury.

At 4K, VRAM requirements are severe. Demanding titles regularly consume 12 to 16GB at 4K Ultra settings. Cards with 12GB VRAM will stutter in these scenarios as the driver evicts and reloads textures. Sixteen gigabytes provides adequate headroom for current titles at 4K Ultra, but is becoming tight at the extremes. For 4K gaming, 16GB is the floor and 20GB or more provides genuine future-proofing.

Ray tracing multiplies VRAM consumption. Enabling ray tracing in titles like Cyberpunk 2077 at 1440p adds approximately 2 to 4GB to the base VRAM draw. Full path tracing (ray tracing at maximum quality) adds 4 to 6GB. A 12GB card that comfortably handles 1440p rasterization may show VRAM pressure immediately when ray tracing is enabled. Plan VRAM capacity with ray tracing in mind if that visual feature matters to you.

Future VRAM trajectory: game releases from 2025 and 2026 consistently use more VRAM than releases from 2022 and 2023, driven by larger texture atlases, more complex shader systems, and more aggressive scene geometry. The trend is unlikely to reverse. When evaluating a GPU today, consider what VRAM usage in the most demanding titles will look like in two years, not just today. A card that barely has enough VRAM now will be a card that constantly struggles in 18 months.

AMD vs NVIDIA Ecosystem Trade-offs

Beyond raw performance numbers, AMD and NVIDIA have meaningfully different ecosystems with genuine advantages and disadvantages on each side. Choosing between them based solely on benchmark frame rates misses important factors that affect the real-world ownership experience — particularly for users who create content, use their PC for work, or care about software compatibility.

NVIDIA's software ecosystem is broader and more mature. NVIDIA Broadcast uses AI processing to enable background removal, virtual backgrounds, and noise suppression in video calls and streaming with no external software required. NVIDIA's NVENC encoder, present in all RTX cards, is regarded as the best quality GPU encoder for streaming and video export — producing smaller file sizes at equivalent quality compared to AMD and Intel encoders. For streamers and video editors who also game, NVENC alone is a compelling reason to choose NVIDIA.

AMD's software ecosystem has improved substantially. AMD Software: Adrenalin Edition now includes Noise Suppression, enhanced streaming features, and AMD's Fluid Motion Frames (their Frame Generation equivalent, available on more titles than DLSS Frame Generation). AMD cards also offer better Open-Source compatibility for Linux users — AMD's open-source driver stack is more complete and better supported than NVIDIA's proprietary Linux drivers, an important consideration for dual-boot users.

Resale value historically favours NVIDIA cards. A two-year-old RTX card typically retains 45 to 55 percent of its original value on the used market. AMD equivalents retain 35 to 45 percent. This difference in residual value partially offsets the higher purchase price of NVIDIA cards when calculating total cost of ownership over a replacement cycle. Buyers who upgrade GPUs frequently should factor this into the decision.

Driver stability: both AMD and NVIDIA have had driver quality issues in recent history. AMD's driver stability has improved significantly from the problems that plagued RDNA 1 launch drivers. NVIDIA has had widely reported driver issues with specific titles in 2025. Neither platform has a monopoly on driver quality in 2026. Read current driver release notes and community forums for the specific GPU you are considering to get an up-to-date picture of driver health at time of purchase.

Display technology compatibility: NVIDIA G-Sync compatible monitors work with AMD cards at their base adaptive sync functionality, but do not access all G-Sync hardware features. NVIDIA cards work with FreeSync monitors through the G-Sync compatible programme. AMD cards work natively with FreeSync monitors without restriction. If you already own a specific monitor, verify compatibility with your chosen GPU ecosystem before purchasing.

When to Buy and When to Wait

The question of when to buy a GPU is perennial, and the honest answer is more nuanced than 'always wait.' GPU releases follow predictable cycles, prices follow predictable patterns, and the performance improvement per generation follows a predictable trend. Understanding these patterns lets you make an informed decision about timing rather than defaulting to perpetual waiting.

The GPU release cycle: NVIDIA and AMD both release new GPU generations approximately every 18 to 24 months. NVIDIA's Ada Lovelace architecture (RTX 40 series) launched in late 2022. The RTX 50 series (Blackwell consumer) has begun launching at the high end in early 2025, with mid-range variants becoming available through 2026. AMD's RDNA 4 architecture launched in early 2025 and mid-range RDNA 4 cards are now available.

Price drop timing: the best time to buy a current-generation GPU is six to twelve months after its launch, once production ramps up and initial demand stabilises. The absolute lowest prices on current-generation cards typically occur in the three months before a major new launch announcement, when manufacturers and retailers discount existing inventory to clear space. This window is predictable but requires patience.

Performance improvement per generation: the raw performance improvement from one GPU generation to the next has moderated from the 30 to 40 percent improvements of earlier eras to 15 to 25 percent for mid-range cards in recent cycles. The larger improvements in current-generation GPUs come from upscaling quality improvements and efficiency gains rather than raw rasterization performance. If upscaling quality matters to you, a generation upgrade provides real benefit beyond the frame rate numbers.

When waiting makes sense: if you currently have a GPU that runs your target games at acceptable frame rates but are considering upgrading, wait until a new generation is announced and current-generation prices drop. If a new generation launched in the last six months and mid-range pricing has not yet settled, wait six more months. If you can see a specific upcoming launch that addresses your needs directly, a short wait is rational.

When buying now makes sense: if your current GPU is failing, insufficient for your target games, or you do not currently have a GPU, buy now. The performance you get today has real value that accumulates every day you wait. 'Better is coming' is always true — it is not a reason to wait indefinitely. If a GPU at your target price today meets your performance requirements, buy it and enjoy it. The returns to waiting diminish significantly once you already have an adequate specification.

Price watch tools: use CamelCamelCamel to track Amazon price history for a specific GPU. Use PCPartPicker's price alert feature for your target card across multiple retailers. Set a maximum price threshold based on your analysis and let the alert come to you rather than checking daily. This removes the emotional element from the decision and ensures you buy at a price you determined was fair before desire to purchase biased your judgement.

Thank you for reading!

Find more guides, reviews, and build lists at

jnrbuilds.com