NVIDIA RTX 4080
First review of Nvidia's newest Lovelace GPU PG.74

BEEF UP SECURITY
Our tips to stay safe, secure, and private online PG.44

RETRO EMULATION
How to play your favorite games from the past PG.52

MAXIMUM PC

MINIMUM BS • JANUARY 2023 • www.maximumpc.com

GET RTX 4090 FOR LESS!

JUST $2,940!
Step-by-step guide to assembling this PC
PG. 16

Get maximum performance for minimal bucks

- Choose the best value components
- 4K gaming on a mid-range budget

AMD RDNA 3 GRAPHICS

Everything you need to know about Radeon RX 7900 XT and XTX PG. 32

FUTURE Digital Edition

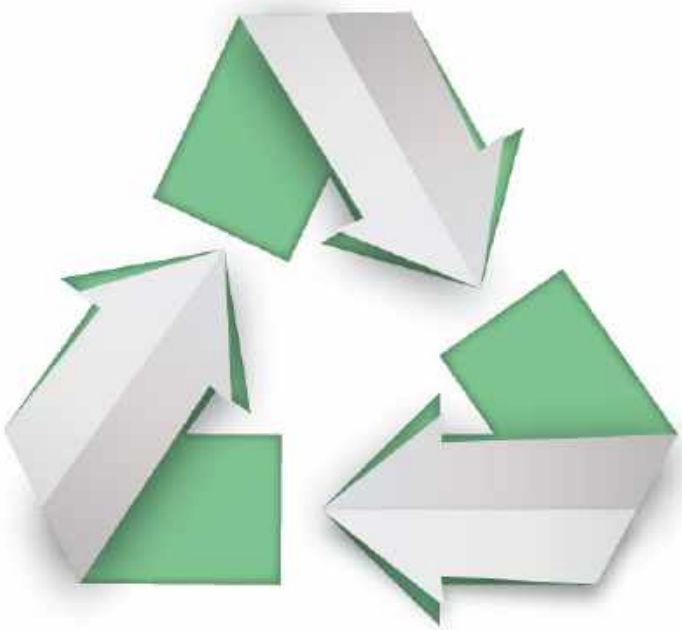

PAPER REVOLUTION

68% of paper and paper-based packaging in the U.S. is recycled into new products, one of the highest recycling rates of any material in the country.

Sources: American Forest and Paper Association, 2022
U.S. Environmental Protection Agency, 2022

Discover the story of paper
www.lovepaperna.org
Scan for paper facts, activities, blogs and much more!

LOVE PAPER®

inside

JANUARY 2023

SCAN TO GET THE TOM'S HARDWARE WEEKLY NEWSLETTER

QUICKSTART

8 THE NEWS
The GPU war hots up; new world record overclocking record set; FakeCatcher combats deep fakes.

12 THE LIST
Best 4K TVs for PC gaming.

This TCL is very much in the budget camp, but it's a world away from earlier cheap TVs.

42 SUBSCRIBE TODAY
Subscribe to *Maximum PC* today to save money and receive a *Book of Amazing Technology*.

R&D

61 HOW TO
Streamline your PC's startup process; Colorize B&W images.

62 AUTOPSY
Tearing down Xbox's adaptive controller with help from iFixit.

LETTERS

14 DOCTOR

94 COMMENTS

IN THE LAB

74 GEFORCE RTX 4080 FOUNDERS EDITION

80 ASUS ROG STRIX Z790-A GAMING WIFI D4

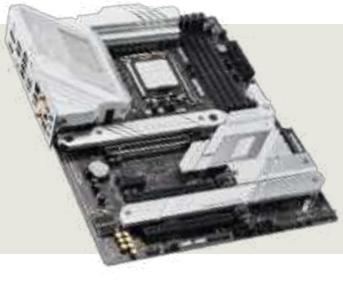

88 SUREFIRE HARRIER 360

16 SLICK & SUBTLE 4090 RIG

32 AMD RDNA 3 ARCHITECTURE
What does RDNA 3 have in store? *Jarred Walton* deep dives into AMD's latest architecture.

44 PC SECURITY AND PRIVACY
With threats to your online security, *Nick Peers* shows how to lock your data and identity.

52 RETRO EMULATION
Jonni Bidwell explains how to play classic retro games on your PC and Raspberry Pi.

90 CALL OF DUTY: MODERN WARFARE II (2022)

techradar.pro

IT INSIGHTS FOR BUSINESS

Best practices and solutions

Up-to-the minute news

In-depth reviews

http://www.techradar.com/pro

MAXIMUMPC

EDITORIAL
Editor-in-Chief: Guy Cocker
Staff Writer: Sam Lewis
Contributing Writers: Jeremy Laird, Chris Lloyd, Ian Evenden, Jarred Walton, John Knight, Nick Peers, Jonni Bidwell
Production Editor: Jon Crampin
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Photography: Neil Godwin, Olly Curtis, Phil Barker
Cover Photo Credits: Magictorch

BUSINESS
US Marketing & Strategic Partnerships: Stacy Gaines, stacy.gaines@futurenet.com
US Chief Revenue Officer: Mike Peralta, mike.peralta@futurenet.com
East Coast Account Director: Brandie Rushing, brandie.rushing@futurenet.com
East Coast Account Director: Michael Plump, michael.plump@futurenet.com
East Coast Account Director: Victoria Sanders, victoria.sanders@futurenet.com
East Coast Account Director: Melissa Planty, melissa.planty@futurenet.com
East Coast Account Director: Elizabeth Fleischman, elizabeth.fleischman@futurenet.com
West Coast Account Director: Austin Park, austin.park@futurenet.com
West Coast Account Director: Jack McAuliffe, jack.mcauliffe@futurenet.com

Director, Client Services: Tracy Lam, tracy.lam@futurenet.com

MANAGEMENT
MD Tech: Paul Newman
Group Editor-in-Chief: Graham Barlow
Group Art Director: Jo Gulliver

PRODUCTION
Head of Production: Mark Constance
Senior Production Manager: Matthew Eglinton
Production Manager: Vivienne Calvert
Production Assistant: Emily Wood

FUTURE US, INC.
10th floor, 1100 13th Street, N.W Washington, DC 20005, USA
www.futureus.com

INTERNATIONAL LICENSING & SYNDICATION
Maximum PC is available for licensing and syndication.
To find out more, contact us at licensing@futurenet.com or view our available content at www.futurecontenthub.com.
Head of Print Licensing: Rachel Shaw, licensing@futurenet.com

SUBSCRIBER CUSTOMER SERVICE
Website: www.magazinesdirect.com
Tel: 844-779-2822
New Orders: help@magazinesdirect.com
Customer Service: help@mymagazine.co.uk

BACK ISSUES
Website: https://bit.ly/mpcsingleissue

Next Issue On Sale 31st January, 2023

[SUBSCRIPTION DELAYS]
We rely on various delivery companies to get your magazine to you—many of whom continue to be impacted by Covid. We kindly ask that you allow up to seven days before contacting us about a late delivery to help@magazinesdirect.com

Guy Cocker

MID-RANGE MONSTER

LAST MONTH, I went all-out on building a system with the latest and greatest (i.e. most expensive) components available. It was an exercise in seeing what was capable at the cutting edge with just-released tech and a money-no-object attitude. And what a system it was, offering a 25 percent performance leap in CPU-bound tasks like Cinebench, and an astonishing 80 percent performance leap in 4K games.

This month's build is more in line with what I'd go for, personally, if I were building a gaming PC right now. It still uses the RTX 4090 for those crazy 4K frame rates, but it's paired with the brilliant Intel Core i5-13600K, which is $300 cheaper than the i9-13900K I used. Our builder Sam has also smartly gone for a Z690 board, still using faster DDR5 memory, but another $300 saving over the Z790 board I used. Does it still perform as well in games, however? Well, Sam has the full lowdown in his feature starting on page 16, where he shows you how to build this machine in detail, and how it compares to last month's system. I think you'll be interested in the results.

Aesthetically speaking, Sam's vision for the build was also to make a more subtle looking machine than your traditional gaming PC. Sure, there's a splash of RGB from the GPU inside, but the CPU is air-cooled and the other components are all very stealthy. For those who want something that's a bit more ostentatious, fear not, as Sam is currently working with Corsair on an all-out water-cooled build for a future issue—if you've always wanted a PC that's worthy of having its own Instagram account, keep an eye out for that.

If you've been holding off buying an RTX 4090 to see how the GPU market shakes out, this issue holds the answer on how to best spend your money. Our resident graphics specialist, Jarred, has put the RTX 4080 to the test in our labs to deliver a definitive verdict on the less-expensive card. Is it the right path to DLSS 3 and ray-traced graphics? Find out on page 74. Jarred has also processed the facts and figures from AMD's recent RDNA 3 announcement to find out whether it's worth buying the RX 7900 XTX or Radeon RX 7900 XT GPU. His deep dive begins on page 32.

We also look at one of the issues I receive the most correspondence about, namely security and privacy. There seems to be a news story every month about a social media data leak, government agencies being able to access your data, or browser 'Incognito' modes not being as anonymous as you'd hope. Nick Peers looks at how best to protect yourself online, from the basic protocols you should all be following, to the advanced practices you can use to truly cover your tracks. It starts on page 44.

We also have our first review of a pre-built Intel Raptor Lake/RTX 4000-series desktop this month in the form of the astonishing Velocity Micro Raptor Z95 (page 76). We have an excellent 'how to' on streamlining your PC's resources (page 64)—free tips to make sure Windows 11 is operating at its best. Plus, we review *Call of Duty: Modern Warfare II*—it's already the best-selling *CoD* of all time, but is it worth playing this year? Find out on page 90.

Enjoy the issue!

Guy Cocker

Guy is Maximum PC's Editor-in-Chief. He built his first gaming PC in 1997 to play Tomb Raider *on 3dfx, and has been obsessed with all things PC ever since.*

↘ submit your questions to: editor@maximumpc.com

THE NEWS

GPU War Heats Up

The next generation of cards finally go head-to-head

AMD'S NEW GPU series, known as RDNA3, has landed in the form of the 24GB Radeon RX 7900 XTX and the 20GB 7900 XT. AMD claims the cards have the most advanced technology and 'breakthrough' levels of performance. Both cards use the 5nm Navi 31 die, the biggest the company has, with the 7900 XT being a slightly cut-down version with 84 compute units against its bigger sibling's 96. It's a chiplet design, with the main 5nm GPU surrounded by six 6nm memory chips with attendant controllers. The game clocks are 2.5GHz and 2.3GHz respectively. Much to everybody's relief, they use two 8-pin power connectors. For a more in-depth look at the technology inside the new RDNA 3 cards, see page 32.

First impressions are that it looks as if AMD is on to a winner. AMD estimates that the RX 7900 XTX is between 1.5x and 1.7x faster than the previous generation's RX 6950 XT. The prices aren't cheap, but they do seem competitive next to Nvidia's next-gen cards. The RX 7900 XT is $899 and the 7900 XTX is $999. Not exactly loose change, but given their expected performance, it seems reasonable. Nvidia's monstrous GeForce RTX 4090 may offer the ultimate performance, but the RX 7900 XTX gets close for $600 less.

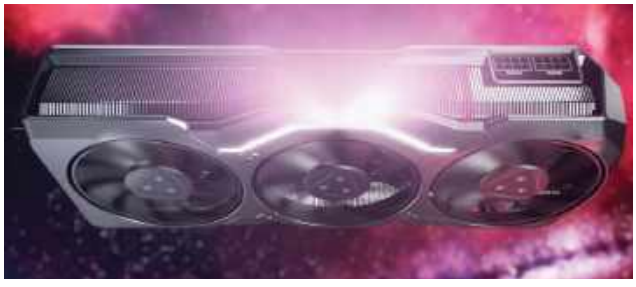

When AMD promised that the RDNA 3 was going to be significantly faster than any other Radeon, it wasn't wrong.

Over in the rival camp, the Nvidia GeForce RTX 4080 has landed (see review, page 74). It's something of an oddity—the 4090 is appreciably more powerful, yet the 4080 is still painfully expensive at $1,199. Many buyers are opting for its big brother. Nvidia seems to agree, and supplies of the 4080 are much lower than for the more expensive RTX 4090. If anything, the release of the RTX 4080 has cemented the 4090 as the real object of desire for most gamers. AMD has thrown a little mud Nvidia's way over the 4080 by sending a series of slides to media outlets claiming its GPUs consume less power (true), are more affordable (true), and have more memory for the money (also true). Plus, you don't need a new ATX 3.0 power supply or case.

Meanwhile, we still have the 'issue' of the RTX 4090's power cables melting. An early tester managed the feat, and there is evidence of a few sub-standard cables in early examples of the card, but most people don't appear to be having any trouble. A few have tried to deliberately replicate the problem too, without success. The chief of Corsair has put the blame on poorly-inserted cables, pointing out that they are difficult to insert fully, and that they need to be.

After a spell of silence from Nvidia, we now have an official response after studying a number of melted cables. It claims to only be aware of 50 cases worldwide, and that a common factor was the cables not being inserted all the way. It has also promised to replace any damaged hardware.

The issue has gained a lot of publicity, which it transpires is out of proportion with the actual problem. This didn't stop AMD from joining in. A senior AMD Director of Gaming tweeted "Stay safe this holiday season" with an image of AMD's twin 8-pin power sockets. It also hasn't stopped one unhappy customer from filing a class action for breach of warranty, fraud, and more. He paid $1,600 for a 4090, installed it, and promptly melted the cable. He claims Nvidia "marketed and sold the RTX 4090 with a defective and dangerous power cable plug and socket", and is backing this up with evidence of 26 other occurrences. So, if you get a 4090, make sure those power cables are pushed in all the way. Then check again.

So, the stage is set for 2023 in the GPU market. In one corner is the green team with a power-hungry beast that will top the benchmarks, but with a less impressive lower model. The red team's two capable cards aren't too far behind and are much cheaper. Whichever side you pick, you'll have a card that is streets ahead of the previous generation. What's going to be interesting is when the mid-market cards arrive. If AMD's pricing continues to undercut Nvidia by decent margins, the company could do well. It appears that using the chiplet design has paid dividends, if not delivered the laurels for making the fastest card ever. **–CL**

> "Whichever you pick, you'll have a card that is streets ahead of the previous-gen.

FASTEST HARD DRIVES EVER

SEAGATE'S SECOND GENERATION of Mach.2 drives, the Exos 2X18 and 2X16, can manage data transfer rates of 545MB/s over a SATA 3 interface, which is comparable to SATA SSD performance. They build on the slower 2X14 that began shipping in 2019 and, while they're no match for an M.2 NVMe SSD, these mechanical drives come in larger capacities of 16TB and 18TB.

The magic isn't achieved by spin rates—those are a modest 7,200rpm—but using two actuators effectively means it can act as two drives simultaneously. One area where they won't match solid state is in seek times, which at 4.16ms are nowhere near the microseconds of an SSD. These drives are aimed at server work and 'hyperscale' workloads, so are robustly built, with SAS and self-encrypting versions too. Prices have yet to be revealed, but as enterprise-grade kit, they won't be cheap. But if you need SATA 3 storage that's both fast and capacious, this is where it's at. **–CL**

ZEN 4'S V-CACHE CHIP DUE SOON

But they will be limited to eight-core

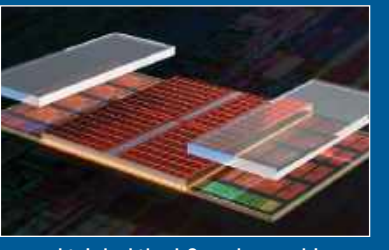

AMD'S SECRET weapon for the Zen 3 range was its V-cache version, the Ryzen 7 5800X3D. It stacked high-speed memory on top of eight cores and tripled the L3 cache, making for the world's fastest gaming chip on launch. Although Intel responded with the Core i9-12900KS to steal some of that thunder, Intel's chip was faster in some games, albeit more expensive. AMD had no trouble selling its V-cache chips, and even now they're still hard to find. The same V-cache editions will be released for the Zen 4 Ryzens and should be due in the first half of 2023.

The hope was that the cache trick would be applied to a higher-end version with more cores, but rumor has it that AMD will start small, with a six-core Ryzen 5 7600X3D and an eight-core Ryzen 7 7700X3D. If pricing follows the pattern, they'll cost $150 more than the standard chips. We're left tantalized by the prospect of even more powerful chips being V-cache enabled. **–CL**

Tech Triumphs and Tragedies

A monthly snapshot of what's good and bad in tech

TRIUMPHS

SUPER-EFFICIENT
The Flatiron Institute's Herni system returns 65.091GFLOP/Watt, making it the world's most power-efficient supercomputer.

MUSHROOM ELECTRONICS
The skin of the Ganoderma Lucidum fungus can be dried and used to make PCB boards.

8K GAMING MONITOR
Samsung's new Odyssey Neo G9 is a 49-inch, 7,680 by 2,160 pixel, curved monitor with a refresh rate of 165Hz.

TRAGEDIES

SLOW YEAR FOR WIN11
Figures from Startcounter put Windows 11's share of the PC market at just 15 percent, one year after its release.

NO MORE HANGING OUT
Launched in 2013, Hangouts, Google's try at a messaging service, is now officially dead.

GAME DISC WITH NO GAME
Buy *Call of Duty: Modern Warfare 2* on disc and you'll find just 72MB of data—there's a huge download instead.

WORLD RECORD OVERCLOCK

Fastest ever desktop processor now 8.8GHz

INTEL'S 13TH-GENERATION Raptor Lake processors are impressive beasts. After struggling against AMD's Zen-based Ryzen chips in the past few years, the blue team now has a winner. We have solid proof of how robust the flagship i9-13900K model is, thanks to a world overclocking record of 8.81285GHz, achieved with the help of liquid nitrogen cooling, a core voltage of 1.325V, a multiplier of 88, and a tweaked bus of 100.15MHz.

Raptor Lake was always going to be a contender—before it was released, engineering samples were reported to reach 8GHz. The eight-year-old record was held by AMD's FX-8370, a 32nm eight-core chip based around the Piledriver architecture, which managed 8.72278GHz. So, Intel took the record with an improvement of 90MHz, just one percent.

Meanwhile, there are rumors of a refresh for Raptor Lake in the second half of 2023, with an announcement expected at CES in Las Vegas in January. Up to 16 new versions could arrive, along with the B790 motherboard chipset. There's no hint of the Core i9-13900KS yet, though Intel has confirmed its existence, with a clock speed of 6GHz out of the box.

After that, we've got Meteor Lake, which follows the same design structure as Raptor Lake and is expected to focus on power efficiency with an IPC increase of 15 percent or more. We could see these as early as the end of the year. These will have Redwood Cove performance cores and Crestmont efficiency cores. By 2024, the first 15th-generation Arrow Lake chips are expected, with Lion Cove performance cores and Skymont efficiency cores. The aim is for an IPC increase of between 22 and 34 percent. Pat Gelsinger's plan to introduce five nodes in four years sounded optimistic, but Intel is going for it. **–CL**

Raptor Lake puts Intel back in the game and is proof Gelsinger's plan is no idle boast.

META LOSES 11,000 STAFF

Layoffs spike across entire tech industry

THE NEWS for Meta's Metaverse goes from bad to worse. Now, Mark Zuckerberg has laid-off 11,000 employees, including many working in his Reality Labs division—a cut of around 13 percent. The global pandemic led to a spike in revenue as people went shopping and looking for entertainment online. Much of that extra revenue was invested in new staff. Things temporarily went back to normal, but then sharply became worse in 2022, which has proved to be a problem. Zuckerberg has admitted over-estimating how long the pandemic boom would last, saying it wasn't "the wisest decision he could have made".

Meanwhile, at Twitter, Elon Musk fired around 3,700 people, or half the workforce, including many contractors monitoring the platform's content. Amazon has also recently announced plans to lay 10,000 people off, although it has over 1.5 million employees worldwide, about a fifth of which are corporate roles. Other companies to shed staff include Cisco, which lost 4,300, Stripe (1,100 or 14 percent of its workforce), Shopify (1,000, 10 percent), Coinbase (1,100, 18 percent), ToGo (1,300), and even Microsoft (1,000).

According to TrueUp's online tech layoff tracker, there have been 202,522 internet-based tech jobs lost worldwide up to mid-November 2022. So what's happening at these big tech companies? The tech industry isn't alone in suffering the effects of a global economic downturn, of course, but laying off staff is an option it often uses to cut costs. Before we get too pessimistic, the unemployment rate for tech workers is just 2.2 percent (below the national average), and it's an industry that finds it easy to employ people in better times. **–CL**

Zuckerberg's virtual world is being built by fewer people, thanks to a round of layoffs.

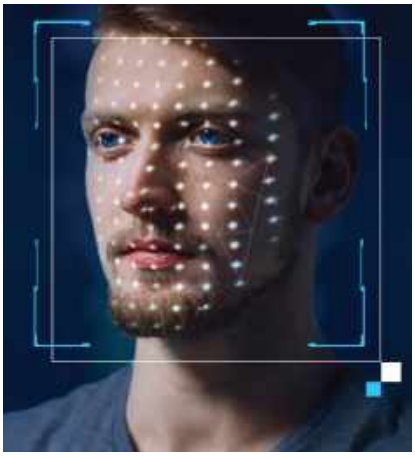

Intel's FakeCatcher

DEEPFAKES, VIDEOS DOCTORED BY AI to appear real when they're not, have been around for years. But they hit the headlines around 2018 when it became possible to create them on modest systems with readily available software. There was something of a panic as it became clear that our world of digital imagery was forever untrustworthy. Editing images is as old as the invention of photography, but deepfakes made realistic videos a reality. At their simplest, deepfakes work by mapping an AI-manipulated face over the movements and expressions of an actor to alter an original clip.

Intel's FakeCatcher software has a claimed accuracy of 96 percent and can spot a fake in milliseconds. It can even uncover the source model with a 93 percent accuracy. How does it manage this? It uses blood. As your blood circulates around your body, your veins pulse and change color, an effect called photoplethysmography (PPG). FakeCatcher looks for PPG signals across the face and creates a map. This is run through some machine learning wizardry to see if it matches reality. Intel also uses eye gaze detection. This all requires considerable power: FakeCatcher currently runs on a Xeon-powered server that can cope with 72 streams at once. The idea is that news outlets and social media platforms can check suspicious content.

FakeCatcher is part of Intel's Responsible AI work. The company is also a member of the Project Origin Alliance and the Content Authenticity Initiative, so is taking this issue seriously. Already legal moves are being considered to hold platforms responsible for spreading deepfakes if they don't make an effort to spot them. Fakes and fake detectors look set to battle it out for the foreseeable future. **–CL**

The Internet on One Fiber

A new data transfer record using a single strand of a fiber optic cable has been set at 1.53 petabits. A petabit is 1,000 terabits (although it's not clear whether they used decimal or binary, which would make it a pebibit). To put this into perspective, the latest study by TeleGeography, an international telecommunications consulting company, estimates that total internet traffic for 2022 averages at 997Tbps, under one petabit, so you could run the entire global internet down one 0.125mm fiber.

The feat was achieved by a team at the National Institute of Information and Communications Technology in Japan. They used light in 55 specific frequencies, a process known as multiflexing. It's not the fastest data transfer (1.84 petabits on an experimental optical chip), but it used standard fiber, so the existing infrastructure could be upgraded. **–CL**

matter

Smart Homes to get Smarter

One problem with smart homes is they aren't particularly smart because the IoT devices can't communicate with each other. However, the Connectivity Standards Alliance (formerly Zigbee) has defined a new standard to address this, called Matter.

The project started in 2019, and while there have been some delays, the standard is finally finished and over 280 companies have signed up. The result is that certified Matter smart devices will be able to work co-cooperatively, so you're not tied to a single brand, and don't have to continually update to new standards. It will take some time to work through all the products—right now, there are around 200 certified devices, most of which are light bulbs. **–CL**

Jarred Walton

TECH TALK

Understanding Power and Efficiency

CREATING A MODERN microprocessor involves myriad decisions, including size and transistor counts, but even with a finalized design, there are still a few knobs that can be used for further tuning: voltage and frequency. These tools are often available to end users and lie at the heart of overclocking. As the saying goes, with great power comes great responsibility.

I mean that literally in this case. Power increases directly in proportion with frequency and with the square of the voltage. That's why higher clock speeds require more power, but if you also need more voltage to stabilize those higher clocks, you get a double whammy. The reverse is also true, naturally—drop the voltage and/or frequency and you can reduce power requirements.

The RTX 4090 is the first of Nvidia's Ada Lovelace GPUs, and there's been a lot of noise about the maximum 450W power rating of the reference design, with custom cards potentially drawing 500W or more. Meanwhile, the RTX 4080 curtails power use (see review on page 74), with a limit of 320W by default. But while both cards hit their power limits in some games, there are also plenty of situations where they don't need as much power—and we can manually set a lower power limit to greatly boost the overall GPU efficiency.

Many previous-generation graphics cards offer the ability to adjust voltage, frequency, and power limits with software such as MSI Afterburner. So far, Ada Lovelace GPUs are more limited and require voltage mods with a soldering iron to unlock more performance. However, we can still set power limits and let the hardware manage things within those constraints.

What happens to performance and power use with tighter limits? I took an RTX 4090 Founders Edition and ran tests on eight demanding ray tracing games at 4K and maxed out settings. I chose ray-tracing titles because they tend to be more power-hungry than traditional games. *Forza Horizon 5*, for example, used around 310W at 4K Extreme. I repeated those tests with a power limit of 90, 80, 70, 60, and 50 percent.

> The more power a game uses at 100 percent, the bigger the impact of a lower power limit.

The more power a game used at the default 100 percent setting, the more the impact of a lower power limit became. That makes sense, but across all the power limits, power use decreased faster than frame rates—or in other words, the RTX 4090 become more efficient. The sweet spot seems to be around the 70 percent mark, at which point the average power use dropped by 22.6 percent while performance only fell by 5.2 percent.

Using eight games in aggregate hides the impact of outliers that might become more common over the coming years, so let's also look at the extremes. The worst-case result was *Metro Exodus Enhanced Edition*, where performance at the 70 percent power limit decreased by 8.2 percent. That's noticeable outside of running benchmarks, but power use was still 27.7 percent lower. The best-case result was in *Minecraft*, which only used 380W at default, so lowering the power limit was less impactful. Its power use still fell by 17.9 percent with the 70 percent limit in place, while performance hardly changed at all with a 3.4 percent dip.

Nvidia's RTX 4090 can use up to 450W of power, but could be more efficient with a lower power limit.

However, overclocking naturally shows different results. Cranked up to a 133 percent power limit and with higher GPU and memory clocks, the RTX 4090 Founders Edition performance improved by a whopping... 3 percent. Power use, on the other hand, jumped up by 16.8 percent, averaging 474W across the test suite compared to 406W at stock—with a worst-case power draw of 515W.

I talked about the early news of melting 4090 power adapters last month, and Nvidia's official response is that it's extremely rare and mostly down down to loose power cables. Overclocking certainly won't help matters, but reducing power draw through the 16-pin adapter can't hurt. If you happen to have an RTX 4090, giving up performance to improve efficiency isn't a bad idea.

Jarred Walton has been a PC and gaming enthusiast for over 30 years.

THE LIST

THE BEST 4K TVS FOR PC GAMING

WITH THE LAUNCH of the RTX 4090 and 4080, (see page 74) we're now able to run games at 4K Ultra settings at 60 frames per second. While that's great if you have a 4K monitor, they're often expensive and you need a larger display (and therefore a big desk) to really make the most of that resolution. It might be more justifiable to spend that money on a 4K TV, which will be great for not just gaming but also sports, reality TV, movie nights... the list goes on. These are the best.

5 TCL 55R617 ROKU TV

www.tcl.com, $530

This is very much in the budget camp, but it's a world away from earlier cheap TVs from TCL. It includes Roku TV for all of your Smart TV functionality and voice control. For gamers, it also has full-array local dimming for great contrast ratios, HDR support, and a decent 6–12ms response time. It's not the best quality display here, but for the price, it's hard to fault.

4 SONY A8H OLED

www.sony.com, $1,900

If you're thinking of wall-mounting your TV, then this Sony is the one you want. It's beautifully thin and you get all the benefits of an OLED here at a reasonable price, including crazy deep black levels, great viewing angles, and a fantastically detailed 4K picture. Be aware though, there's no variable refresh rate (G-Sync, FreeSync) or HDR10+. For those, look to the number 1 entry on this list.

3 HISENSE 55H8G

www.hisense.com, $800

If you're looking at the price of the other TVs here and scoffing, Hisense has been making quality low-cost 4K TVs for years. There are no compromises—it supports HDR10 and Dolby Vision HDR, Black Frame Insertion helps reduce motion blur, and there's a dedicated game mode.

2 SAMSUNG Q9F

www.samsung.com, $2,500

LG may have the inkiest blacks around but Samsung's QLED tech produces the most vibrant colors and vivid contrasts we've seen on a television, thanks to its HDR 2000 tech. The Q9F also does a great job with low-quality sources, meaning you don't need a GPU capable of 4K to make the most of it.

1 LG OLED48CX www.lg.com, $1,500

This 48-inch LG OLED works as well on your desk as the living room wall. LG's black levels are unsurpassed, it has a 120Hz refresh rate, and there's AMD FreeSync Premium and Nvidia G-Sync compatibility. The CX is as close to perfect as it gets.

Jeremy Laird

TRADE CHAT

The GPU pricing paradox

SOMETHING ODD IS GOING ON in the PC graphics market. Nvidia and AMD have launched their latest GPUs, but they seem priced to fail. I'm not talking about the $1,600 Nvidia GeForce RTX 4090—the only one of the new generation that offers value compared to its predecessor. No, it's every other graphics card launched by the dynamic duo that's the problem.

In case you've been living under a rock, that's the Nvidia GeForce RTX 4080 16GB, and the AMD Radeon RX 7900 XT and 7900 XTX. The GeForce card is priced at $1,200, while the AMD boards are $899 and $999 respectively. Of course, those are MSRPs, rather than real-world street prices.

Recent history has seen prices of pretty much all GPUs outstrip any MSRP. Indeed, I recently saw an Nvidia GeForce RTX 3080 from the previous generation priced at $699 for the first time. So, it's taken two years from launch to see the inflated price of the RTX 3080 fall back to MSRP. Incredible.

With that in mind, the crazy high prices of those new GPUs makes superficial sense. Over the last couple of years, gamers have been willing to spend megabucks on graphics. The genie is out of the bottle—$1,200 for a second-tier GPU like the RTX 4080 is the new reality. And the only reason why AMD's boards top out at $1,000 is that the new 7900 simply can't compete at the high end.

The latter point is certainly true. AMD hasn't attempted a GPU to take on the mighty RTX 4090. Or rather if it has, it's either yet to come, or it suffered some kind of calamity and has been canceled. I happen to think the latter is more likely. But right now, the big question for PC gamers is whether that narrative of a new reality is realistic.

The good news is that there's every reason to think it's wrong. While gamers have proved willing to pay extortionate prices over the past two years, it doesn't automatically follow that they'll continue to do so. Market conditions of late have been little short of freaky, what with a huge spike in demand thanks to ethereum mining and the influence of a global pandemic.

> Who will want to pay $1,000 for a GPU with sub-standard ray-tracing performance?

AMD's latest Radeon RX 7900 boards could be a tough sell.

Both of those factors have now essentially evaporated. Of course, some lingering pandemic-related supply constraints still prevail. And it's unlikely that we've heard the last of crypto-fueled demand for graphics cards. But the majority of the upward price pressure provided by those exceptional externalities has vanished.

In the meantime, both AMD and Nvidia have cranked out a huge number of those previous-gen GPUs. Many are as yet unsold and far more again are being unleashed onto the used market as the whole crypto thing unwinds following the great ethereum merge and the currency's shift from proof of work to proof of stake.

As a consequence, this year has seen GPU prices tumble dramatically. That talismanic RTX 3080 is getting on for one-third of the price it was at the beginning of the year, with prices still falling. At the same time, widespread reports of the RTX 4080 selling poorly have emerged. For sure, the RTX 4080 did not rapidly sell out after launch, something that had become the norm in recent years.

Of course, retailers can't immediately slash prices on the RTX 4080s they already have in stock to get the cards shifted as they'd be making big losses. So, it will take a while for pricing to reflect the reality of much lower demand. But that is the reality. Demand for GPUs is a fraction of what it was when gamers were paying megabucks. Eventually, that will be reflected in prices.

I've got a feeling that the AMD Radeon RX 7900 launch is going to be even worse. Who's going to want to pay $1,000 in this market for a GPU with sub-standard ray-tracing performance? Anyway, although I don't expect prices to immediately implode, 12 months from now, I reckon the GPU market is going to look an awful lot like it did before the Ethereum-Covid double whammy.

Six raw 4K panels for breakfast, laced with extract of x86... Jeremy Laird eats and breathes PC technology.

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> Free data recovery
> SATA controller
> New Win11 features

Drive stopped working

I have a 3.5-inch hard drive that has been attached to my PC via an external SATA adapter for over a year now. Recently, I booted into Windows only for the computer to hang.

I eventually traced the problem to this drive—powering it off 'fixed' the problem but introduced a new one. Now, when I power the drive back on, its drive letter appears in File Explorer but shows no drive space information.

If I double-click the drive to try and open it, I eventually get a 'Location not available: The file or directory is corrupted and unreadable' message. If I open Disk Management, the drive shows up as healthy, but as a RAW partition. Does this mean the drive is toast?

—John Arteaga

THE DOCTOR RESPONDS:

It sounds like the drive's partition header has become corrupt somehow—John mentioned that the drive was running hot when it malfunctioned, but as drives routinely run 35-50°C when under load this shouldn't necessarily indicate there's a problem.

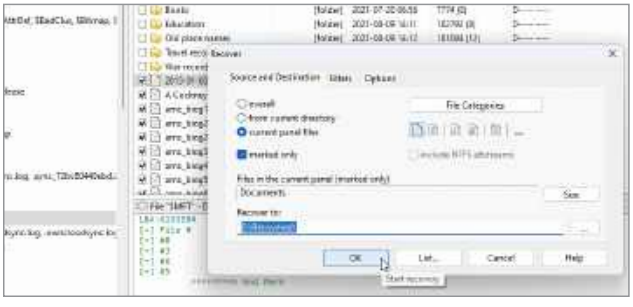

DMDE can recover files from lost and deleted partitions.

We instructed John to download and extract the 64-bit GUI version of the free portable DMDE data recovery tool (https://dmde.com/). Once that's installed, launch dmde.exe and select a drive from 'Physical Devices' from the 'Select Disk/Task' page and click OK.

DMDE will list all volumes it can see, including those not currently visible to Windows. In John's case, his NTFS partition was displayed. Double-clicking this opened a new tab, revealing empty $MetaData and $Root folders, plus an [All Found/Virtual FS] entry that he needed to click. This provides an option to build a read-only virtual file system. Untick 'include deleted' and leave 'Pure FS reconstruction' selected before clicking OK.

John received an 'Error reading MFT#16-38' message, indicating where the corruption lay—clicking OK provided a virtual reconstruction of the underlying NTFS partition, complete with a list of all missing files and folders.

The free version of DMDE allows you to recover up to 4,000 files per session from selected folders—select each folder in turn, then press Ctrl + U. Follow the prompts to select all the files within that folder, then choose where to recover them to before clicking OK.

If you'd like to select everything on the drive and recover it in one go, then purchase a license. DMDE Express costs just $20 for a one-year license, or you can purchase a Standard lifetime license for $48.

Once your files have been recovered, you can then simply create a new NTFS partition using Disk Management, transfer the files back if required and continue to use the drive as before. However, before pressing the drive back into service, use a tool like HDDScan (https://hddscan.com/) to make sure it isn't showing any signs of imminent failure.

Launch apps together

After following your advice last issue and adding yet more apps, MetaX and RenameMyTVSeries, to my media-editing toolkit, I've realized it's now a real chore having to open them one by one when I come to use them. I've no desire to have them all start with Windows as I use them once a week, so is there another way I can automate the process of launching them with a single click? **—Marco Eberhardt**

THE DOCTOR RESPONDS:

The Doc has just the tool for this—a free portable program called Splat. Go to www.dcmembers.com/skwire/download/splat/ to download and run it. First, click the 'Create new

↘ submit your questions to: doctor@maximumpc.com

profile' button underneath the File menu and give your profile a suitable name. Click OK, then right-click the profile and choose 'Edit launch hotkey' to assign it a keyboard shortcut.

Press the [Insert] key or choose 'Edit → Insert Entry' to create your first action. Leave 'Run (if not running)' selected and click the '...' button next to the file icon to select the first program to launch. Add /user:Administrator to the Arguments box if you need to launch it with administrator privileges. If you don't want to clutter your desktop with another open window, set the Launch state to Minimized. Click Save when done.

Repeat the process for any other files or programs you wish to launch. Actions are executed in the order they are created—to introduce a delay after one is executed, add an action using the 'Wait/delay' action to set the pause time. Test your action at any time by pressing the shortcut combination key you created in step 1. Drag and drop actions to change their running order, and examine other available actions too, such as closing programs or starting services.

Disk errors in Linux

I recently added a generic ASM1062 PCI-E SATA card to my Ubuntu-powered server to add an extra drive, but since doing so, I keep getting errors that are linked to the 3TB drive (dev/sdc1) attached to the PCI-E card.

The errors aren't consistent—sometimes they appear on boot and at others, they pop up later. I suspect they're linked to attempts to write to the drive, although the data appears to be on there. None of my other drives (all 6TB) are affected. Is this likely to be a faulty card? **—Harry A Ellis**

THE DOCTOR RESPONDS:

Harry was getting two different sets of errors. The most critical were a slew of read/write errors, indicating a problem with the drive. They were accompanied by the occasional WRITE FPDMA QUEUED and related ATA bus errors.

We suspect the two errors are related. Harry had already tried two fixes: replacing the SATA cable and changing the port into which he plugged the drive on the two-port PCI-E card, all to no avail. Our priority was to get a SMART reading on the affected drive to see if imminent drive failure might be the underlying cause. This is done in an Ubuntu Server environment with the help of a package called smartmontools:

```
sudo apt install smartmontools
```

Once installed, use the smartctl tool to discover a drive's SMART status:

```
sudo smartctl --all /dev/sdc
```

This command should output a comprehensive SMART readout, but Harry reported another error message: "A mandatory SMART command failed: exiting. To continue, add one or more '-T permissive' options". We ran smartctl on the other drives that were all connected directly to the motherboard's four onboard ports. Unsurprisingly, these yielded no such errors and delivered the comprehensive readouts we were expecting.

Given the problematic drive was connected to the PCI-E card, we asked Harry to swap connections with another of his internal drives after powering down his server. This can be done by swapping cable connections between drives or, in Harry's case, making use of his server's hot-swappable drive bays. After rebooting, he got a full SMART report from both the troublesome 3TB drive and the 6TB drive he'd swapped it with.

Miraculously, the write errors also ceased. We asked Harry to check the Logs for warnings as well as errors—sure enough, the 3TB drive required a disk check using e2fsk. This can't be done from within Ubuntu Server; instead, boot into a live Linux environment using a bootable Ubuntu flash drive. Once at the Ubuntu desktop, launch GParted to verify the drives are assigned the same file (/dev/sd?) as they are in Ubuntu Server. Once done, open a Terminal window and issue the following command (/dev/sdd1 was the new location of his 3TB drive):

```
sudo fsck -f /dev/sdd1
```

The check was completed with no errors, and Harry rebooted into the Ubuntu Server proper. It seemed the 3TB's write error problems were fixed but the original ATA bus errors showed up again. A deep dive online revealed these could be due to a faulty SATA cable, lack of power, or the controller.

Having eliminated the cables, we examined Harry's PSU, an InWin 265W PSU (www.in-win.com/en/pc-power-supply/ip-s-series-au/USA). A closer examination of both that and his installed drives revealed there should be plenty of power to spare across both +12V and +5V rails. Harry also confirmed that all power cables were connected.

This left us pointing firmly at the controller. While Harry could try replacing the PCI-E card with a known brand, such as Startech's ASM1061-based card ($48, www.startech.com/en-us/cards-adapters/2p6g-pcie-sata-card), there's no guarantee it will work. We came across a reference to users reporting the errors occurring with drives plugged into the second SATA port on the controller, so we asked Harry to switch cables back to port 1, which seemed to fix the problem.

If it comes back, Harry may want to disable NCQ (Native Command Queuing), though this may come at the expense of performance.

Early access to features

I fancied trying out some of the new features in Windows 11 22H2 but was disappointed to find many, such as File Explorer tabs, weren't available. I've read that Microsoft often splits its releases into A/B groups, with some getting certain features first. How can I get myself into the right group for accessing the tools I want to preview as opposed to what Microsoft thinks I should? **—Kenny Upton**

THE DOCTOR RESPONDS:

While the A/B feature is good for testing purposes, it can be frustrating missing the cut if you upgraded for one specific feature. If you're willing to take risks, take a drive image using a tool like Macrium Reflect Free (www.macrium.com/reflectfree.aspx) then you can bypass Microsoft's decision-making measures with the help of a free tool called ViVeTool-GUI.

Go to https://github.com/PeterStrick/ViVeTool-GUI/releases to download the latest version as a portable tool or installer. Once installed or extracted, open the tool (ViVeTool_GUI.exe) and if prompted, click 'More info' followed by Run anyway.

Once administrator access is granted, it should scan for the latest available set of features. After this seemingly inexhaustible list loads, identify the feature you want to switch on, then enter its name into the Search box.

Enabling tabs in File Explorer, for example, is handled by a feature called TIFE (feature ID 37634385), while enabling the Task Manager entry when right-clicking the Taskbar can be found under Servicing_TaskbarTaskManager (feature ID 36860984).

Once located, select the feature, click the Perform Action button and choose 'Activate feature', then reboot to see if the feature is now available.

THE SLICK AND SUBTLE 4090 RIG

Building with the best for less

MORE OFTEN than not, we try to use our head, rather than following our heart. Last issue, we blew the budget on an RTX 4090 build, but this issue, we're keeping that GPU and pairing it with much better value components. That includes an Intel i5-13600K CPU, a much cheaper CPU than its top-end i9-13900K cousin, but not one we're expecting to hobble the Nvidia GPU in any games. We've also tried to keep costs down in other areas—the idea being you can get up and running with the best GPU currently available, and still be able to upgrade your rig down the line.

Using the NZXT H7 case and without RGB (except a little splash on the GPU), this system is a minimalist approach to an RTX 4090 build— certainly not the typical showboat PC you'd expect when including the current top GPU monster. We're not knocking the overall design and aesthetic here, though, as the H7 case from NZXT is a reputable and strong chassis to use for such a PC.

Because we're using a top-end GPU, this build will still be in the mid-range price bracket. But it's considerably cheaper than a most other RTX 4090 builds. We haven't chosen a fancy cooling solution, instead opting for the stock Intel cooler. We've also used five NZXT stock fans to keep the temperature down, and made the decision to incorporate a Z690 motherboard to save money over using a newer Z790 variant. Based off our tests last issue, there shouldn't be a performance hit.

We'll be putting this build head-to-head with the aforementioned machine, and we're excited to see what differences there are. Let's get cracking with the build! –SAM LEWIS

BENCHMARKS

Part		Price
Case	NZXT H7 Mid Tower	$130
Motherboard	MSI MAG Z690 Tomahawk Wi-Fi - ATX DDR5 LGA1700	$280
CPU	Intel Core i5-13600K	$330
CPU Cooler	Intel Laminar RM1	$N/A
GPU	PNY XLR8 GeForce RTX 4090 Gaming VERTO EPIC-X RGB 24GB GDDR6X	$1,600
Memory	32GB (2x 16GB) Corsair Vengeance DDR5-4800MHz	$155
SSD	1TB WD_BLACK SN850 w/o Heatsink PCIe 4.0 NVMe SSD	$230
PSU	NZXT C1000 1000W 80+ Gold PSU	$180
OS	Windows 11 Home 64-bit OEM	$32
Total		**$2,937**

PRICES CORRECT AT THE TIME OF PRINTING

BUILD IT!
Step-by-step guide to assembling this PC **PG. 22**

WHAT'S IN THE BOX

CPU
Intel Core i5-13600K

The 12600K from the last generation was an impressive chip for the price. So with the release of the 13600K, we had high expectations for the Raptor Lake variant. It packs 14 cores using a hybrid architecture splitting the cores into six performance foreground task cores and eight efficiency cores for things behind the scene. It's a speedy CPU with enough power to tackle the majority of tasks you throw at it with a max turbo frequency of 5.10GHz too. This should be a CPU that has enough grunt to keep up with the RTX 4090, but on a budget.

Price: $300, www.intel.com

MOTHERBOARD
MSI MAG Z690 Tomahawk Wi-Fi

Sure, we could have opted for a cheaper motherboard here to keep costs down further, and there's even a DDR4 variant of the same mobo. Yet, when we were deciding on where to spend our budget wisely, in terms of longevity, we picked a board with great compatibility for the next generation of components.

We've chosen the DDR5 option for this reason and also to accommodate our Corsair Vengeance sticks. However, we have stuck with the Z690 chipset instead of the latest Z790 to save some extra bucks too as the Z690 chipset is still a great platform to build on.

Price: $280, www.msi.com

CPU COOLER
Intel Laminar RM1

To cut costs but keep our CPU cool, we decided to steer away from an AIO. Although that option would make for a much better cooling solution and look more striking in the build, the stock cooler still performs well, although we wouldn't use it on a higher-performing more resource-hungry CPU such as the i9-13900K. Intel's latest stock coolers also have visually improved to be on a par with AMD's alternative offerings. For a while, Intel's looked outdated.

To fill the gap of an AIO CPU cooler, we've introduced three NZXT Aer F 120mm fans found in some spare NZXT cases we had in our studio. These are around $18 each and will pull in plenty of air sitting at the front of the case.

Price: $N/A, www.intel.com

BUILD IT! Step-by-step guide to assembling this PC **PG. 22**

CASE
NZXT H7

We'll be building this rig in the minimal-looking NZXT H7 case. As we've mentioned before, this is our least favorite of the three H7 case variants. The H7 Airflow offers the best airflow (unsurprisingly) and the H7 Elite is the showstopper of the selection with its RGB and glass front paneling.

However, the minimalist style suits our system the best and it's still a well-built clean design. The build quality is superb, exactly what you'd expect from NZXT, and there are plenty of great cable channeling options to keep things nice and organized.

Price: $130, www.nzxt.com

PSU
NZXT C1000 1000W 80+ Gold PSU

For this system to run properly, it's recommended to use a PSU of at least 1,000W. After all, if you're spending that much money on the GPU, it's better to be safe than sorry. As we said earlier, our GPU is power-hungry, requiring either three or four individual PCIe cables per graphics card. We'll touch on this later as it was a slight inconvenience but the bottom line is this C1000 has enough capability to safely manage this power and was a necessity for the build. Plus it's rated at 80 Plus Gold efficiency. Price: $180, www.nzxt.com

GPU
PNY XLR8 GeForce RTX 4090 Gaming VERTO EPIC-X RGB 24GB GDDR6X

Here it is, the star of our rig. This is a huge combination of metal, plastic, and other wizardry that somehow results in some jaw-dropping eye candy. Yet, it comes with quite the price tag—$1,600 if you're lucky enough to grab one at retail price. It features a whopping 24GB of GDDR6X memory and it packs a mighty punch. Of course, the RTX 4090 is currently the most powerful graphics processor on the market, and this PNY XLR8 is a great partner card. However, it's power-hungry so you'll likely need a new PSU to accommodate this behemoth. Price: $1,600, www.pny.com

SSD
WD_BLACK SN850 w/o Heatsink PCIe 4.0 NVMe 1TB SSD

We'll be honest here, we couldn't get an SN850X in time for our photo shoot, due to some issues with the delivery. Thankfully, we had a spare SN850 available, so that found its way into our build instead. That said, this SSD is hardly a slouch and can easily be swapped out or upgraded down the line.

This drive has a 1TB capacity, which is a healthy amount of storage, and it also has a speedy PCIe 4.0 interface. We eagerly await PCIe 5.0, but the WD's speeds are impressive—with up to 7,000MB/s read and up to 4,100MB/s write. Those figures certainly float our boat.

Price: $105, www.westerndigital.com

BUILD IT! Step-by-step guide to assembling this PC **PG. 22**

MEMORY
Corsair Vengeance DDR5 @4800MHz (2x 16GB)

We're glad that DDR5 no longer costs an absolute fortune to attain. Of course, proce drops were inevitable, but it's makes a lot more sense to now choose it over DDR4. We haven't got the quickest speeds in comparison to other DIMM sticks of this latest generation, but 4,800MHz is nothing to turn your nose up at. These are some sleek sticks that easily match the subtle look of the build, though we'd have preferred the white variant (which are sold by Corsair) to match our case in this instance.

Price: $155, www.corsair.com

A WOLF IN SHEEP'S CLOTHING?

LENGTH OF TIME: **1-2 HOURS**

LEVEL OF DIFFICULTY: **EASY**

BUILD IT!
Step-by-step guide to assembling this PC

1 THE STRUCTURE

TO MAKE our build run smoothly, it's better to take things apart right from the start rather than some time further down the line. This way, we can plan cable routes and it makes it easier to install and change components around as we build. NZXT cases are easy to dismantle and this one, in particular, doesn't even feature any thumbscrews to release the front, back, side, and top panels—they simply pop off with a tiny amount of force. That's pretty welcome in this case, as the last thing you want is a smashed glass panel (believe us, we've done it). When you pop the panel, gradually lift it out and up to remove it. We also removed the filters to gain access to the fan mounts.

3b HEATSINK COVER

REMEMBER we put a heatsink to the side? We need it now. As we have opted for the stripped-down cheaper version of the SSD without the pre-installed heatsink, we need to protect it and cover it with the motherboard's offering. To install, peel back the protective covering, line up the screw holes and gently rest it in place. Once happy, tighten it to secure it down.

2 INSERT THE BRAIN IN THE MAINFRAME

A PC WITHOUT a CPU is like a car without an engine, so to kick things off, we'll be adding the i5-13600K onto the motherboard. Before grabbing the CPU, we need to prepare the mobo by opening up the CPU bracket. Unhook the retention bar by pushing down on it and slowly guiding it sideways to release it. Then lift this to open up the socket.

After that, we need to make sure we align the arrow on the corner of the CPU with the arrow on the motherboard and slowly place it on the socket, making sure not to bend any of the pins. It's always recommended to hold the CPU by the sides rather than the bottom. Lastly, rest the CPU bracket back down and tuck the retention bar under the hook to secure it into position.

3a THE TOP SLOT

WHEN INSTALLING any PCIe interface, always try to use the uppermost available slot on the motherboard—that's what we have done here for our SN850 SSD. Using a small screwdriver to begin with, unscrew the heatsink on the motherboard. On our specific board, this was labeled number 1. We'll just put this to one side for now, as we will need it later.

Tilting it at a 30° angle, slot the SSD into the port and lower it down. Thankfully, this board has a rotating plastic screw clip that holds the SSD down, no need for that awkward tiny screw that we always fear we might lose here.

4 SATISFYING CLICK OF CONFIRMATION

IF YOU HAVE four available DIMM slots and aren't using all of them, give your motherboard manual a quick check to see what slots you should use first. In our instance (and most instances), this is slots 2 and 4, so we pushed back the clips on the top and bottom of these slots to open them up. The RAM sticks will only go in one way, so make sure to line up the notch on the motherboard with the cutout on the stick and push down to fix in place.

You should be greeted with a glorious click to let you know that it's in the right position. We'd have loved some white RAM sticks to match this minimal build, which we didn't have, but Corsair does sell. The front-facing fans act as our intake with the top fans as exhausts, so make sure the latter face the inside of the case.

5 COOL DOWN TIME

BEFORE WE GET ahead of ourselves and add the motherboard, we have to finish things off with our stock cooler. This is far easier than installing an AIO with a radiator as these can get fiddly. Stock coolers are a breeze to install, especially so with this Intel Laminar RM1. You don't even need an LGA1700 back bracket on the motherboard as there are no screws or standoffs needed.

First, add a pea-sized amount of thermal paste to the center of the CPU. Then align the CPU cooler on top and guide the pins into the motherboard holes. To fix it down, turn and press each corner down and you'll hear another satisfying click once it's in. Tighten from one corner to the diagonal corner to maintain an even pressure and to make life a little easier for yourself.

6 THE ANCHOR

ONE LAST THING we need to do before placing the motherboard inside the chassis is to add the PSU to the case. This step doesn't necessarily need to be done now, it just has to be in before we connect all of the cables. Installing the PSU is one of the easiest steps in most builds and only requires four case screws. Working around the back of the case, slide the PSU into position with the fan facing downwards. Align it with the screw holes on the case and tighten it, again working diagonally. To give us more room for the PSU cables towards the end of the build, we also removed the HDD cage as we won't be needing this.

7 INTO THE MOTHERSHIP

LET'S GET CRACKING with the rest of this build. We'll start by installing our motherboard into the case. First, we rest the chassis down on the back panel to make the job easier. Once it's laid flat, carefully pick up your motherboard by the sides and place it onto the standoffs in the case with the I/O shield going in first. This sits in the rectangular gap at the back of the case. Once in place, you can start using the standoff screws that came with the case to secure the board down. We'll be screwing it down on every standoff except the bottom right and the standoff that sits above it. We'll need these to support our GPU, believe it or not!

10 FRONT I/O CABLES

BEFORE WE take up a lot of room with our GPU, we should get some cables connected to our motherboard, starting with the fan and front I/O cables. Thankfully, our motherboard has plenty of headers so we don't need a daisy chain to use all of our case fans. Most of these sit along the bottom of the board so we ran the cables around the back and fed them up from the bottom to keep things organized. We then connected the USB cables and audio cables to the motherboard as seen in our step image below. Again, we kept on top of cable management here by tucking the wires behind some of the paneling of the case.

11 4090 O'CLOCK

NOW THAT ALL of the necessary cables around the bottom of the motherboard are connected, we can fill this space with the party piece of this system, our GeForce RTX 4090. It's best to keep the case flat on its back when installing this, especially considering the hefty weight of it. First, we need to remove some of the shroudings on the rear of the case, which easily screw off using a Phillips screwdriver. Make sure you don't lose these screws as we'll be using them for the GPU.

We took off the second, third, and fourth shrouds from the top. Then carefully line up the I/O panel with the rear of the case and the PCIe connector with the motherboard, and push it down gradually until it clicks into place. Once in, screw it down on the case using the screws from the shrouds. We can then adjust the anti-sag bracket and tighten that to keep the GPU upright.

8 ANTI-SAG BRACKET

SO THIS IS the much sleeker alternative to a standard GPU bracket that came with our PNY XLR8 RTX 4090 GPU. Instead of attaching directly to the case where the GPU mounts, it connects to the case through the motherboard using two large standoffs. So to start with, using the two larger standoff screws in the kit, screw these into the two remaining holes we left on the motherboard. Then we will rest the bracket over these and secure it on using the tiny screws in the pack.

The last thing we have to do here is to connect the L-bracket to the main bracket, which is what will prop up the GPU. As we clearly don't have the juggernaut-like GPU in place right now, we can only loosely tighten it in place here. We added it on the fourth screw hole up from the bottom with the foam padding facing upwards, as you can see from the red O-ring in the image below. When our GPU is in place, we can raise and tighten this bracket so it holds the graphics card level.

9 NO FANCY COOLING HERE

YES, THAT'S THE TRUTH, there's no water-cooled or AIO magic going on here, just some old-fashioned fans. We've opted for quantity over quality here, and to bring as much air into our case as possible, we've added three spare NZXT Aer F 120mm fans from other NZXT cases we had in the labs.

These are easy to install, and they simply use four case screws to secure onto the front fan bracket. To help with cable management, we positioned the fan cables so they face toward the back plate of the case. It's also important to remember that faces suck, so position the face of the fan outwards so it pulls air inwards towards your components.

As we added these three fans to the front, we moved the initial front fan that came in the H7 case to the top. This means we have three intake fans and two exhaust fans for our build.

12a CABLE CRAZINESS

ALL THE TECHNOLOGICAL greatness that comes with our mighty GPU unfortunately also brings some negatives. One of which is the PCIe power situation. This specific GPU requires four—yes four—PCIe power cables that connect to a new adapter included with the graphics card. Unfortunately, each connector needs its own individual power cable, which means there are a total of four connectors needed from the PSU. To help with organizing the cables in the case, we pulled the spare PCIe connectors and cable-tied them to make everything as tidy as possible.

12b INTO THE GPU

AS YOU CAN SEE, it's not the cleanest of solutions, but it's the best we could do with what we had. Thankfully, there was a convenient hole in the case where we could thread our big chunk of PCIe cables through. It's also a good job we have the anti-sag bracket in place as the combined weight of these cables is pretty hefty.

BUILD IT!
Step-by-step guide to assembling this PC

13 PSU CABLING.

OH, THE PROBLEMS from the 4090 don't end there. In fact, this is where we came across a major bump in the road. Our PSU, although capable enough of running this build, had only five available CPU/PCIe ports on the back. We've just used four of these in connecting the GPU, leaving us with only one for the CPU. However, our motherboard had two dedicated CPU ports on the top to use—ideally, when a motherboard has an option for two, you should use both of them.

Yet as our CPU isn't too power-intensive, we've only used the first CPU port as we didn't really have an alternative. We've done this so you don't have to, because getting hold of an ATX 3.0 PSU with all the available six CPU/PCIe ports at 1000W is currently just as hard as getting hold of a 4090 in the first place. So, we've done the best we can with the next best thing. As for the rest of the cables, thankfully they gave us no issues. Always refer to the manual when installing the PSU cables just to ensure you are installing them correctly.

14 THE FINAL TOUCHES

USING NZXT'S cable routes was a great way to keep this build tidy around the back, in case of any further component changes down the line. This case has allocated channels to guide your cables through and hide a lot of the mess from the front. We've used as many of them as we could, but still have a few cables that weren't quite long enough to run along these routes. The main thing here is to take your time choosing where you want your cables to go and to make the front of the case as clean as possible.

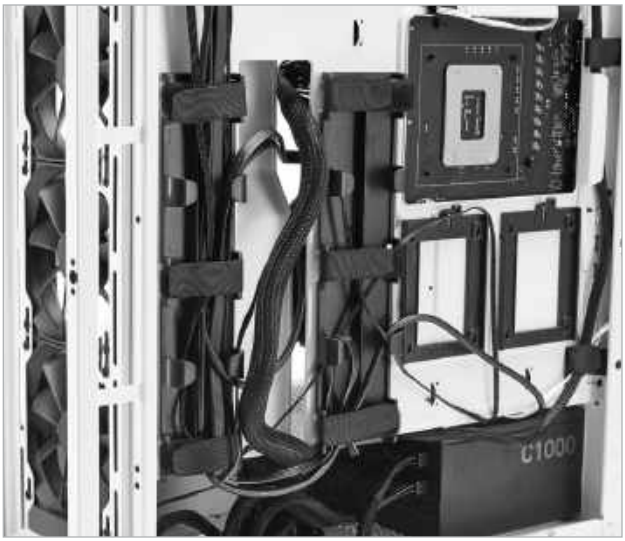

15 BIOS FLBK UPDATE

WE HAD OUR FINGERS CROSSED that everything would boot up after getting it all connected. So it was a little disappointing when the machine wouldn't post—at least it turned on all of the fans and the GPU. Due to the age of the motherboard and using a brand new CPU, we had to update the BIOS, something we always do as a matter of course anyway. Our MSI motherboard has a dedicated BIOS FLBK USB port and button to update the BIOS without the need to enter the BIOS manually. This again is handy seeing as we have more chance of adding another 4090 to this build than it posting at this stage.

To download the update, use another computer or laptop and format a clean USB drive to FAT32. To do this, right-click on it in file explorer and select format. Within this pop-up window, under file system, select FAT32 and click start. Back up your drive first, as formatting will wipe any files on there. After you've formatted, head to the motherboard manufacturer's website and go to the support page of your specific mobo. Here you'll find a download section, so grab the latest BIOS version.

For the MSI file, you have to find it and extract the file. Within the extracted folder, ignore the txt file, as you will need to rename the other file on the drive to 'MSI.ROM'. Once renamed, you'll get a prompt asking if you're sure you want to change the file extension, select OK and move this file to your clean FAT32 USB drive. It needs to be at the root of the USB drive so make sure there's nothing else on it.

Now we can safely eject the drive and put it into the BIOS FLBK USB port on our new build. With the power supply switched on, press the flash BIOS button once and you should see a small red light start to flash indicating that it's reading the USB stick. The system should then boot up and it will begin updating. The whole process should take around 10 minutes and your system will boot up and restart a few times during this. When the light behind the FLBK button goes out, everything should be finished.

16 GET GOING WITH WINDOWS 11

BEFORE PUTTING AWAY the other system we used to download the BIOS and the USB drive we used to install it, we still need them for the operating system. Head over to the Download Windows 11 page on the Microsoft official website. On this page, choose the second option labeled 'Create Windows 11 Installation Media'. This download will allow us to create a bootable USB drive with Windows 11 on it. Use the USB drive from earlier, though it needs to be at least 8GB in size, wipe it clean and make sure it's still formatted to FAT32.

After the download has finished, open it and run through the program to get ready to install it onto the USB drive. When you get to the 'Choose which media to use' page, make sure to select the USB flash drive option. Click the clean USB drive from the list and then hit next. This may take some time so kick back for a while, watch an episode of your favorite TV show, take your dog for a walk, or do whatever else you need to do to kill a bit of time. After this has finished, eject the USB drive safely and insert it into a USB port on the motherboard of the new build. Then boot your PC and it should automatically load into the installation wizard for Windows 11.

If it doesn't, then while the PC is booting up, mash the delete key to enter the BIOS. Find the boot priority section on the BIOS and re-order this so that your USB drive with Windows 11 on it is at the start of the list. Then save and exit and reboot your system. Hopefully, now you can get to the Windows 11 installation wizard and follow all the necessary steps to set up Windows to suit you, and then you can get cracking with your system.

BUILD IT!
Step-by-step guide to assembling this PC

17 UPDATES AND DRIVERS

BEFORE WE CAN dive headfirst into 4K RTX 4090 glory, we have to make sure our system is running on the latest software and has all the necessary drivers to get things running smoothly. These drivers and updates ensure any compatibility and security issues are fixed, giving us peace of mind that our PC is running properly.

We always head straight to the Windows 11 settings first and check if there are any OS updates. If so, install them and restart your system afterwards. Next, head over to the specific motherboard page on the manufacturer's official website, go to the support page and install the latest Intel chipset driver. Open this file and follow the onscreen instructions, you'll likely have to restart your system too, once installed.

One handy tool to ensure everything on the Intel side of your machine is running well is to install the Intel Driver and Support Assistant. At the risk of sounding like a broken record here, we recommend opening this and letting it check through your system. Install any updates that might be available, which tend to be Wi-Fi and Bluetooth driver updates.

The last important driver to install is for the PNY XLR8 GeForce RTX 4090 we have. You'll need to download the GeForce Experience application from Nvidia. In the application, click on drivers in the top left-hand corner and check for updates. Then install it and follow the onscreen instructions to update to the latest graphics card driver. After this, restart your computer again to make sure the installation is complete. Now, finally, we should be able to start having some fun with this new rig.

BUILD IT!
Step-by-step guide to assembling this PC

1 As our RTX 4090 is one resource-heavy component, it requires a lot of cables. It's a shame the new adapter isn't longer so that one singular cable could run to the bottom of the case so we could hide the bundle under there.

2 The white variants of this fan would have looked better to match the rest of the machine. In any case, they are hidden by the solid front panel.

3 Although the Intel stock fan aesthetically fits the build, an AIO CPU cooler would be better suited for the power we have on offer here.

4 To help with the cooling duties, two additional exhaust fans could be added along the top to properly fill this case out.

SUBTLE 4090 BUILD BENCHMARKS

	ZERO POINT	
Cinebench R23 Single-Core (Index)	1,691	1,994 (18%)
Cinebench R23 Multi-Core (Index)	28,692	20,687 (-28%)
CrystalDisk QD32 Sequential Read (MB/s)	7,089	6,967 (-2%)
CrystalDisk QD32 Sequential Write (MB/s)	6,743	5,304 (-21%)
3DMark Fire Strike Ultra (Index)	12,785	24,250 (90%)
Cyberpunk 2077 (fps)	48	87 (81%)
Cyberpunk 2077 RTX (fps)	23	46 (100%)
Metro Exodus (fps)	77	123 (60%)
Metro Exodus RTX (fps)	52	88 (69%)
Total War: Three Kingdoms (fps)	61	103 (69%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of the "Ultra" 4K build from our July 2022 issue, featuring an AMD Ryzen 9 5950X, Zotac Gaming GeForce RTX 3090 AMP Extreme Holo 24GB, MSI MPG X570S Carbon Max Wi-Fi, 32GB of G.Skill TridentZ DDR4-3200, and 2TB MSI Spatium M480 M.2 SSD. All games tested at 4K 'Ultra' graphics presets with DLSS and V-sync turned off and XMP for RAM speed turned on. No manual CPU overclocking.

13600K OR 13900K?

IS IT WORTH STRETCHING FOR THE MOST POWERFUL INTEL CHIP?

THE AIM OF THIS BUILD was to piece together a machine that includes Nvidia's latest and greatest graphics processor, but keep the costs down by cutting corners elsewhere. In this instance, that was on a more modest Intel 13th-gen chip and a motherboard that, while not cutting edge, is still equipped for the future. Likewise for the rest of the build, we chose components that we didn't think would have a major impact on performance but would crucially save us some money.

Instead of spending at least $150 on an AIO, we opted to dig around in our stock room and reuse some spare fans from unused NZXT cases. These come in at around $18 each, so for a total of under $60, this is a good solution to pair with the Intel Laminar stock cooler we used. Of course, an AIO cooler would be a more efficient and effective way to keep the temperatures down, but this alternative is good for the time being if your budget doesn't stretch to an AIO when you start building your PC.

Other than our mishap with a lack of CPU/PCIe ports on the back of our PSU, the build went smoothly and resulted in a system that runs impressively. To avoid making the same mistake we did, we highly recommend trying to find an ATX 3.0 PSU if you're going down the 4090 route. One of these PSUs should have plenty of longevity with this and future generations. Another thing to note about this build is that it doesn't run silently, but that's just the nature of our cooling solution. Be sure to grab your headphones!

We were also hesitant about our CPU being a bottleneck for the GPU. This wasn't the case, however, and we were more than happy with how the CPU and GPU worked in synergy.

It only takes a quick glance at our benchmark table to demonstrate that this is a monster of a machine—even compared with the 8K gaming monster build that Guy put together in our previous issue (Holiday 2022). That housed an Intel Core i9-13900K and an MSI GeForce RTX 4090 Gaming X Trio GPU. Apart from the CPU benchmarking, the rest of the benchmarking is similar considering we managed to build this machine for around $800 less. Although let's be honest, Guy's machine does beat this one in the looks department.

The bottom line is that, from the outside, this build doesn't look like anything crazy, but that's exactly why you shouldn't judge a book by its cover. Running on ultra at 4K resolution, this sleeper of a build manages to achieve over 85fps on triple-A titles such as *Cyberpunk 2077*, and trust us, these games look stunning on this machine.

So that leaves us with one last question—should you opt for the 13600K over the 13900K? If your main priority is games then you can absolutely go for the cheaper 13600K and Z690 motherboard to achieve impressive results and save a bit of money. Overall, we love this minimalist gaming rig and hope it's shown you don't have to blow the budget to reach gaming paradise.

SUBTLE 4090 BUILD VS. 8K GAMING MONSTER BENCHMARKS

	8K GAMING MONSTER	
Cinebench R23 Single-Core (Index)	2,247	1,994 (-11%)
Cinebench R23 Multi-Core (Index)	35,457	20,687 (-42%)
CrystalDisk QD32 Sequential Read (MB/s)	7,057	6,967 (-1%)
CrystalDisk QD32 Sequential Write (MB/s)	6,743	5,304 (-21%)
3DMark Fire Strike Ultra (Index)	24,528	24,250 (-1%)
Cyberpunk 2077 (fps)	78	87 (12%)
Cyberpunk 2077 RTX (fps)	43	46 (7%)
Metro Exodus (fps)	140	123 (-12%)
Metro Exodus RTX (fps)	105	88 (-16%)
Total War: Three Kingdoms (fps)	103	103 (3%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of the "8K Gaming Monster" build from our Holiday 2022 issue, featuring an Intel Core i9-13900K, MSI GEFORCE RTX 4090 GAMING X TRIO 24GB, ASUS ROG MAXIMUS Z790 HERO, 32GB of CORSAIR VENGEANCE DDR5-5200, and 1TB PNY XLR8 CS3140 M.2 SSD. All games tested at 4K 'Ultra' graphics presets with DLSS and V-sync turned off and XMP for RAM speed turned on. No manual CPU overclocking.

PRODUCTIVITY STARTER TOOLS

WANT SOME HANDY WEBSITES AND APPS TO GET YOUR CREATIVE JUICES FLOWING?

SO, YOU'VE JUST BUILT YOUR NEW PC, you have a catalog of games that are sitting waiting in the download queue ready to be installed, and life is looking pretty good. But games aren't the only thing you should be using your PC for. There'll be plenty of times when you'll need a productivity suite for work, a personal project, or just to have some fun. And what's better than productivity applications, we hear you ask? Free ones, of course! Here are some of our favorites to add to your newly-constructed machine.

GIMP

IT'S NOT THE GREATEST name for a piece of software, but GIMP stands for GNU Image Manipulation Program, which explains exactly what this piece of software does. GNU means free software ('free' referring to freedom within the application as in open-source, not free as in price—although this is) As for the rest of the name, that's much more self-explanatory, it's a powerful and detailed tool to edit images.

Although GIMP isn't as polished as the likes of Adobe Photoshop or Affinity Photo 2, it's a capable piece of software that will make a great addition to your new PC. It does much of what the other two image manipulation applications can do without either a hefty price tag or subscription fees.

GIMP is a great starting program to bring out your inner creative. You can certainly learn a lot about photo editing from just playing around, and who knows, your hobby could even earn you some money. $free, www.gimp.org

GOOGLE OFFICE SUITE

IF YOU HAVE A GOOGLE ACCOUNT, you have access to a plethora of great services at your fingertips contained within Google's productivity suite. These are essentially web-based tools and alternatives to Microsoft's Office 365, which many people will be familiar with. Google Docs, Slides, and Sheets are direct alternatives to Microsoft Word, PowerPoint, and Excel, except you don't need to pay $7 per month or more to Microsoft.

Google, like Microsoft 365, is cloud-based with offline functionality for when you don't have an internet connection. All your files are easily accessible from PC to mobile and can be saved to Google Drive, so you can easily adapt and work on the go, which is perfect for staying productive. Each application has plenty of functionality and like the rest of Google's products, the design is intuitive. And if you have previous experience using Microsoft Office, you'll no doubt pick things up even quicker. $free, www.workspace.google.com/features

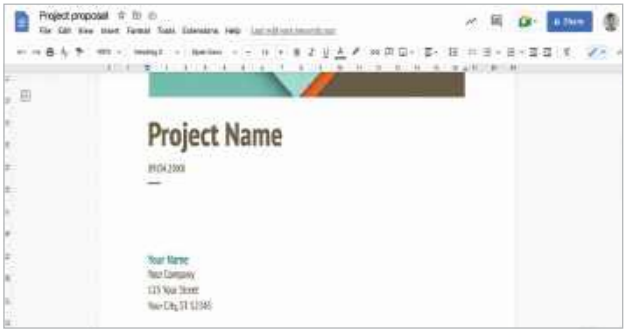

PDF CANDY

EDITING PDF FILES is one of life's great annoyances, but this web-based application makes it quick, easy, and free. There's a catch with the free version, in that you can only edit one PDF per hour, but with PDF Candy Pro, there are no limits and the file size per task increases to 500MB. At $6 per month or $48 per year, this is a great tool at an affordable price. Even with the free version there are plenty of tools, including compress, edit, merge, delete pages, watermark, crop, rotate, resize, and more.

PDF Candy also offers an array of conversion tools, so you can convert files such as Word, Excel, jpg, docx, png, and others to PDF. You can also do the reverse conversion by changing PDFs to other formats including Word, Powerpoint, Excel, jpg, tiff, and more, so there's plenty of capability here all in one site. It's also easy to navigate and understand, making it a must-have for your new PC. PDF Candy: $Free, PDF Candy Pro: $6 per month, $48 per year, www.pdfcandy.com

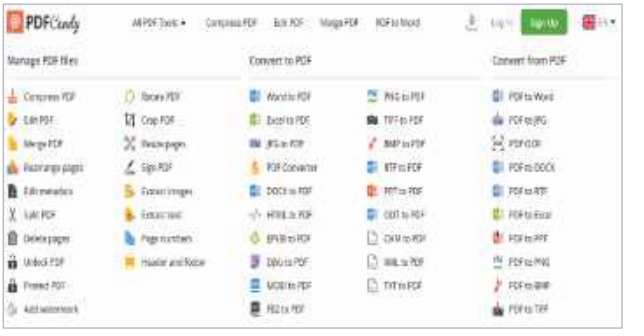

UNSPLASH

WHEN CREATING ANY DESIGN WORK, whether it's for personal use or in a commercial environment, finding the right image—and making sure you're legally covered to use it—can take up a lot of time. Unsplash handles both of these issues in a tasty-looking, free-to-use, stock image library website, which has been linked to Getty Images since 2021. Unsplash is a community-driven platform where photographers upload their work to the site for others to use creatively under the Unsplash license. All photos can be downloaded and used for free. No permission is needed, though credits are always appreciated.

These factors make Unsplash a great site for creatives. When you download an image, you have multiple options with different file sizes to choose from, and the creator's link pops up afterwards, so you can correctly credit their work. Pair Unsplash with GIMP and you can practice your design and image manipulation skills for free. $free, www.unsplash.com

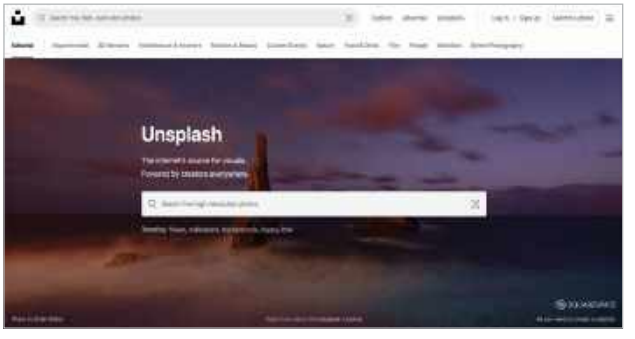

NOTION

IF ORGANIZING ISN'T YOUR FORTE, then Notion is a no-brainer. More than simply a note-taking app, Notion is an all-in-one hub where you can project manage and keep track of databases, calendars, and checklists. There are plenty of templates to choose from to plan your creative ideas and share them with others. The platform also has cloud storage, so files can be accessed wherever you are—always a plus for productivity apps.

The note-taking side of the app offers lots of customization to help you lay out your thoughts clearly and it's easy to embed images, videos, and other media into your notes. With the free Personal plan, you have access to the Notion API to build custom integrations for your team, unlimited pages, and the ability to sync files across devices. The Personal Pro plan offers the same benefits with added features, such as unlimited file uploads, unlimited guests, and a 30-day page recovery. Notion Personal: $free, Notion Personal Pro: $4 per month. www.notion.com

AMD RDNA 3 ARCHITECTURAL DEEP DIVE

Taking a dip in the pool with the upcoming RX 7900-series

Get ready for our third and final round of graphics architectural deep dives. Two months ago, we spilled the beans on Nvidia's Ada Lovelace architecture. Last month, we did the same for Intel's Arc Alchemist. Last to the party, but perhaps the most interesting in its approach in the realm of GPU design is AMD's RDNA 3, premiering with the Radeon RX 7900 XTX and Radeon 7900 XT.

Thanks to its use of chiplets, the RDNA 3 architecture fundamentally changes several of the key design elements. AMD pioneered chiplets in the consumer realm with its Zen 2 (Ryzen 3000) CPUs. That has enabled AMD to deliver more cores and better performance, giving Intel some much-needed competition. Now it's looking to do the same with GPUs and graphics cards.

It's not an easy transition to make and, of course, GPUs are fundamentally quite different from CPUs. Whereas CPUs typically come with a few dozen general-purpose cores at most—up to 96 with the latest AMD Genoa EPYC processors—GPUs already have thousands of specialized shader 'cores', each of which needs access to the large data sets involved with graphics work, such as the textures and the geometry.

So, it's time to don your swimsuit as we dive into the latest RDNA 3 architecture to see what makes it tick. –JARRED WALTON

THE GPU CHIPLET ERA BEGINS

For now, AMD has only revealed one specific design for RDNA 3, the Navi 31 GPU. It consists of two separate pieces of silicon, the Graphics Compute Die (GCD) and the Memory Cache Die (MCD). The names already provide a good indication of how AMD chose to split up the normal monolithic chunk of silicon that makes up a GPU into multiple chiplets.

The GCD houses the Compute Units (CUs) that provide the raw number crunching capabilities of AMD's GPUs. It also houses other core functionality like video codec hardware, display interfaces, texturing units, render outputs, and the PCIe connection. So far that's no different than previous GPUs, but with RDNA 2 as an example, the GPU also housed a large chunk of L3 cache—the Infinity Cache—along with memory controllers that link to external GDDR6 memory.

With RDNA 3, or at least Navi 31 (see the side panel on Future RDNA 3 GPUs), the GCD has up to 96 CUs, which is where the typical graphics processing occurs. It also has the Infinity Fabric along the top and bottom edges of the chip that provide an extremely wide and fast link to all the MCDs. It can link with up to six MCDs and, in that configuration, the total Infinity Fabric bandwidth is a staggering 5.3 TB/s.

Meanwhile, the MCDs do exactly what their name implies: Each houses a large 16MB chunk of L3 cache, along with the physical GDDR6 memory interface. They also contain Infinity Fabric links to connect to the GCD, which you can see in the die shot along the center-facing edges of the MCDs.

What's the point in splitting up the functionality between the GCD and MCDs? We're glad you asked, as it's the key point of a chiplet approach. The GCD will use TSMC's 5nm N5 node, basically, the latest and greatest full production node that's currently available—4N, N4P, and N4 are all derived from N5.

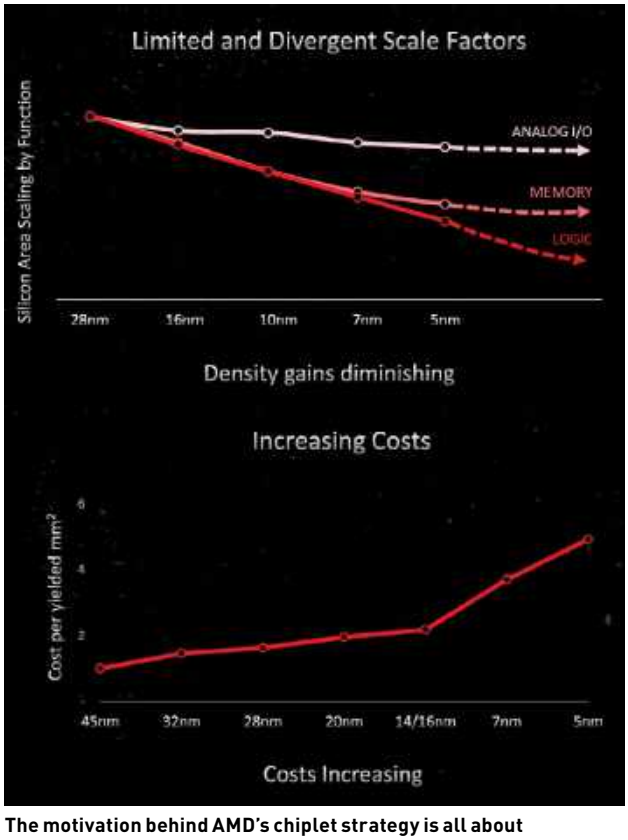

The motivation behind AMD's chiplet strategy is all about improving yields and flexibility while reducing costs.

POWER REQUIREMENTS OF INTERFACES

Interconnect	Picojoules per Bit (pJ/b)
On-die	0.1
Foveros	0.2
EMIB	0.3
UCIe	0.25–0.5
Infinity Fabric (Navi 31)	0.4
TSMC CoWoS(MB/s)	0.56
Bunch of Wires (BoW)	0.5–0.7
NVLink-C2C	1.3
Infinity Fabric (Zen 3)	1.5

The MCDs, on the other hand, use TSMC's N6, a derivative of the previous 7nm N7 process node.

There are several important reasons for the split. First and foremost, N6 is a lot cheaper than N5. TSMC doesn't publicly disclose contract pricing, but there are indications that 5nm-class wafers cost about twice as much as 7nm-class wafers. Second, the scaling of certain elements has slowed down or almost entirely halted with smaller process nodes. External interfaces like those for GDDR6 would be just as large on N5 as on N6. The scaling of cache cells has also slowed down, so 16MB of L3 cache on N5 would be nearly as large as 16MB on N7.

Beyond cost and scaling, there's another benefit to the split: chip yields. Creating a larger chip means there's a much higher potential for any given die to contain a defect—there are relatively consistent impurities with silicon wafers. To avoid having a single defect result in a useless die, large chips need the ability to disable portions of the die. That can result in more complex routing and other redundancies, which in turn adds to the cost.

The Navi 31 packs 45.7 billion transistors into a 300mm^2 die. The MCDs meanwhile each pack 2.05 billion transistors on a chip that's only 37mm^2 in size. That works out to 152.3 million transistors per mm^2 on the GCD and 55.4 million transistors per mm^2 on the MCD. And that's still including the Infinity Fabric sections on both chips—about nine percent of the GCD die area and 15 percent of the MCD die.

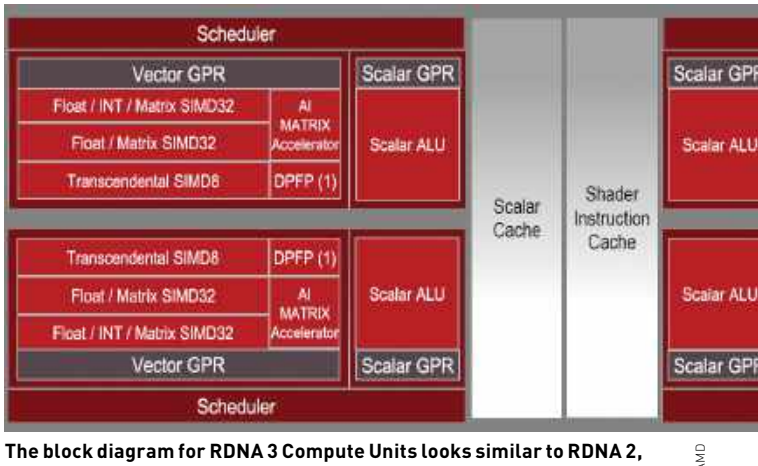

The block diagram for RDNA 3 Compute Units looks similar to RDNA 2, but Float/Matrix SIMD32 units and AI Matrix Accelerators have doubled.

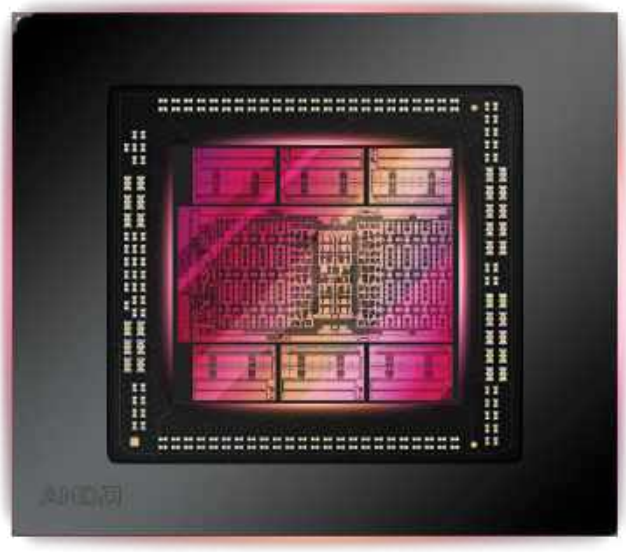

AMD's Navi 31 GPU will be the first consumer graphics chip to utilize chiplets, separating the core compute functionality from the memory controllers and L3 cache.

INFINITY FANOUT INTERCONNECT

All the chiplet talk might make it seem as though that approach is a design no-brainer, but there are areas of potential concern. On-die interfaces tend to be the most efficient in terms of power requirements, requiring around 0.1 picojoules per bit (pJ/b). External interfaces such as the Infinity Fabric link used on AMD's Zen 3 processors on the other hand can require 1.5 pJ/b or more. Scaling the CPU chiplet links up to the number used by Navi 31 would have quickly consumed a huge chunk of the power budget.

Just how wide is the Navi 31 Infinity Fabric link? AMD doesn't give an exact number but claims it transfers 2,304 bytes per clock and has a speed of 9.2Gbps. To reach the 5.3TB/s figure, that would require a 2.3GHz Infinity Fabric speed transferring four bits per wire per clock.

The CPU Infinity Fabric uses an organic substrate, but with the high bandwidth demands of a GPU, AMD wanted something far more efficient. It created a new interface that it calls the Infinity Fanout Interconnect. It says it only uses 20 percent of the power of the CPU variant of the Infinity Fabric. In fact, AMD says all the Infinity Fanout links combined only account for less than five percent of the total graphics card power budget.

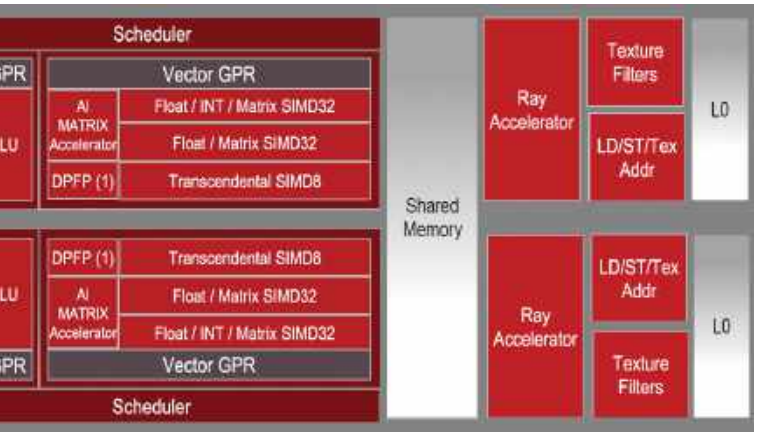

FUTURE RDNA 3 GPUs

Navi 31 has been revealed, but it certainly won't be the only member of the RDNA 3 family. AMD ultimately had four different RDNA 2 GPUs for desktops, plus the same core architecture is also in multiple other products like the Steam Deck APU, Ryzen 6000 mobile processors, and now the Ryzen 7000 family of CPUs. There are rumors of at least two more desktop graphics card GPUs based on RDNA 3: Navi 32 and Navi 33.

AMD won't publicly discuss these future parts, but at least one source revealed key specifications for Navi 31 a few months back that match up almost exactly with what AMD has now disclosed. As such, we can be reasonably confident with the details that the same source revealed about the other designs.

Navi 32 will follow the same pattern established by Navi 31, just with a smaller GCD and with support for fewer MCDs. The rumored die size is 200mm², about two-thirds as large as the Navi 31 GCD, with up to 60 CUs and the ability to link to four MCDs. That gives it up to a 256-bit interface and 64MB of Infinity Cache. Like Navi 31, it can use less than four MCDs, and we suspect we'll see versions with fewer CUs enabled and with three MCDs.

The full or nearly full GCD will likely be used for RX 7700-class products with four MCDs, possibly both an RX 7700 XT as well as an RX 7700 — or AMD might use the XTX suffix on other parts. The same GCD with fewer active cores and only three MCDs could then end up as the RX 7600-class, and again we expect AMD will do both higher and lower tier variants of this card.

Meanwhile, Navi 33 will be a monolithic die, with up to a 128-bit memory interface. Current rumors are that it will have 32 CUs and up to 4,096 shaders, with up to 32MB of L3 cache and a die size of around 200mm². It will be a mobile-first design, taking over from Navi 24 and the RX 6400/6500 class of graphics cards.

Those are the two known designs, though we can't help but wonder if AMD might consider doing another chip that's larger than Navi 31's GCD at some point. All the building blocks are there, and 300mm² isn't particularly large for a GPU. Imagine a 450mm² GCD with support for eight MCDs. Such a design could close the gap with Nvidia's RTX 4090, though there might not be sufficient demand to justify such a design.

RADEON RX 7900 SPEEDS AND FEEDS

The Radeon RX 7900 graphics cards will be the first to use AMD's RDNA 3 architecture. There will be two variants, the faster and more expensive 7900 XTX, and the slower and cheaper 7900 XT. The XTX uses the fully enabled GCD and delivers a theoretical 61.4 teraflops with 960 GB/s of memory bandwidth, while the XT disables 12.5 percent of the CUs and clocks slightly lower, with 51.6 teraflops and 800 GB/s of memory bandwidth.

Both use GDDR6 memory clocked at 20 Gbps, but there's one less MCD on the XT model. Total power is 355W for the XTX and 300W for the XT, similar to the existing RX 6950 XT and 6900 XT—20W higher on the XTX compared to the reference 6950 card, but we tested third-party cards that pulled 400W.

AMD says it designed RDNA 3 to scale to 3GHz and seems to have targeted maximum efficiency, which likely limited the clocks on the reference card. There will be overclocked third-party cards, and we're looking forward to seeing how high they can scale in performance without sacrificing efficiency.

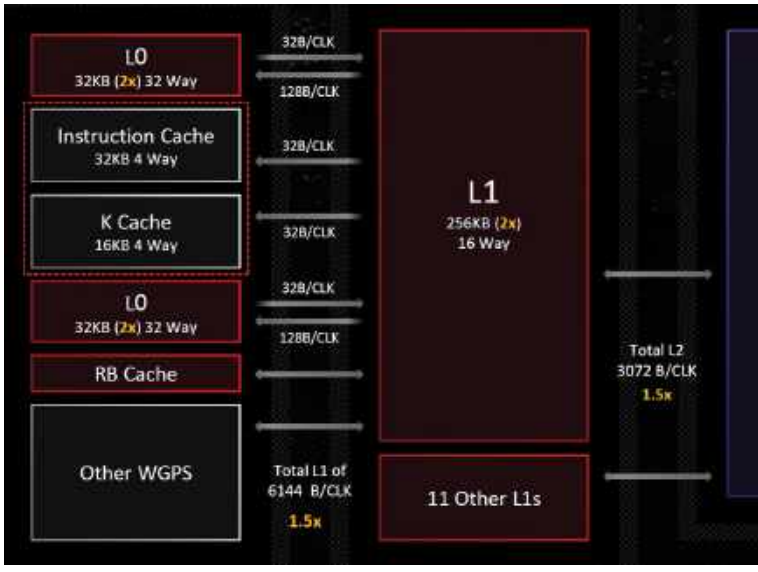

But there's still a cost to the chiplet approach. If this were a monolithic die, built entirely on TSMC's N5 node, there would be no need for the extra Infinity Fabric links. Wipe those out and the total die size would likely be in the 400mm^2 range, as there would be some scaling of the cache and logic sections of the MCD. That gives some insight into the cost difference between N5 and N6. 5nm wafer costs are so much more that it was worth adding about 27mm^2 on the GCD and 33mm^2 combined on the MCDs.

RX 7900 XTX AND RX 7900 XT COMPARED

Graphics Card	RX 7900 XTX	RX 7900 XT
Process Technology	TSMC N5 + N6	TSMC N5 + N6
Transistors (Billion)	58 (45.7 + 6x 2.05)	56 (45.7 + 5x 2.05)
Die size (mm2)	300 + 222	300 + 185
CUs	96	84
SPs	6,144	5,376
Ray Accelerators	96	84
Boost Clock (MHz)	2,500	2,400
VRAM (GB)	24	20
VRAM Bus Width	384	320
Infinity Cache	96	80
Render Outputs	192	192
Texture Units	384	336
TFLOPS FP32	61.4	51.6
Bandwidth (GB/s)	960	800
TBP (watts)	355	300
MSRP	$999	$899

PRICES CORRECT AT TIME OF PRINTING

REWORKING THE ARCHITECTURE

Breaking the GPU into chiplets represents a major change in the underlying design, but that's just the start of the architectural updates for RDNA 3. There are general design updates, enhancements to the GPU shaders, changes to the ray-tracing hardware, and more.

For the GCD, it all starts with TSMC's 5nm N5 process that promises a substantial boost in transistor density and better power characteristics. According to AMD, RDNA 3 GPUs can hit the same frequency as RDNA 2 GPUs while using half the power, or they can hit 1.3x the frequency while using the same power. Ultimately, AMD will want to balance frequency and power to deliver the best overall experience, but the top RX 7900 XTX will push to slightly higher power limits than the previous-generation RX 6950 XT. That should mean it will deliver a decent bump to clock speeds and even higher performance.

With RDNA 3, AMD split the clocks of the GPU core into two domains, one for the front-end hardware (instruction prefetch and decoding) and another for the shaders. On the previous RDNA 2 architecture, AMD found the shaders were outrunning the front end, meaning they were often not fully filled with work and were wasting power. By separating

AMD's latest GPUs feature four levels of cache with significantly wider links between each level, plus up to 24GB of GDDR6 memory.

the clock domains, AMD says it was able to boost silicon utilization by 20 percent while at the same time improving efficiency. The lower-clocked shaders save up to 25 percent in power use, and the higher front-end clock boosts frequency by up to 15 percent.

Speaking of frequencies, AMD lists a boost clock of 2.5GHz for the TXT model and 2.4GHz for the XT (see Speeds and Feeds sidebar). These are not the maximum clocks, and each chip has its own performance characteristics. AMD's Navi 31 GPUs, just like the Zen 4 CPUs, are now designed to use every bit of potential headroom to achieve maximum performance. 2.5GHz is a conservative clock, in other words, and we could see real-world gaming clocks that are quite a bit higher. In fact, AMD says that RDNA 3 was designed to reach clock speeds of 3GHz or more—a speed that we've now achieved with manually overclocked RTX 4080.

What about Nvidia's RTX 4090? Apparently, that wasn't part of AMD's plan for RDNA 3, at least not once it saw the level of performance—and the price—of the 4090. A few people at AMD even indicated that they were surprised with how big Nvidia went with AD102 on the 4090. With RTX 4080 now out, we should find out how AMD's counterattack lands in the next couple of months.

NEW COMPUTE UNITS

Getting into the lower-level details, many of the biggest changes with RDNA 3 occur within the Compute Units (CUs) and Workgroup Processors (WGPs). The CU has been AMD's main building block for GPUs going back multiple generations of GPU architectures. Each CU contains a bunch of GPU shader ALUs, L0, and other cache, shared memory, and in RDNA 2 and RDNA 3, it also houses the Ray Accelerators used for ray tracing.

PERFORMANCE PREVIEW

Talking about all the architectural advances of RDNA 3 is fine, but how do the new graphics cards perform in the real world? We don't have hardware for testing just yet, but AMD provided some benchmarks showing generational improvements when compared with the RX 6950 XT.

Comparing the upcoming 7900 XTX to the current fastest 6950 XT, performance has improved by anywhere from 44 percent (*Resident Evil Village* with ray tracing) to 100 percent (*Dying Light 2* with ray tracing). Across the eight benchmarks, AMD provided, the XTX part is 60 percent faster on average, while the 7900 XT is 37 percent faster.

That's a healthy bump in performance but it doesn't provide any context for how AMD might stack up against Nvidia's hardware. We can help with that. Based on our own tests using different games but still with a mix of ray tracing and non-RT titles, Nvidia's RTX 4090 comes out 120 percent ahead of the RX 6950 XT at 4K. If we skip the RT games where Nvidia traditionally has an advantage, the gap narrows to 68 percent. But AMD suggests that the real competition will be the RTX 4080.

Based on our current results, the RTX 4080 is just over 60 percent faster than the RX 6950 XT, but if we drop RT games that drops to 30 percent. AMD might be more likely to show games where its gains are larger, so we can't issue a verdict just yet, but early indications are that RDNA 3 could prove competitive with Nvidia's penultimate 40-series GPU.

AMD 4K BENCHMARKS

Game	RX 7900 XTX	RX 7900 XT	RX 6950 XT
Call of Duty: Modern Warfare 2	139	117	92
Cyberpunk 2077 (No RT)	72	60	43
Resident Evil Village	190	157	124
Watch Dogs Legion	100	85	68
Cyberpunk 2077 (RT-Ultra)	21	18	13
Dying Light 2 (High RT)	24	21	12
Hitman 3 (Ultra RT)	38	34	23
Resident Evil Village (Max RT)	135	115	94

MAKE THE SWITCH FROM WINDOWS TO THE MOST POPULAR LINUX DISTRIBUTION AROUND!

Discover the best features of Ubuntu's default software, the best FOSS apps for you to download from the Software Centre, and then go beyond the Unity interface.

ON SALE NOW

FUTURE

Ordering is easy. Go online at:
magazinesdirect.com
Or get it from selected supermarkets & newsagents

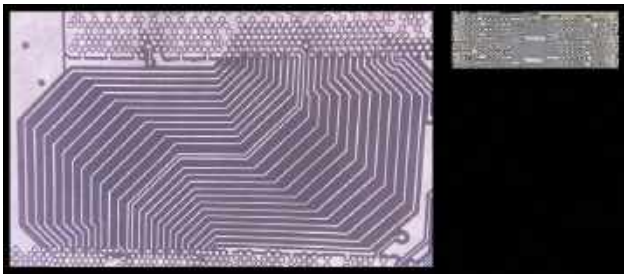

Zen 3's Infinity Fabric link with 25 wires on an organic substrate compared to RDNA 3's Infinity Fabric with 50 wires using a high performance fanout bridge.

RDNA 3 comes with an enhanced Compute Unit pair. With the previous two RDNA architectures, AMD included 64 Stream Processors (SPs) per CU. The SPs are the shader ALUs (arithmetic logic units), and they operate on chunks of similar data using SIMD32 (Single Instruction Multiple Data) operations.

With RDNA 3, the computational throughput has been significantly boosted. Where RDNA 2 contained two SIMD32 units capable of working with FP (floating-point), INT (integer), or matrix data, RDNA 3 adds a second SIMD32 unit per CU that's limited to FP and matrix workloads. In other words, the 32-bit floating-point (FP32) and matrix throughput has been doubled. That's important as the primary data used in consumer graphics applications (games) is FP32.

If this sounds familiar, Nvidia took a similar approach with its Ampere architecture in 2020, doubling the FP32 CUDA cores per Streaming Multiprocessor (SM)—Nvidia's functional equivalent of AMD's CU. Nvidia at the time said that only about 30 percent of the computational workload done with graphics used INT32, so it saw some great benefits in doubling the amount of FP32 compute.

In an odd change of pace, while this should mean that the number of shaders per CU has increased from 64 to 128, AMD still chooses to quantify the number of SPs per CU as 64. However, each SP can now do twice as many FP32 operations as before. If you look at the specifications for the RX 7900 XTX and see 6,144 SPs and 61.4 teraflops, it's because each SP can now perform four operations (two FMA, fused multiply add) per clock.

Another update to the CUs is that the vector units can now issue either a single Wave 64 FMA instruction that operates on 64 FP32 pieces of data, or can dual issue two different Wave 32 instructions. This can improve the utilization of the hardware resources, as there may not always be 64 chunks of data that need the same calculation performed.

AMD shares some of the execution resources in its SPs between vector operations and matrix operations, so with the doubling of the FP32 units, it also doubled the matrix throughput. Matrix operations tend to be used in AI workloads, and AMD calls these AI Matrix Accelerators. These now also support BF16 (brain-float 16-bit) data types along with INT8, and there's now support for INT4 WMMA Dot4 instructions (Wave Matrix Multiply Accumulate).

BIGGER AND FASTER CACHES AND INTERCONNECTS

We already discussed the massive bandwidth of the Infinity Fabric that links the GCD with the MCDs. It's 2.25 times wider than the RDNA 2 link between the Infinity Cache and the L2 caches. All accesses to the GDDR6 memory go through the Infinity Cache, so only cache misses get passed on to the GDDR6 memory, therefore improving the effective bandwidth of the architecture.

The number of GDDR6 chips and the total memory interface width on the full Navi 31 implementation is also 50 percent wider

FIDELITYFX RESOLUTION REVISITED

One of Nvidia's continued advantages over AMD comes in the realm of AI-assisted upscaling technologies. DLSS got off to a slow start in 2018, as the original implementation proved cumbersome, but DLSS 2 greatly improved the overall quality, performance benefits, and ease of implementation. Now Nvidia has DLSS 3 with the RTX 40-series, potentially offering greater performance gains.

AMD has countered with FSR, which has also gone through some updates. FSR 1.0 was easy to integrate and can also be leveraged via Radeon Super Resolution (RSR) in the drivers for a quick performance boost. But the spatial upscaling of the original algorithm, while easy to implement with games or drivers, is lacking in overall quality.

FSR 2 greatly improved the quality by switching to a temporal upscaling algorithm, which also increased the complexity and reduced the performance gains. FSR 2.1 and now 2.2 have focused on reducing the amount of ghosting and improving the overall image quality. At present, FSR 2 has been released in 46 games and there are 55 announced upcoming releases. That's quite a bit less than Nvidia's current 255 games, but AMD continues to gain ground.

AMD has also started work on FSR 3 with Fluid Motion Frames—a vendor-agnostic take on Frame Generation. AMD says it's already seeing some impressive results in early work, with up to double the performance of FSR 2. What we don't know yet is how it will compare with DLSS 3 in terms of image quality, particularly on the generated frames.

Nvidia claims it needs 305 teraops of processing power provided by its Optical Flow Accelerator in RTX 40-series GPUs to do Frame Generation with the desired quality and performance. A simpler form of interpolation is certainly possible, but the artifacting could prove severe. We'll find out how AMD's FSR 3 solution fares later this year.

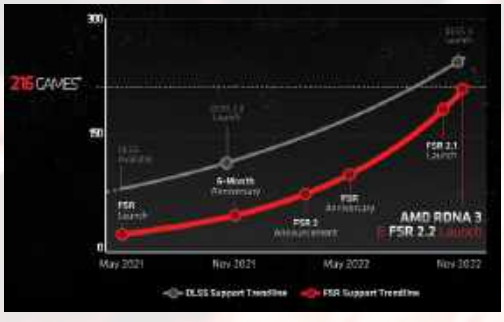

AMD's FSR has enjoyed success in the past 18 months and is the company's most popular graphics library.

VIDEO AND DISPLAY UPGRADES

Like Intel and Nvidia, AMD will add hardware AV1 encoding and decoding support to RDNA 3 GPUs. Also like Intel and Nvidia, AMD will have dual media engines to improve the performance and flexibility with video encoding and transcoding.

We've done some previous testing with the various GPUs to check quality and performance, and AMD's hardware consistently came in well behind Nvidia and Intel, but AMD claims RDNA 3 GPUs will offer significant quality improvements as well—including some AI enhancements on the decoding side.

RDNA 3 will also bring a new Radiance Display Engine into play. This includes full coverage of the REC2020 color space and 12-bit HDR color with up to 68 billion colors. Perhaps more importantly, AMD will offer full support for DisplayPort 2.1 UHBR (Ultra-High Bitrate) 13.5 connectivity. That's something missing from Nvidia's RTX 40-series, and even Intel supports UHBR 10 on its Arc GPUs. In practice, it might not actually matter that much.

DisplayPort 2.1 monitors aren't quite here yet, although we should start to see the first publicly available displays arrive in early 2023. AMD's upcoming GPUs will be able to support resolutions of up to 8K at 165Hz or 4K at 480Hz, thanks to UHBR 13.5. Nvidia's continued reliance on DP1.4a means that it can only provide support for 8K 60Hz and 4K 240Hz.

That's mainly thanks to Display Stream Compression (DSC), which provides "visually lossless" compression rates of up to 3:1. Without DSC, Nvidia would be limited to 4K 98Hz on DP1.4a. With DSC, however, we're still waiting for monitors that approach the limits of DP1.4a. Chalk this up as at least a minor victory for AMD.

RDNA 3 will support DisplayPort 2.1 with up to 54Gbps of bandwidth, providing support for higher resolutions and refresh rates.

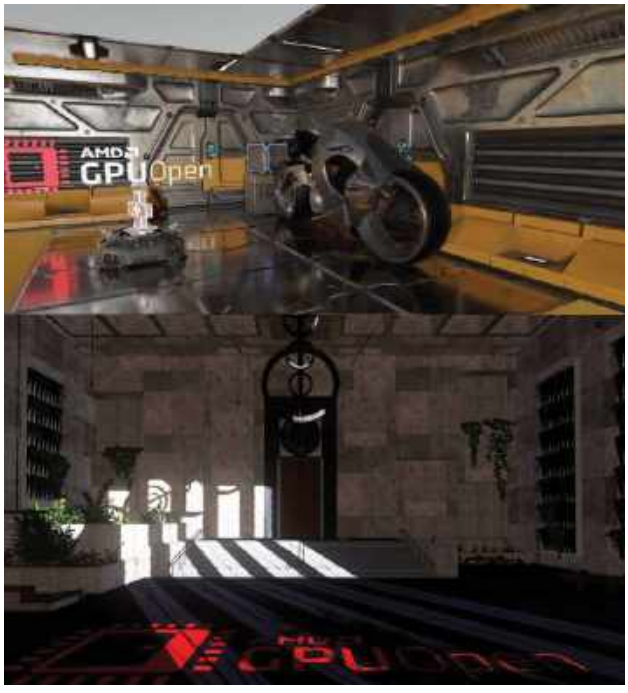

AMD has improved ray-tracing performance while minimizing the amount of die space required for the extra functionality.

than on Navi 21—384-bit compared to 256-bit. Combined with a 20 Gbps GDDR6 speed, the total bandwidth for Navi 31 has increased to 960 GB/s, 67 percent more than the RX 6950 XT and 88 percent more than the RX 6900 XT.

Naturally, that means all the other interconnects need to be wider as well—there would be no sense in boosting the Infinity Cache bandwidth without also boosting the L1 and L2 cache bandwidths. The L2 cache on Navi 31 is now 6MB, 1.5x larger than the 4MB L2 cache on Navi 21, and the interconnect has likewise been increased by 1.5x to 3072 bytes per clock.

The L1 cache size has doubled to 256K and the L0 cache to 32K. Meanwhile, the L1 cache interconnect is 1.5x wider and can transmit 6,144 bytes per clock. All of this goes into feeding data into the GPU computational units as efficiently as possible.

While we can expect many of these changes to apply to all variants of the RDNA 3 architecture, it's worth noting that most of the above specifically applies to the full Navi 31 implementation used in the 7900 XTX. The 7900 XT has one fewer MCD and thus has a 320-bit memory interface and 80MB of Infinity Cache, while other future GPUs will have different configurations—though we expect them to also have larger caches and wider interconnects compared to the equivalent RDNA 2 variants.

AMD'S 2ND GENERATION RAY TRACING

The ray-tracing hardware of the RDNA 2 architecture always felt like something that was tacked on to meet a required feature checklist rather than something that was an integral part of the design. For example, AMD's RDNA 2 GPUs lack dedicated BVH traversal hardware, instead sharing some of the calculations required for ray tracing with other hardware units. That partially accounts for the RX 6000-series' weak ray-tracing performance.

AMD appears to be putting more emphasis on the Ray Accelerators in RDNA 3, which should close the performance game in ray tracing games—at least when compared with Nvidia's Ampere GPUs. The new Ada Lovelace architecture improved RT

performance yet again, and based on what AMD and Nvidia have revealed, it's reasonable to assume that the RTX 40-series GPUs will continue to lead when it comes to ray tracing.

Some of the Ray Accelerator updates are focused on reducing the number of calculations required through early culling of subtrees. This is facilitated through the use of DXR Ray Flags and can reduce the number of instructions required for each iteration of BVH traversal. RDNA 3 also benefits from the larger VGPRs (Vector General Purpose Registers), which are 1.5x larger and thus allow for 1.5x as many rays in flight.

Another new approach for RDNA 3 ray tracing is the addition of three different BVH box sorting algorithms to optimize different kinds of ray tracing workloads. Sorting nodes using the closest first provides a good general-purpose approach. Sorting by largest box first can boost performance for RT shadows and sorting by closest midpoint first improves RT reflection performance.

There are some other improvements to the RT hardware, but overall, thanks to the new features, higher frequency, and increased number of Ray Accelerators, AMD says RDNA 3 should deliver up to a 1.8x performance uplift for ray tracing compared to RDNA 2. What does that mean in terms of relative performance? Check our performance preview sidebar on page 37.

RDNA, TAKE THREE

AMD has made tremendous progress over the past several generations on both its CPU and GPU offerings. The first-generation Zen CPUs showed promise but also had some teething pains—memory compatibility and stability, for example, weren't great at launch. Generation two added new features and started to close the gap with Intel's processors, but it was really Zen 3 and Zen 4 CPUs that challenged Intel's undisputed supremacy.

RDNA 3 feels like that same Zen 3 moment, only this time for GPUs. AMD has taken the key learnings from its CPU chiplets and adapted them for graphics work. The result looks quite different in the details, but the core concepts are all there: improve yields, optimize costs, and allow a wider range of scaling without having to build multiple completely different chips.

Besides breaking away from monolithic GPU designs, AMD also reworked many of the building blocks for its architecture. There are wider and faster interconnects, larger caches, and more computational resources than before. All of these have been tuned to achieve maximum efficiency, delivering another 50 percent or greater improvement in performance per watt.

Nvidia hasn't stood still and Intel has joined the GPU fray, so we need to spend hands-on time with the upcoming graphics cards to see how they stack up. But RDNA 3 feels like the most revolutionary architecture to come out of AMD in a long time and it may give Nvidia the competition we've been hoping to see.

Since total board power hasn't significantly increased over the previous generation, AMD's upcoming 7900-series cards will still be roughly the same dimensions as the 6900-series parts.

GOING ALL-IN ON AMD ADVANTAGE

AMD has been working with OEMs for a couple of years now to promote AMD Advantage laptops. The idea is simple enough: any laptop that uses a Ryzen CPU and a Radeon graphics card can potentially become an AMD Advantage system. AMD will now be extending the program to include desktops with Zen 4 Ryzen 9 CPUs and RDNA 3 Radeon 7900-series graphics cards.

While any enthusiast can put together a PC that uses AMD hardware, the Advantage program specifically focuses on system integrators. Certified desktops will also need an AM5 X670E motherboard, a premium case with tool-less entry, liquid cooling on the CPU, at least 32GB of DDR memory with AMD EXPO memory profile support, and a 2TB or larger NVMe SSD. In short, you'll get a high-end PC with all the bells and whistles—and likely a price to match.

AMD claims there's no cost to the program, which involves close collaboration with the system integrators to ensure it meets quality standards and that users get the best experience. That includes tuning and optimizing BIOS settings so that features such as Smart Access Memory (SAM) and EXPO profiles are enabled.

While SAM leverages PCIe's resizable BAR (Base Address Register) functionality, AMD claims that it gets more benefit than Nvidia because it works on the CPU, motherboard, and GPU sides of the equation. Incidentally, this may also explain why Intel's Arc GPUs seem to benefit more from enabling ReBAR.

Ultimately, the goal of AMD Advantage is to provide people with the option to get a PC that's guaranteed to make the best use of AMD's various technologies. At present, AMD is working with seven different system integrators, mostly established players that cater to the higher end of the market, such as Falcon Northwest, Maingear, and Origin PC.

Pre-built PCs from vendors using the appropriate AMD hardware can be AMD Advantage certified, showing they are tuned for optimal performance.

magazinesdirect.com

Over 100 brands to choose from

3 GREAT REASONS TO SHOP WITH US

- From tech and gaming titles to fashion and celebrity magazines there's something for everyone.
- Immerse yourself in our specialist one-off publications about your favorite hobby or interest.
- Exclusive Offer - Get 15% Off Guides and Specials - **MDMAG15**

VISIT OUR ONLINE STORE AT

www.magazinesdirect.com

✓ Convenient home delivery ✓ Save on shop price ✓ Never miss an issue ✓ 1.4m subscriptions sold

SUBSCRIBE & SAVE 76%

PLUS, GET THE FREE DIGITAL EDITION OF

BOOK OF AMAZING TECHNOLOGY

GREAT REASONS TO SUBSCRIBE

Save money ● Never miss an issue ● Get every issue delivered to your door

To order go to **www.magazinesdirect.com/mpc/B2IMAG**

or call **1-844-779-2822** and mention code **B2IMAG**

TERMS AND CONDITIONS Savings based off the US newsstand cover price of $9.99 per issue for 13 issues, and the retail price of *Book of Amazing Technology* digital edition of $6.99. Offer only valid on one- or two-year print, or print + digital subscriptions. *Book of Amazing Technology* will be emailed three to five days after payment is processed. Visit us at http://magazinesdirect.com or call 1-844-779-2822 for other subscription options. Your first issue will arrive four to six weeks after your payment is processed—please allow additional time for overseas deliveries. Offer valid until February 28th 2023.

BOOST SECURITY & PRIVACY ONLINE

Nick Peers explains how to lock down your identity and data online

THESE DAYS your privacy and security have never been under greater threat, particularly when venturing online. Whether you're worried about malicious individuals, gangs of hackers, or even your own government, there's never been a better time to take the steps necessary to lock down your data and modify your online behavior to protect yourself and your loved ones from the less salubrious side of the internet.

We won't go over old ground in this feature, we hope you've taken our previous advice and have strong security software in place on your computer, and are already protecting your online accounts with strong, randomly generated passwords. What we will do is help you take the next steps to tighten things further.

You'll gain practical advice on strengthening network security and employing VPNs to not only secure your online traffic but allow you to dial into your home network securely when on the road. You'll discover how to generate disposable email addresses to protect your actual address from future data breaches, plus we'll show you how to browse more safely, encrypt data for more secure cloud storage, and tighten email and messaging security. It's all here, so turn the page and start taking back control.

LET'S START BY LOCKING down your online accounts. You should already be using some form of password manager to generate and store strong passwords—at least 12 characters in length combining alphanumeric characters and symbols. But if you're truly looking to lock down your online accounts, there's more to do.

First, make sure you only store minimal personal information in your online account—if possible, no more than email and password. It may be fiddly having to input your card details each time you check out, but this slight inconvenience far outweighs having all your personal data stolen through a data breach.

In fact, you can go one step further. While it's almost certain your email address has already been exposed by a password breach somewhere down the line (don't believe us? Visit https://haveibeenpwned.com/ and don't panic if or when your email address is flagged up), you can minimize its future exposure by switching to disposable email addresses. These special email addresses are designed to be used specifically for online accounts (see the box for details on how they work).

While you can manually create and use aliases via an email relay's own online portal, look to see if your password manager can link into your chosen service to generate random email addresses when setting up new accounts. The following example shows how to combine Anonaddy with our favorite password manager Bitwarden (or VaultWarden, its self-hosted variant).

First, sign up for a free account at https://app.anonaddy.com/register. Once you're verified, you'll find yourself at its main dashboard. Click your username in the top right corner, select Settings, and scroll down to the API section to generate a new token. Copy this to the clipboard, then

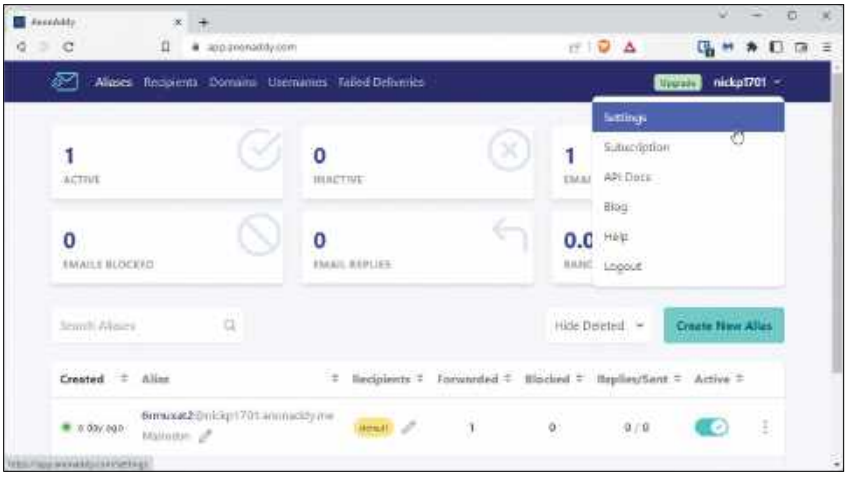

Email relaying services like AnonAddy help protect your email address.

open the Bitwarden browser plugin and click the Generator button at the bottom. Make sure 'Username' and 'Forwarded email alias' are both selected, choose AnonAddy, and paste the token into the API Access Token field. Finally, change the Domain field so it reads username.anonaddy.me (substituting username with your AnonAddy username) to allow you to set up unlimited email aliases with a free account.

Every time you come to sign up for an account, simply go into Bitwarden first to generate a random AnonAddy alias and strong password. Once saved, paste or autofill this into the relevant fields and your email address will remain safely hidden. Logging into your AnonAddy account should confirm the email has been generated—click the edit button under the Alias field if you want to provide a description that explains exactly which account the alias is linked to.

STORE YOUR DATA ONLINE

There's one compelling reason for storing your personal data in the cloud: it's secured in an off-site location that would survive fire or theft in your home. The problem lies with the security and privacy of that data. As with all cloud-hosted material, you must ask yourself: who am I trusting my data with, and what steps have they put in place to keep it private?

Start by consulting your cloud provider's support to find out what protection exists for your data. Take Microsoft's OneDrive for example, it provides detailed information about how your data is encrypted, not just in transit, but also 'at rest', namely while it's residing in Microsoft's data centers. In addition, OneDrive includes a Personal Vault where data can be stored with additional safeguards in place, such as some form of two-factor authentication (2FA) such as a fingerprint, face pattern, or PIN code.

This is all well and good, but data breaches do occur, and can you really trust any cloud provider to guarantee that your data is genuinely out of reach to everyone from hackers to government agencies? The simple answer is no, but although one option is to keep your data away from third-party providers altogether through self-hosting using services like Nextcloud (www.nextcloud.com) or Vaultwarden (https://github.com/dani-garcia/vaultwarden), you still have

> "It may be fiddly having to input your card details, but it outweighs having your personal data stolen through a data breach.

© ANONADDY

DISPOSABLE EMAIL ADDRESSES

The best way to keep your email address private is to surround it with a slew of disposable (or temporary) email addresses.

These addresses pass validity tests for allowing you to sign up for services without revealing your true address because they're little more than aliases that forward all mail sent through them.

While you can do this by hand, setting up unique addresses for each online account you sign up to isn't practical, but thankfully there are plenty of online services offering such throwaway email addresses. First, check if your email provider offers such a feature—

to tackle the question, and expense, of where to safely store redundant copies of your data off-site.

You could house a second server or NAS with a trusted friend or family member, then use a tool like Syncthing to keep your data in sync between the two. In this scenario, you've transferred trust from an unknown third-party provider to someone you know, but risks remain: can they be trusted? What happens if the remote server is stolen, or breaks down?

A more practical solution is to stick with cloud storage but take additional precautions to protect your data. This is done by wrapping anything you upload to the cloud within another layer of encryption, so even if your cloud provider's encryption is broken or they're forced to hand it over, your data is still hidden inside that secondary layer.

BACK UP WITH DUPLICATI

We recommend two tools for the job. For utilizing your cloud storage as an off-site backup location, look for a dedicated backup tool that connects directly to your cloud storage, but offers you additional encryption for your data prior to uploading it. One such free tool is Duplicati (www.duplicati.com). You can install it in Windows or on your Linux-powered server—either way, once set up, everything is administered through your web browser.

When the main screen appears, click 'Add backup' to get started. The wizard is straightforward to follow—the default 'AES-256 encryption, built in' option is simplest, but you can generate and use stronger keys if you wish by installing GNU Privacy Guard on your PC (you'll need to configure GPG under 'Options → Advanced' at the end of the setup process in Duplicati). The key thing with either option is to choose a strong password or passphrase—your password manager can help to both generate and store these via a secure note.

When you come to pick a destination, click the Storage Type dropdown to reveal a vast array of options, from generic (FTP, WebDAV) to the most popular cloud storage providers (choose OneDrive v2 if you're a OneDrive user; Google Drive, Dropbox, and Box.com are also among those supported). Once connected to your cloud storage, choose what files and folders to back up to the cloud, pick a schedule, and finally set various options. Backups are split into 50MB chunks by default, which speeds up the process and provides more redundancy against possible file corruption. You're also able to manually set 'backup retention' to balance keeping multiple backups versus available drive space, but the default smart option is usually the best.

Once done, your first backup—a full one by default—will always be the slowest. If you find your internet connection is a bit sluggish, you can click the speedometer icon on the main screen to throttle the upload speed to say 50 percent of your upload speed. From then, backups only update any changes to your files, so will be much quicker.

ENCRYPT ON THE FLY

Duplicati helps provide you with a secure online backup, but if you've integrated your cloud provider into your file system for backup and sync purposes, you'll need to try a different tool to keep sensitive files private. The best tool for this job is Cryptomator (https://cryptomator.org/). It's free and open source for Windows, Linux, and Mac, while paid-for iOS and Android apps give you access to these files on the road if required.

Cryptomator works by creating one or more 'vaults', special folders inside which you store your most sensitive files. Simply set these up inside your cloud storage folders, and they'll encrypt your files on the fly prior to uploading them, putting them beyond the reach of any malicious actors. Again, each vault is protected with its own password, so make use of your password manager's generator to create strong ones that can't be guessed.

Not all threats to the security and privacy of your files come from the cloud. If you're regularly on the go with your laptop or transferring sensitive data using a flash drive or other removable storage, or even simply worried about the physical security of your desktop, then it pays to encrypt all or part of your hard drive.

The obvious choice here is Microsoft's BitLocker technology, which is built into all Windows versions since Win10. Home editions support 'device encryption', which basically encrypts your entire drive, but if you don't trust Microsoft to look after your data, or want more flexibility, take a look at VeraCrypt (www.veracrypt.fr/en/) instead. It's open source and can be used to encrypt individual partitions, USB

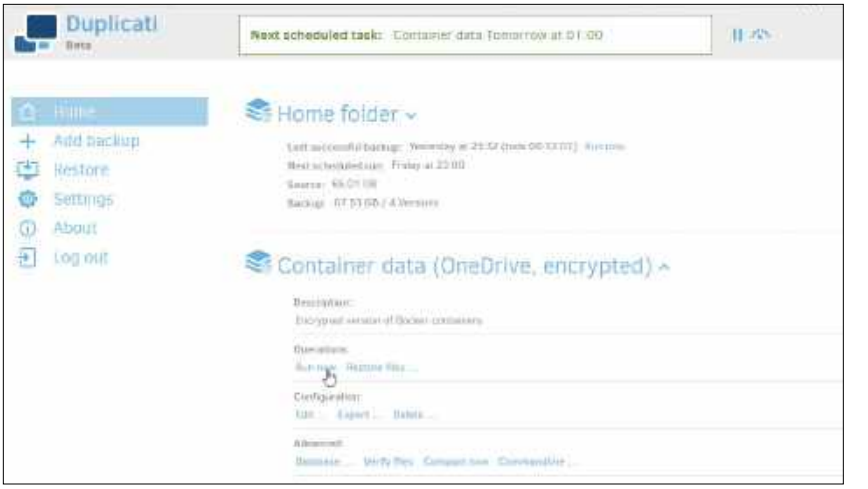

Store encrypted backups of your files online using Duplicati.

Proton Mail (https://proton.me/mail) is one example of a provider that does.

If yours doesn't, then three services to consider, all of which can be integrated into Bitwarden to easily generate disposable email addresses on demand, are Firefox Relay (https://relay.firefox.com/), SimpleLogin (https://simplelogin.io/), and AnonAddy (https://anonaddy.com/). All three services offer free and paid-for tiers, but only AnonAddy's free tier is practical for our needs. Not only does it allow you to create as many disposable email addresses as you need, but you can also connect each address to one of two 'identities', allowing you to choose which actual email address their mail should be forwarded to.

storage, and virtual drives, which work in a similar fashion to Cryptomator's vaults. For our guides on file- and disk-based encryption, see the March 2020 issue.

USE A VPN

Next, let's examine the security and privacy of your physical connection to the internet. Data is transmitted over the internet in 'packets'. Each packet is tagged with your public IP address—that's the address allocated to your modem router by your Internet provider—allowing both it and anyone who intercepts it, such as a hacker or government agency, to trace your traffic back to your door.

Some additional protection may be afforded those packets of data if they've been encrypted, such as secure (https) website connections and data encrypted prior to transmission, but if not, you could also be potentially exposing sensitive information such as passwords or credit card numbers.

One way to hide the data being transmitted, at least until it reaches its destination, is through a Virtual Private Network. The footer on page 50 explains how VPN data encryption and transmission works. There are two ways to use a VPN. The main use is to encrypt your connection to the internet, which requires you to use a third-party VPN provider. You install client software on your PC or mobile, log into your VPN account and then choose which of its worldwide servers to connect to. Once done, your traffic is scrambled between source and destination. There's a performance penalty while your data is encrypted and decrypted, but it's a small price to pay for security and privacy. It's important at this point to understand that using a VPN means transferring trust from your ISP to the VPN provider.

Using a VPN means your data is scrambled and descrambled from source to destination.

When you log into your VPN account, the provider gets your account details and can see everything you're doing online, making it theoretically easy to log all your activities. Mindful of this, VPNs have always emphasized their privacy credentials.

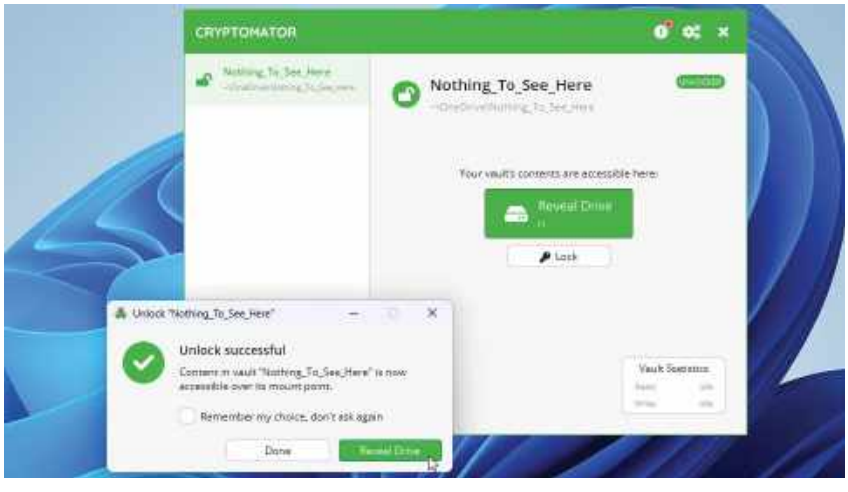

Use Cryptomator to store sensitive files directly in the cloud.

HOW TO PICK A VPN PROVIDER

When evaluating VPN providers, what should you be looking for? Obviously, privacy and security are the key factors: check their promises and in particular their no-logs policy, plus see if they are subject to independent security audits. While not compulsory, membership of the VPN Trust Initiative (https://vpntrust.net/), which establishes baseline practices for security and privacy, is a bonus.

Crucially, find out where each VPN is headquartered. Avoid companies based in countries that are members of one of the 'Eyes' alliances. These nations have signed up to agreements linked to global surveillance of citizens, and include the Five Eyes Alliance nations of the US, UK, Canada, Australia, and New Zealand.

LOCK DOWN YOUR HOME NETWORK

There's little point protecting yourself online if you leave your home network unprotected—and that's doubly true if you plan to dial in using a VPN server. Obviously, you start by ensuring all your devices are secure and fully up to date—also make sure Windows is protected by strong anti-malware software; Windows Security does a good job of covering the basics but consider upgrading to the likes of Bitdefender or Norton for additional layers of protection.

The focus of your attention should be your router, which is your main line of defense against drive-by hackers looking to get into your network. If it's getting long in the tooth, now is the time to upgrade to one that supports the latest security protocols and receives regular firmware updates to fix security holes. Both Asus and Synology routers fit the bill here.

Regardless of whether you upgrade or not, there's plenty you can do to tighten security further. Start

Over the years, VPN technology has evolved with the result there are multiple security protocols you can connect through. The rule of thumb is that the newer the protocol, the more secure it is. Newer protocols like OpenVPN and most recently Wireguard (sometimes offered under a different name, such as NordLynx in the case of NordVPN) also improve encryption performance to reduce the speed penalty.

You should also look at the number of servers offered by a VPN, and which countries they operate in. By connecting to a server in a specific territory you appear to be in that country, which allows you to bypass geo-restrictions, such as those preventing you from watching Netflix while abroad. It also allows you to perform certain activities where there's a grey area over legality—for example, torrenting is safer through a Swiss server because it's illegal for Swiss Internet service providers to see what their users are downloading.

Although free VPN services do exist, they come with many limits—there's usually a capped bandwidth, for instance. Windscribe (https://windscribe.com/) is one of the best here, offering a generous 10GB monthly allowance, plus other restrictions too. You may only be able to use a limited number of servers and protocols, for example, and performance may be capped too.

That means you'll need to pay for secure, reliable, and fast access. Avoid monthly subscriptions as these are prohibitively expensive—typically up to $15/mo. Instead, look for deep discounts on longer periods—usually around two to three years where you can enjoy discounts up to 85 percent. Before signing up, see how many computers and mobile devices are covered by your subscription as the numbers tend to vary.

Most VPNs offer a trial period or a 30-day money-back guarantee, so it's a good idea to road-test several before settling on one. Take the time to benchmark your internet speed before and after connecting to the VPN using a tool such as Speedtest.net to see what effect it has on performance, but note the protocol and server you connect to will influence the results (the further away the server is, the slower your connection will be). Finally, look for useful add-ons—some services offer kill switches that terminate your net connection if the VPN connection is lost, while others support double NAT, which effectively encrypts your connection twice for even greater privacy.

If you're looking for a quick recommendation, then at the time of writing NordVPN (www.nordvpn.com) tops most VPN charts due to its combination of privacy, features, and speed. The Standard plan costs under $3 per month when purchased as a two-year subscription.

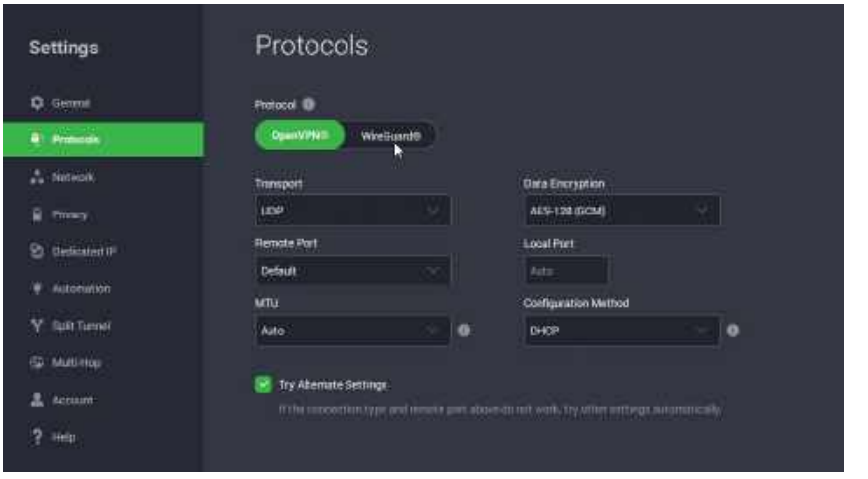

OpenVPN and Wireguard are recommended VPN protocols.

HOME VPN SERVERS

Another use for VPN servers is as a secure means of dialing into your own local network while on the road. Doing so allows you to access network resources from shared folders and printers—plus any self-hosted services you've not opened to access over the internet—as if you were sat at home directly connected to the network itself.

You don't need to sign up for, or trust, any third-party provider, but you will need to set up the VPN server on your network first. Check to see if your router or NAS has a built-in VPN server or offers it as an add-on package – look for VPN Plus Server in the case of Synology routers and NASes, for example.

You'll also find VPN server software for most OSes, although it's best deployed on something that's running 24-7, like the low-powered NAS server we built

> Connecting to a VPN server lets you bypass geo-restrictions, such as those preventing you watching Netflix abroad.

by changing the administrator account password to something strong and disable remote administration of the router, to prevent people from attempting to log on to it from outside your network. Also, take the time to explore your router's security section—look for additional features such as DoS (Denial of Service) Protection and switch on any firewalls, if they exist.

Enable parental controls to manage your children's access and keep them away from unsavory and potentially dangerous content. Also, disable any services that you don't use, such as FTP, to close the ports to outside access.

Next, look to strengthen your Wi-Fi network: check to see if your router offers WPA3 encryption, with WPA2 fallback for older devices, and enable it if it does. Switch on the guest network for visitors and, again, protect it with strong encryption and a difficult-to-guess password or lengthy passphrase.

Make sure the guest network is isolated from the rest of your network and either disable it when not in use or transfer any potentially insecure devices, such as smart lighting systems, to this network for increased security.

in the September 2022 issue. Here we recommend trying Linuxserver's docker implementation of the Wireguard protocol (https://github.com/linuxserver/docker-wireguard).

When you come to set up the server, the choice of security protocol is crucial. If compatibility is crucial, then the L2TP/IPSec option is a good compromise between security and performance. You should be able to connect to this through your device's native VPN connection tools ('Settings → Network & internet → VPN' in Windows 11, for example). Otherwise, you'll need to visit either OpenVPN (https://openvpn.net/vpn-client/) or Wireguard (www.wireguard.com/install/) to download and install the client.

When configuring your connection, you'll need to set up a pre-shared key to act as a password and configure access via a username. Typically, this involves selecting an existing user account on your router or NAS. You may also have to enable a DHCP server to allocate a local IP address to your client when it connects.

It's a good idea to set up and test the connection before you go away. Also, note that you'll need to know your current public IP address—get this from www.whatsmyip.com before you travel. In most cases, this address changes over time, so don't skip this step before leaving your home you'll find it impossible to connect.

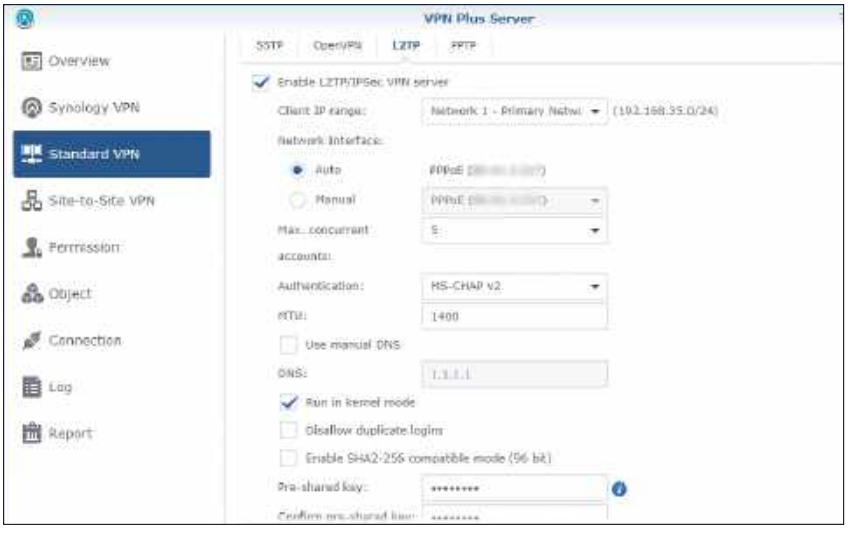

Dial into your home network securely with a VPN server.

BROWSE THE WEB PRIVATELY

Your VPN may hide your internet connection while in transit, but it can't hide it from the services you connect to. With that in mind, you need to take additional steps. When it comes to browsing the web, there are all kinds of threats to look out for, from the well-known, such as cookies and browser history, to the more exotic-sounding, including browser fingerprinting and crypto-mining.

The simplest way to tackle these is by choosing a more privacy-focused web browser. If you're one of the 66 percent of users relying on Chrome, then it's time to change. The obvious alternative is Firefox (https://getfirefox.com), but if you want to push the boat out, there are browsers designed specifically for privacy.

If you can't bear the thought of abandoning Chrome, then try Ungoogled Chromium (https://github.com/ungoogled-software/ungoogled-chromium), which is Chrome minus its dependency on Google's web services. Firefox aficionados should explore LibreWolf (https://librewolf.net/), which hardens Firefox's privacy tools further by removing all telemetry, deleting all cookies and website data every time you close the browser, and including a content blocker by default.

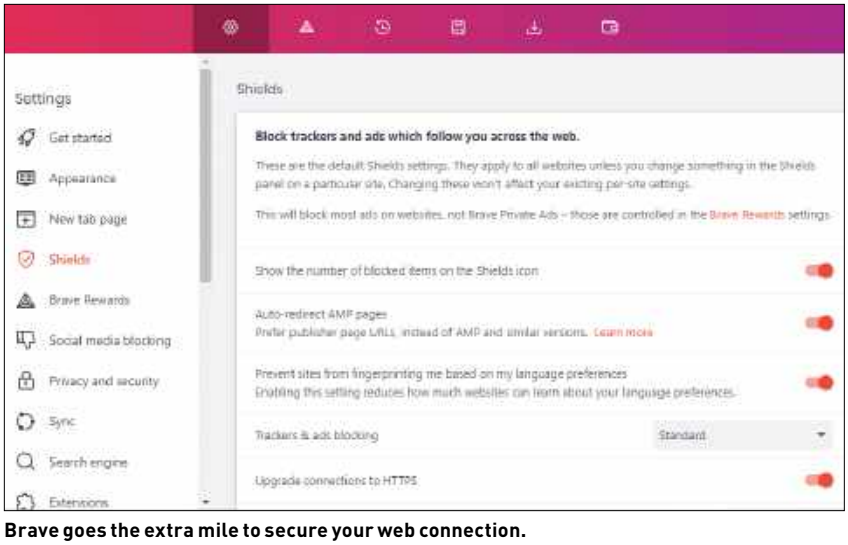

Brave goes the extra mile to secure your web connection.

© BRAVE

HOW DO VPNs WORK?

VPNs serve three specific functions: first, they encrypt data over Wi-Fi networks, including insecure public ones. Second, they prevent your ISP from seeing what you're doing online, and third, they can mask your precise location from any internet services you connect to.

VPNs do this by creating a 'secure tunnel' through which your data passes. Every single packet of data is encapsulated within an encrypted outer layer that masks the data (including what type of data it is) inside. So your ISP knows data is being transmitted, but it can't identify it—at least until it reaches its own servers where it's decrypted to connect to whatever online services you're using.

To prevent this, VPNs bypass your ISP's servers and instead connect to their own. The secure

If you're willing to embrace something less familiar, Brave (https://brave.com/) currently rules the roost among security and privacy researchers. This Chromium-based browser claims to block all privacy-invasive ads and trackers, third-party data storage and browser fingerprinting that identifies you online and results in 'relevant' ads being served to you.

> The simplest way to tackle threats involved in browsing the internet is to choose a more privacy-focused web browser.

LOCK DOWN YOUR EMAIL

We've seen how to protect your email address from unwanted attention, but what about your actual emails? Email is an inherently insecure medium—never mind the phishing attempts and fake email accounts, many emails still aren't sent encrypted, making them easy prey for third-party attackers. That's why you should never send any sensitive information via email. Choose a secure messaging system such as Signal (see below) instead.

Securing your email is a fiddly process—if you want to encrypt emails so they're secure, you'll need your recipient to buy into the same concept. The simplest solution is to sign up to a secure email service, such as ProtonMail (https://proton.me/mail), which uses end-to-end encryption and protects your inbox from unauthorized access. If you want secure communications with a trusted contact, this is the simplest approach for both of you to follow.

If you want to be able to secure an existing email account, then you need an email client that supports OpenPGP (see www.openpgp.org/software/). Our recommendation is the latest version of Thunderbird. OpenPGP secures your email using a pair of cryptographic keys: one public and one private. Your public key identifies you to others and is shared with them. Messages to you are encrypted using this public key, while your private key, which remains in your possession, is used to decrypt the message when it arrives.

This way, OpenPGP not only validates your identity (as well as that of your contact), it ensures that any email conversations are encrypted in transit.

SOCIAL MEDIA AND MESSAGING

When it comes to other forms of messaging, you should employ a similar approach to trust as you do with other online services. Does the messaging system employ end-to-end encryption that's beyond the reach of even the companies or organizations running the platform? And what about any backups?

If you're a WhatsApp user, for example, your messages may be encrypted during transit, but any backups you make to either iCloud (iOS) or Google Drive (Android) aren't by default. To rectify this, open the app on your mobile and navigate to 'Settings → Chats → Chat Backup', then check the 'End-to-end Encrypted Backup' setting, making sure it's set to On.

What platforms you use will in part be determined by your friends, colleagues, and other contacts. But if you're able to give them a steer, we recommend moving over to Signal (https://signal.org/), which promises even tighter encryption than its rivals, is run by a non-profit organization and opens its code to the open-source community for independent auditing.

When it comes to social media platforms, their nature is to be public, so unless you're hiding behind a pseudonym, it's more important to simply be aware of what you've posted and assume it's publicly accessible (even if you've explicitly set visibility to a select group of people).

Instead, focus on making sure your accounts are locked down with strong passwords, random email addresses, and 2FA to help keep hackers at bay. If you're concerned about specific platforms, try to source alternatives. The recent uproar over Elon Musk's acquisition of Twitter has resurrected interest in Mastodon (https://joinmastodon.org/), the less user-friendly but open-source and decentralized alternative, for example.

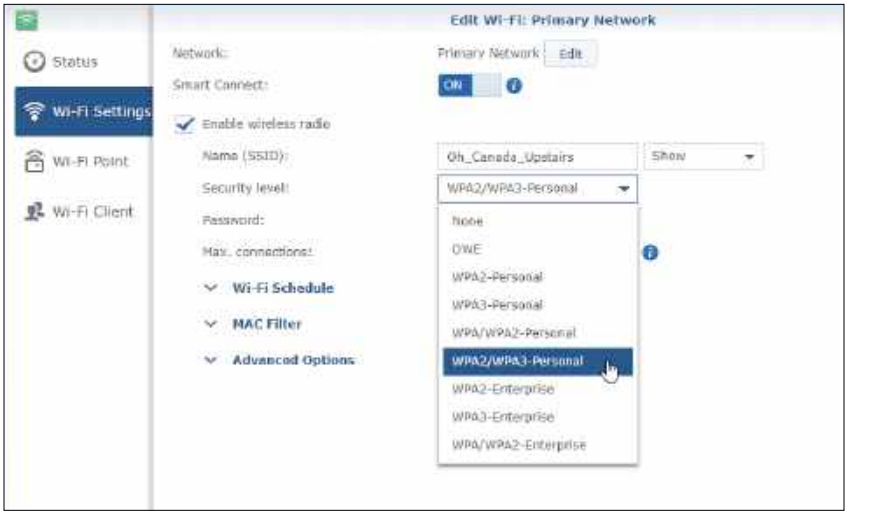

If your network supports it, it's time to switch on WPA3.

tunnel transmits your encapsulated data to the VPN's server, which is the only part of its journey your ISP can track. Here, the outer packet is decrypted and the original data is sent on its way to its destination. Any data sent back goes through the VPN server, where it's encapsulated again before being sent back to you.

The encryption-decryption process requires both client and server possess the correct keys for encrypting and decrypting the data packets, which is why no one else can see the data inside. It's an intensive process that can result in a slower internet connection when a VPN is used.

The re-routing of data via the VPN's server also hides your true location from the service you're connected to because the VPN server is identified as the originator of the data rather than your own public IP address. It can also mask your location from other trackers, such as browser cookies.

RETRO EMULATION

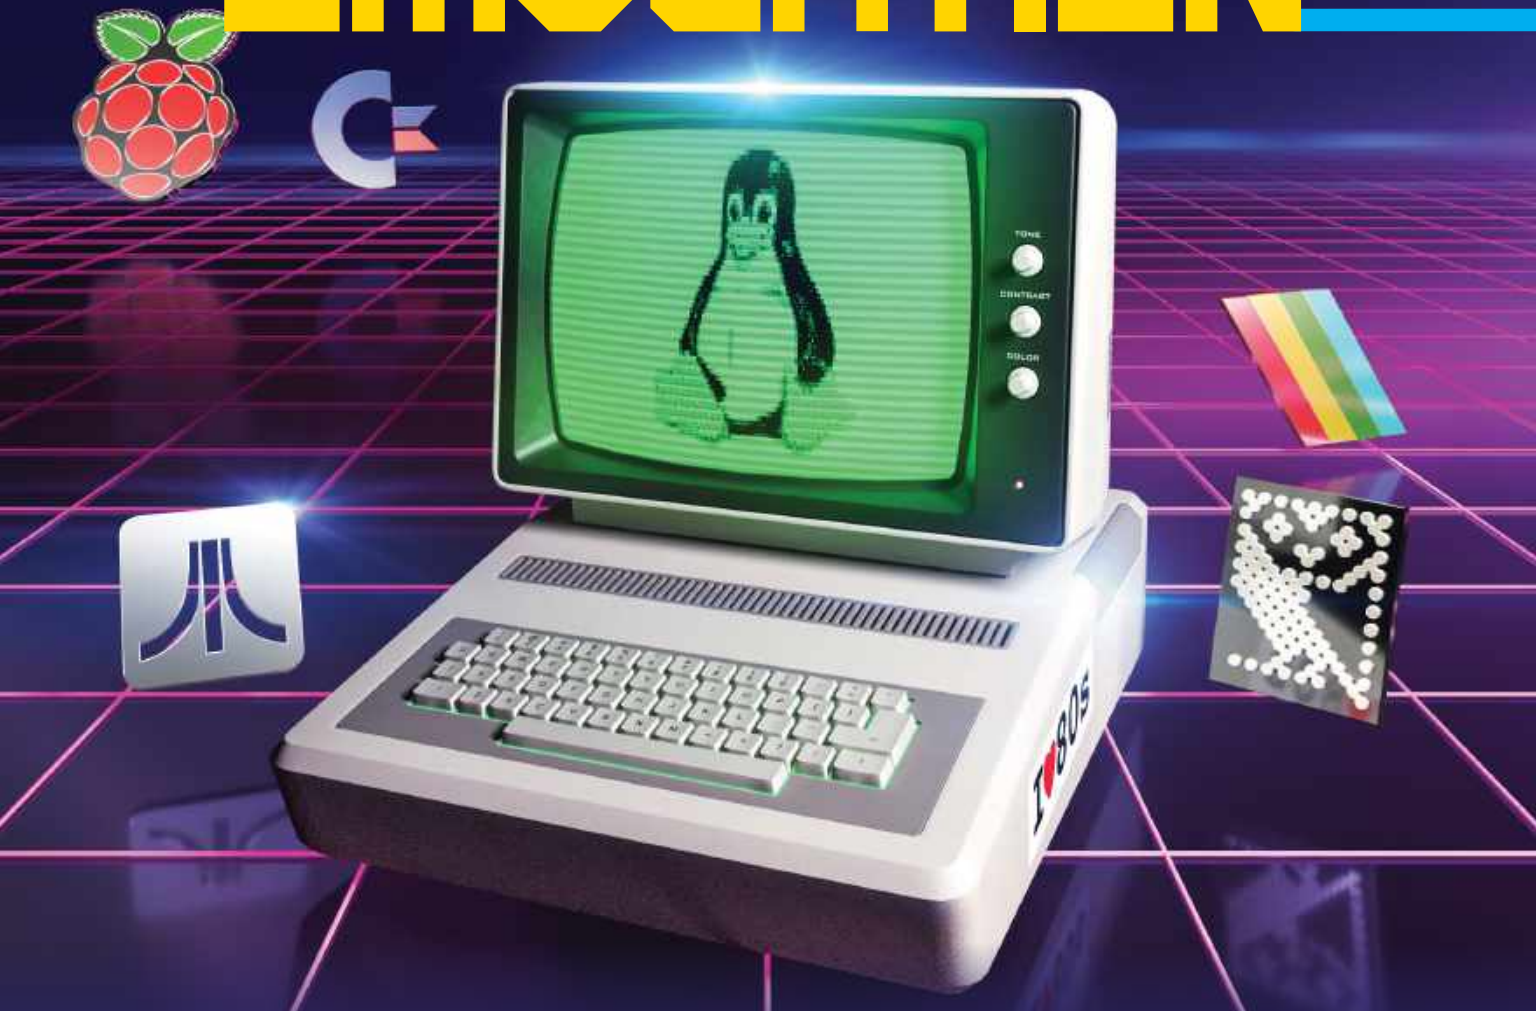

The *Maximum PC* team all have their own favorite retro memories. *Jonni Bidwell* helps you relive yours

We're sure everyone reading this has at least one fond memory of computing in a past era. Maybe it was playing the classic platform game *Crash Bandicoot* on the original PlayStation, or was it discovering the SAY() and TRANSLATE$() functions in Amiga Basic? Perhaps it was even reverse engineering a printer driver for a PDP-11 microcomputer. Whatever your rose-tinted memories of bygone computing eras, it's fun to indulge them.

There are thousands of open-source emulators and tools to help us on this trip into nostalgic reverie. Retro gaming isn't new—the first version of MAME (Multiple Arcade Machine Emulator) appeared in 1996 and enabled PC users to relive hours in the arcade. MAME is better than ever and can now play over 7,000 titles. For the ultimate experience, we'll show you how to build your own mini arcade cabinet with Pimoroni's Picade.

Emulators abound for the machines that defined home computing, such as the ZX Spectrum, the BBC Micro, and the Commodore 64. Moving on to the 16-bit generation, the Amiga and its rival the Atari ST complete with its ugly mouse.

These machines had some great games, which we'll show you how to play. But they also had great operating systems and applications, so we'll relive some of those histories, too.

RELIVE PAST COMPUTERS

A roundup of some of the retro computers that made us smile (and occasionally swear)

RETRO GAMING has long been a popular activity on Raspberry Pi. Soon after its 2012 release, coders worked to get the popular MAME running and other emulators followed suit. The Fuse emulator for the ZX Spectrum series was another early arrival, and Fuse development is alive and well today.

Anyone wanting to recreate the BBC Micro experience, which was particularly popular in the UK in the 1980s, should look at BeebEm (http://beebem-unix.bbcmicro.com). It's one of the oldest emulators, having been around since 1994, and will have you reliving your galactic trading days in *Elite* in no time. You can even run it from a web browser at https://beeb.webassembly.link.

C64 ON A Pi

One of the best-selling home computers of all time was the Commodore 64, which sold somewhere between 10 and 17 million units between its launch in 1982 and its discontinuation in 1994. Official numbers for the Pi family state that over 45 million units had been sold by March 2022, usurping the C64 and becoming the best-selling computer family in history. And if you want to emulate the C64 on a Pi (or a PC), then look no further than VICE (Versatile Commodore Emulator).

After the C64 came the Amiga. According to Amiberry developer Dimitris Panokostas, Amiga emulation on the Pi began in earnest in 2015, when UAE4ARM was ported to it. Dimitris started work on Amiberry a year later and, today, it's probably the best way to get your Pi to pretend to be an Amiga.

The Internet Archive's Arcade Collection allows you to play classics such as *Q*bert* in your browser.

Amiberry uses much of the same emulation code as the venerable WinUAE, and in turn FS-UAE (a Linux port), but with optimizations for lower-powered boards. These all have their roots in the original UAE, which originally stood for 'Unusable Amiga Emulator' because it couldn't do much besides displaying the disk logo. Back then though, most people didn't think Amiga emulation would be remotely possible, so even this was a big step forward.

If we mention the Amiga, some people will be upset if we don't also mention the Atari ST—you can emulate that on a Pi too. Unlike Amiga emulation, where you can't do a whole lot without copyright-encumbered ROMS, the emulation community at https://emutos.sourceforge.io has engineered its own. Combine this with the Hatari emulator and those 520ST memories will come flooding back.

RETRO GAMING ON THE PC

If you're struggling to source a Pi, or don't want to do retro gaming on one, you'll be pleased to hear that all of the emulators and platforms we mention here are available for your PC, too. Emulating newer platforms, such as high-end Amigas and PlayStation, will perform better there too.

If you have an old PC lying around, then you might want to skip the emulation part and run some DOS classics natively. You don't even need a dubiously obtained copy of MS-DOS: the open-source FreeDOS project (https://freedos.org) saw a new release in February 2022. That might seem like a while ago, but it's been feature complete for many years and doesn't have the security (or breakage) concerns of modern software. Alternatively, you can run them on modern Linux via DosBox. You'll find a list of open-source games at https://hellricer.github.io/2020/11/01/opensource-msdos-games.html.

Many classic games have been reworked for modern platforms. Check out *Widelands* (a port of the classic *The Settlers*), *OpenMW* (*Morrowind*), *ET: Legacy* (*Wolfenstein*), *OpenRA* (*Red Alert*, *C&C*, et al), *OpenLara* (*Tomb Raider*), and, of course, *ScummVM* (classic Lucasfilm games including the *Monkey Island* series. These days, it supports a huge number of non-SCUMM games too—see https://scummvm.org). The id Software classics *Doom* and *Quake* (and sequels) were open-sourced a while ago, but more eye-candy is available through ports such as Doomsday Engine, PrBoom+-RT, and RTXQuake (the last two with glorious ray tracing). If you're a fan of the 1980s, look out for *Thatcher's Techbase*, a *Doom 2* add-on (https://thatchers-techbase.github.io).

***Thatcher's Techbase* is a *Doom 2* mod dedicated to the former British Prime Minister of the 1980s.**

RETRO GAMING ON THE Pi

Set up RetroPie and set a course for nostalgic reverie.

Graphics weren't really up to much in 1981, but shaders can even make *Super Tank* look more beautiful.

THERE'S BEEN A RESURGENCE in retro gaming lately. The C64 and NES (launched in 1982 and 1983, respectively) both saw mini editions launched back in 2018 and both enjoyed plentiful sales. A SNES edition appeared last year too, priced rather steeply at $265 (that's relatively more than the original's RRP). This was followed by the Amiga A500 Mini at the start of 2022 (RRP $130).

All of the above are just open-source emulators running on commodity hardware in a designer shell, so you might think the prices are a little overboard. What you're paying for though, mostly, is the license to run these retro ROMs and disk images legally. Nintendo, in particular, will strike down with great vengeance if it catches you distributing any of its intellectual property. Just ask RomUniverse or LoveRetro.co. They might sue us for saying that, so we won't mention their name again.

But there are a huge number of titles spanning the past five decades that you can run for free on cheap, modern hardware. If you have access to old tapes and disks, there are open-source solutions (of varying complexity) for extracting data from these, too. Running the emulators as they are will give you the best gaming experience and maximum customization opportunities. And the Raspberry Pi is a great platform on which to do this.

There are plenty of OSes dedicated to retro gaming on the Pi, such as Batocera, Lakka, and Recalbox, however, we're going to focus on the excellent RetroPie. You don't need much beyond a Pi, SD card, keyboard and monitor to get started, but for an authentic feel, we'd recommend some sort of controller. You can even get one styled after your favorite console from any respectable Pi peripheral peddler.

NO 10 CENTS REQUIRED

If you want to go all out though, get yourself a Picade mini arcade cabinet from Pimoroni. The good pirate-monkey-robot-ninjas at the UK-based Pi emporium were good enough to send us one, and we can honestly say it's a marvel of engineering. It will take a couple of hours to build, but there are comprehensive instructions (https://learn.pimoroni.com/article/assembling-your-picade). Once you've built it, you can play games on its eight- or 10-inch screen with a quality joystick and six brightly colored buttons for bashing. The act

SETTING UP RETROPIE

1. DOWNLOAD PI IMAGER
We usually recommend Balena Etcher for these tasks (and it's a reasonable choice here too), but this time we'll use the official Raspberry Pi Imager. You'll find it in the Ubuntu Software Centre (or repos under `rpi-imager`), Flathub, or straight from the Pi Foundation website at www.raspberrypi.com/software.

2. DOWNLOAD RETROPIE
Open the Imager program and select Choose OS. You'll find RetroPie in the Emulation and game OS category. If you're using a different imaging program you can download the image directly from https://retropie.org.uk/download. You may need to gunzip it, depending on said imaging program.

3. BOOT THE SD CARD
Wait patiently for the image to be written and close whatever you used to write the media. The newly minted SD card will probably show up in your file manager. If so, it's safest to click the eject button there (even if it results in an error message) before removing the card. Pop it in your Pi and you're ready to go!

of putting in coins is replaced by pushing another button. There's a power button too, which can initiate a graceful shutdown if things go weird.

Whatever your setup, the first task is to write a RetroPie to an SD card (see steps, below left). Note that there are different versions for different Pi models (as well as Odroid boards). You don't need a keyboard plugged in if you have a controller, but you do need a screen of some sort.

If you're planning on using a funky kind of screen (such as a Hyperpixel HAT) you'll need to set that up with a regular display and keyboard first. If you're using the Picade you'll need a USB keyboard to set up the controls (press F4 to access the terminal and download the magick script).

To download setup scripts you'll need to either plug in an Ethernet cable or set up Wi-Fi, too. To set up the latter, skip to the terminal and run

```
sudo raspi-config
```

Then skip to 1 System Options→Wireless LAN. You'll need to set the wireless regulatory domain (regdom) to your country before networks are detected. On the first boot, the filesystem is resized to occupy the whole SD card, necessitating a reboot. On the second boot, you should see the RetroPie splash screen and soon after you'll be invited to configure controllers. Hold down any button on the controller (like it tells you) to start. Controls are based on a SNES-like controller, but if you have more (or fewer) buttons it's not the end of the world. The Picade controls are detected as a keyboard, so don't be alarmed (the inputs are converted by the Picade HAT into emulated keystrokes).

The final button, described as Hotkey Enable is important because (in the absence of a keyboard) it, in combination with other buttons, will enable you to configure RetroArch. The 'home' type buttons (such as the X on an Xbox-style controller) are useful here. If you're all out of buttons, the docs (https://retropie.org.uk/docs/First-Installation) recommend using the Select button here.

If you mess up the controller configuration (and save it), you might find yourself unable to navigate the menus. You can plug in a keyboard and use that to reconfigure the controller. Push Esc to display the main menu, then go to Configure Input and press Enter.

LEGAL ROMS

RetroPie doesn't include any ROMs out of the box, so initially, you're only presented with configuration options. RetroPie does helpfully set up Samba though, so you can copy content over without removing the SD card. The ROMs share should appear in your File manager on another PC connected to the LAN. Inside you'll see folders for the many platforms supported by RetroPie.

To test things out we'll grab a ROM (legally!) from the MAME site. Go to https://mamedev.org/roms and find *Super Tank* (in the Other Games section at the bottom). Or choose something else, everything on this page should work well. Check the box (assuming you agree to not use the title commercially) and download the supertnk.zip file. Don't unzip it, but place it in the arcade/ directory

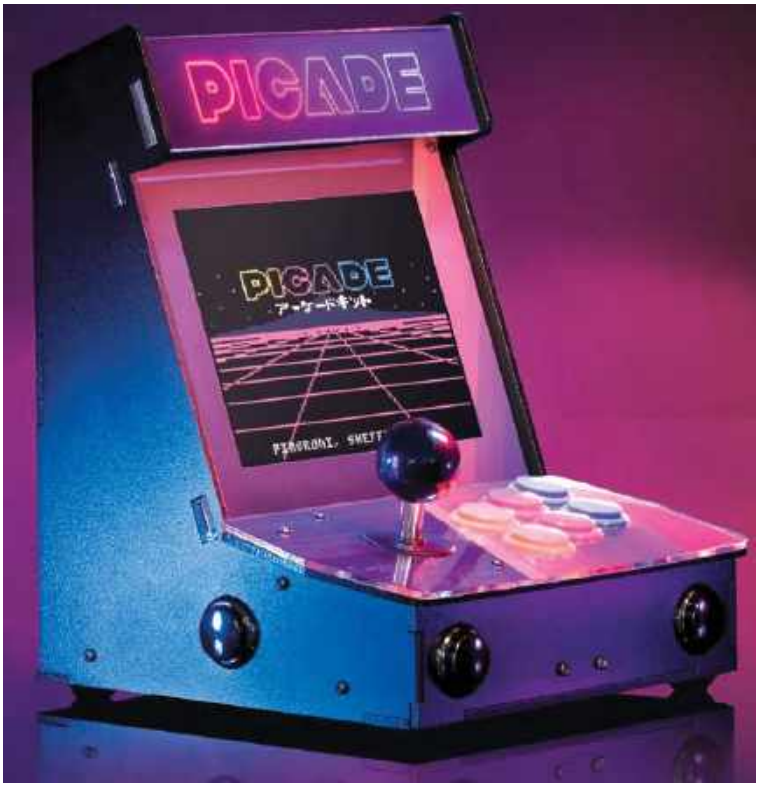

Pimoroni's Picade is the best way to enjoy Pi-based retro gaming.

in the ROMs share. Restart EmulationStation by pressing Start to bring up the main menu, then selecting Quit, then choose the first option. You should now see a new Arcade option, which is for ROMs that run under MAME or FBNeo (FinalBurn, suitable for older arcade titles). The default RetroPie install includes multiple versions of both of these, and when you select the Arcade category and choose *Super Tank,* you'll be asked to choose one as the default. A reasonable choice here is

```
lr-mame2003
```

This should provide good overall compatibility for any arcade games you come across.

All going well, you should see the glorious, square-screened *Super Tank* introduction. Hit Select several times to furnish yourself with credits, then push Start to begin. The goals are simple: collect the yellow Xs, avoid the red ones, eliminate the red tanks and try not to get shot. A two-player mode is available if you have a friend, enemy, or if you simply enjoy shooting defenseless targets.

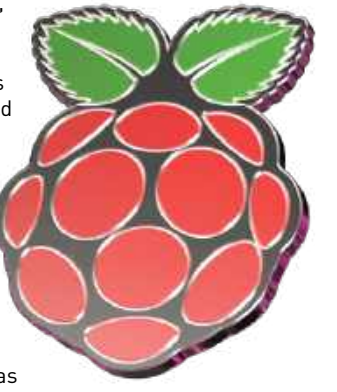

Other MAME versions (besides the three listed) are available in RetroPie. Some of these need to be installed separately. For example, the 2010, 2015, and 2016 editions are more suited to 'modern' games but may struggle on anything less than a Pi 4, whereas Mame4all should work even on a first-edition Pi (although it's considered outdated).

These and other emulators can be added from the RetroPie configuration menu (in RetroPie Setup→ Manage Packages→opt). You can read more on the multitudes of MAME editions by visiting https://retropie.org.uk/docs/MAME.

THOSE MAGNIFICENT MACHINES OF THE 1980s

The eight-bit machines managed to do a huge amount with little memory and CPUs that could barely cope with long division

THE PI A AND PI B monograms are a tribute to the BBC Micro, which also came in a modest A form (with 16KB of RAM) and a beefier model B (with 32KB). That seems reason enough to devote time to emulating those bastions of British engineering.

BeebEm, mentioned previously, isn't included in RetroPie, and nor is it available through its experimental package repos. There have been a few attempts to write custom installers to get it to work. The easiest one is detailed at https://bit.ly/lxf295-beebem. You only need to install the MAME standalone and LR-Mess (non-arcade) emulators to get a virtual BBC B working.

You'll also need to put source some ROM files to make this work. These go in /home/pi/RetroPie/BIOS/mame. And no, we don't know where to get them. You can find disk and tape images at www.bbcmicro.co.uk/index.php, however. You'll need to boot these manually, and seeing as modern keyboards don't have a break key an old-school command line is required:

```
*exec !boot
```

To further confuse things, the asterisk is generated with Shift and '. There is a standalone BBC emulator, BeebEm, and is available for many platforms. But this doesn't build, or at least we couldn't get it to, on modern Linux. You can, however, emulate the BBC on another platform. One of the first alternative operating systems to appear on the Pi was a port of RISC OS, originally written for the first ARM processor on the Acorn Archimedes. You can find out all about it at www.riscosopen.org. We don't have space to cover it here, other than to say it can run the Beebit emulator without too much fuss.

If you can get the BBC emulator working through Mess, the Repton soundtrack will be stuck in your head.

SPECTRUM

Spectrum emulation is more straightforward. The Fuse emulator is included in RetroPie and you can find games to download at www.infinitefrontiers.org.uk/zx-spectrum-downloads. Put the .z80 files in /home/pi/RetroPie/roms/zxspectrum and you should be able to start them from the menu. Spectrum games require keyboard control, but you can map joystick buttons to the keyboard from the RetroGUI in-game menu (summoned by hotkey-X).

COMMODORE 64

People are still creating demos and games for the Commodore 64. Often these use imaginative techniques to produce impressive effects that, according to conventional thinking, wouldn't be possible on such modest hardware. Even in its heyday many of the demos exploited undocumented features (okay bugs) in the 6502 CPU or VIC-II chip to do weird and wonderful things. Even if used within operational parameters, the VIC-II enabled inexpensive shifting of sprites, pixels, characters, bitmaps, and other regions. It was an improvement on previous systems where these things would, for the most part, have to be redrawn by hand.

IF YOU SEE SID, TELL HIM

The SID (Sound Interface Device) chip was ground-breaking at the time. Previous audio hardware wasn't much use for anything besides blips or beeps. SID introduced fine frequency control, which enabled 16-bit resolution and much more sonorous background music. Besides this perfect(ish) pitch, it had previously unseen synthesis abilities. The chip was capable of making four waveforms (square, sawtooth, triangular, and white noise) across

ANATOMY OF RETROPIE

RetroPie is itself just a bundle comprising Pi builds of Emulation Station (ES – the GUI frontend), RetroArch (a frontend which some of the emulators plug into, the one with the funky green retro menu in-game), and the emulators themselves. The emulators that integrate with RetroArch (or rather its libretro library) are known as "libretro cores".

The glue that binds everything together, in a sense, is the RetroPie Setup Script. This is what handles all the package management, including adding custom emulators as you might have done for the BBC Micro above, and as you might do for the Amiga and Atari ST over the page. It can also be used to augment a standard Raspberry Pi OS installation with RetroPie. That's what the first option ("Basic install") is for.

Another important script is Runcommand, which tells RetroPie which emulator and options with which to run a particular ROM. You can configure this by pushing any button when launching a game. Options can be set for a particular system, or in a ROM-specific way.

three channels. Each channel could be modulated, multiplexed, or otherwise transformed, with the effect that there seemed to be many more than three channels.

Each channel has its own envelope, determined by Attack, Decay, Sustain, and Release (ADSR) settings. This makes it possible to simulate, say, organ key presses. There's a musical genre dedicated to including SID output in conventional music. The Swedish band Machinae Supremacy (www.machinaesupremacy.com) describes itself as "SID-metal". If you want to make your own SID tunes, check out GoatTracker (https://sourceforge.net/projects/goattracker2 or see https://retro64.altervista.org/blog/making-commodore-64-music-the-sid-and-goattracker for a guide.

The VICE emulator isn't included in the standard RetroPie installation but is easy to install. Go to the RetroPie configuration, select RetroPie Setup and select "Manage packages", then "opt" and scroll down to find Vice in the list of emulators. Select it and choose to install a pre-compiled binary. Installing from source is possible too, which may be a newer version, but may also take a long time. This is also true for any of the emulators listed.

The C64's predecessors had the PetSCII character set instead of graphics. Try your hand at PetSCII art at https://petscii.krissz.hu.

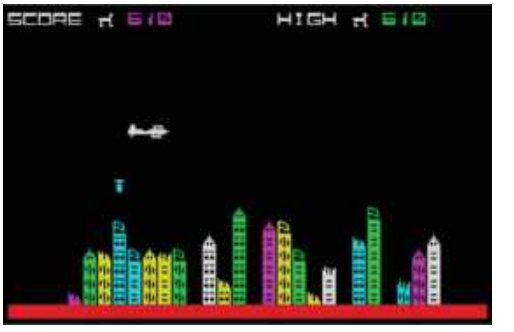

Jeff Minter released all his titles as freeware, including *City Bomb* for the ZX Spectrum.

It can emulate not just the legendary C64, but the PET series (first launched in 1977), the VIC-20 (1980) as well as the post-C64 models (the C64DTV, C128, the Commodore Plus/4), and more.

It's easy to get started with VICE. You can find some impressive demos at https://c64.ch. Pick one and pop the zip file in /home/pi/RetroPie/roms/c64. Multi-disk affairs will need to be unzipped. You can configure VICE from its menu by pressing F12: swap disks, increase emulation speeds, or quit the emulator. You might also want to try the VICE libretro core, which you can find in the "Manage packages" menu. This one can be configured from Retroarch's menu. Both of these emulators have some useful options, for example, "Warp mode" loads games quicker, at the cost of some emulation accuracy. Modern, disk-based, demos in particular take a long time to load (as did disk titles on the original C64), so be patient. There's an "auto warp on load" option if warp speed affects the program. You may need to configure a keyboard controller for some titles, though this may be a little haphazard. You can summon an on-screen keyboard by pressing right shift or Select.

MODERNIZING GRAPHICS

Some of the machines we've emulated here did a lot with not many pixels. There's a certain charm to playing games at their original resolution, but on a powerful system, it's possible to do high-quality image resampling and upscaling.

Summon the in-game RetroPie menu and scroll down you'll find a submenu entitled Shaders. Select Load and you'll see a list of all kinds of graphical customizations (implemented as OpenGL Shaders) that can be applied. Try selecting the 2xScaleHQ preset. Now return to the quick menu (no need to restart the game) and select Resume. Those jagged edges are a lot less jagged, and the fonts are nicer too.

There are all kinds of filters available, and they're not all about improving the look. Some of them are quite demanding, especially on early model Pis. Some of them, however, have been designed to be sparing with resources. In particular, the crt-rpi* and crt-blurPi shaders were written to make authentic-looking cathode ray tube effects that even an original Raspberry Pi can manage. We noticed that our emulated Spectrum keyboard stopped working after changing shader settings, but returning to the controls menu (or restarting the game) seemed to fix it. If it's any consolation, much worse things happened on the original hardware.

It's worth messing around with Retroarch shaders to improve the look of your games.

16-BIT HITS

For many gamers, the golden age of computing was characterized by the Amiga 500 and Atari ST

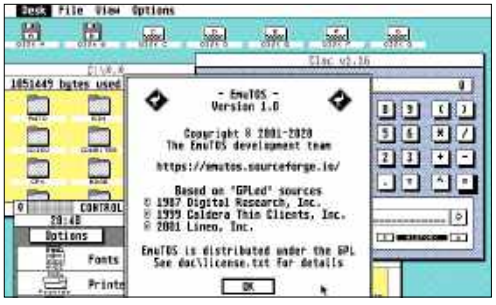

Atari TOS was superceded by MINT. Those are still copyrighted works, so here's EmuTOS.

OUR FAVORITE HOME computer was the Commodore Amiga, in particular the A500, which launched in 1987. As the successor to the immensely popular C64, the machine had big shoes to fill. And fill them it did—right up until questionable management effectively killed it off.

That said, Hyperion Entertainment produced a version of AmigaOS (www.amigaos.net) for PowerPC-based machines. A-EON technology, in collaboration with Hyperion, led to the AmigaOne X5000, the latest in a series of machines capable of running AmigaOS. These can emulate the Motorola 68000 at the heart of the first Amigas without batting an eyelid, though, at $2,500, you might say the X5000 is eye-wateringly expensive.

RetroPie 4.8 (the latest version at the time of writing) was released in March 2022, so some of its components, in particular some auxiliary pre-built binaries, are a little out of date. This can be worked around by building those components from source instead, which is what we're going to do to get Amiberry working. Go to RetroPie Setup→Manage Packages→Manage optional packages and select Amiberry→Install from source. It doesn't take long to build. On a Pi 4, it takes less time than making a cup of coffee. Once it's done, you'll want to reboot, using Restart System from the Emulation Station menu.

Amiga Format's 1993 Christmas Special featured a glorious *Cannon Soccer* mashup.

ABOVE BOARD ROMS

You should now see an Amiga option in the main ES menu. The Amiga Kickstart ROMs are still copyrighted but can be legally purchased from Cloanto's "Amiga Forever". You need these files to do almost anything with an Amiga emulator (although some demos can run with the free AROS ROM included with Amiberry). As well as that, you'll need some disk images (most commonly in the form of ADF images).

Some websites have been permitted by the original developers to distribute their games. For example, Team 17 titles are available from http://dream17.abime.net and Gremlin games are available at http://gremlinworld.emuunlim.com/amiga.htm. And of course, *Amiga Format*/*Power* cover discs from back in the day are easy to find.

Place your Kickstart ROM(s) in /home/pi/RetroPie/BIOS and ADF/IPS disk images in /home/pi/RetroPie/roms/amiga (these paths are the bios/ and roms/amiga/ Samba shares if you're copying over the LAN). These paths were also set incorrectly on the first boot when we tried, so you may need to plug in a keyboard and hit F12 to set them correctly and then hit the Rescan Paths button to pick up the changes.

Navigating the Amiberry configuration is a little tricky with just a game controller, and even using a keyboard it's not ideal, so you probably want to plug in a mouse too. If you're going to be playing *Cannon Fodder*, *Populous*, or other such classics you'll need one anyway.

One of the most exciting developments in the Amiga's history was Blitz BASIC. An edition appeared mounted to the cover of *Amiga Format* in November 1993, along with a zombie apocalypse game programmed in the language. Blitz BASIC soon unseated AMOS as the best way to program games, thanks to its simple(ish) BASIC-like syntax combined with advanced graphical capabilities.

THE SUPERIOR ATARI ST

The Atari 520ST (and in particular the STFM variant) was a fine machine indeed. It just wasn't quite as cool as the Amiga. However, it had some

neat features, such as the built-in MIDI interface, and the STFM's 720K disk format being almost compatible with MS-DOS. It was popular enough to warrant a magazine with a title ending in "Format", too.

You can install the Hatari (Swahili for "danger" and the title of a John Wayne movie, incidentally) emulator in the same way as you did for Amiberry on the previous page. Again it's worth installing from source to get the latest version, but the binary package worked fine in our limited experiments.

Disk images go in RetroPie's roms/atarist/ folder. As with other systems, you'll need a BIOS ROM too. The Atari's OS was called TOS (The Operating System, or Tramiel Operating System) and is still a copyrighted work.

The Bitmap Brothers classic *Xenon 2: Megablast* featured a soundtrack by Bomb the Bass and a rather pushy sales alien.

However a legal and open-source equivalent (based in part on sources made available by Digital Research) named EmuTOS is available, which you can find at https://emutos.sourceforge.io. There are different ZIP archives for different models, and we'd recommend the emutos-512k-1.2.1.zip version to get started. Inside you'll find country-specific images. For example, the US edition is named etos512us.img. Extract this file (or whatever locale best suits you), rename it tos.img, and put it in the bios/ share. To get the controls working in-game, press F12 to bring up Hatari's main menu. In the Joysticks section select either the "keyboard" or "real joystick"

In addition to the ST series, Hatari can emulate its more advanced, 32-bit successors—the TT or Thirty-Two (Atari chose all the best acronyms) and the Falcon. The TT was intended to be a Unix machine, but it took some time for the Unix port to arrive, so that didn't work out so well. It wound up being used as a developer machine for the similarly ill-fated, Atari Jaguar.

The Falcon, by comparison, was a beast of a machine, which included a high-end digital signal processor (DSP) capable of wrangling signals above CD-quality sample rates. Not only that, but it featured a true-color video chip, the VIDEL. Later versions included a 500MB hard drive, too.

And that ends this journey into computing history. Who knows, maybe in 30 years we'll be emulating *SuperTuxKart* and wondering why Wayland and PipeWire were ever a thing. Do let us know about your retro tech memories, and what steps you've taken to recreate them on modern hardware.

RETRO-FITTING RETRO HARDWARE WITH RASPBERRY PI

Thanks to cheap hardware and people's unabashed inventiveness some really amazing peripherals are now available if you still have any working (or almost working) hardware laying around.

First of all, there's PiStorm. This fantastic (and open hardware) board, the work of Claude Schwarz, sits in between the 68000 slot on the Amiga mainboard and a Raspberry Pi. So instead of chugging along at 7MHz, your Amiga will be running an emulated 68EC030 running at around 70-80MHz. Furthermore, the Pi can share up to 128MB of its memory, which the Amiga sees as "Fast RAM", as opposed to the Slow RAM expansions we all had to save up $50 for in the 1990s. You can track down PiStorms on eBay for around $50.

Another fantastic device is the Pi1541. C64 fans might recognize those four digits from Commodore's 1541 external disk drive, which is exactly what it emulates. Unlike the original, it's neither noisy, slow nor cumbersome. It enables you to load 170kb disk images from an SD card, ideal if your 1541 has chugged its last. Read more about the Pi1541 at https://cbm-pi1541.firebaseapp.com.

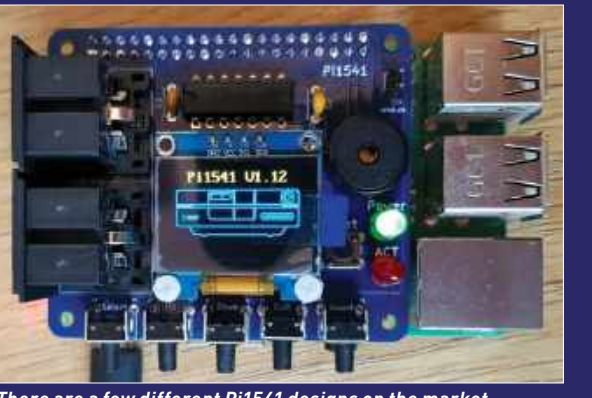

There are a few different Pi1541 designs on the market. This one comes with a rather fetching OLED display too.

DISCOVER THE ALTERNATIVE TO APPLE AND WINDOWS

This is a treasure-trove of Linux and open source knowledge, from the evolution of Ubuntu and its mobile platform, to projects like automating your home with Raspberry Pi.

ON SALE NOW

FUTURE

Ordering is easy. Go online at:

magazinesdirect.com

Or get it from selected supermarkets & newsagents

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

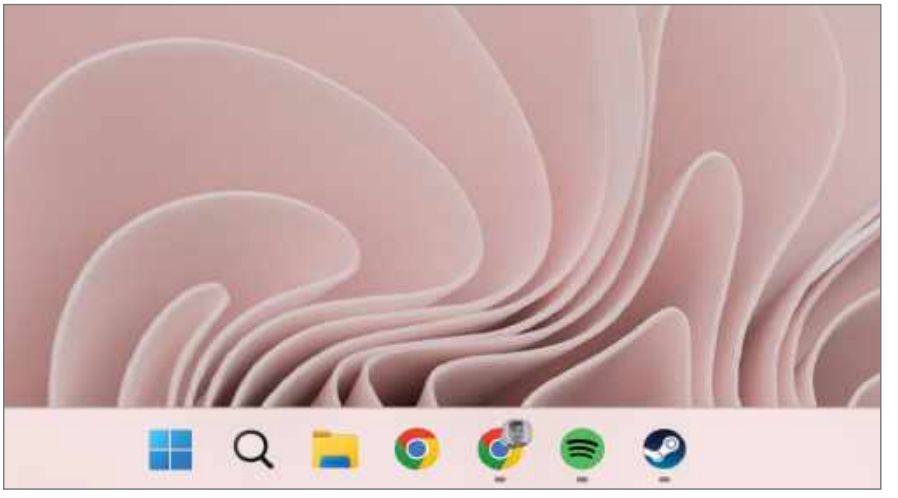

PINNED ITEMS

If you have lots of pinned applications in your taskbar, there's a simple way of getting to them that doesn't involve reaching for your mouse, meaning you can easily cycle through your taskbar apps with this tip. Whilst holding down the Windows key, press a number to open one of your pinned apps. The number corresponds to its location on your taskbar. Easy, right? You don't even need to touch your mouse!

MAKE – USE – CREATE

62
Tear down Xbox's adaptive controller with iFixit.

64
Streamline and speed up your startup process.

68
Colorize black-and-white images in Photoshop.

SAM LEWIS
STAFF WRITER

ABOUT TIME!

I've been trying to move house for a while, but the time has finally come—it's been so long that a new generation of CPUs and GPUs has been announced and launched in that time!

The past six months have felt like being in limbo, waiting to settle down permanently again. A few move dates have fallen through, so most of my possessions have been boxed up and ready to go for a while. The novelty of a minimalist lifestyle soon wears off and sleeping on the floor is no fun, I can tell you!

Anyway, the most exciting part is unloading all my tech into my new place. As good as laptops have been, using a desktop PC again is a relief. Having a widescreen display is a must-have for productivity and, of course, the occasional bit of gaming. My Noblechair has been in storage, but my back is certainly feeling the benefit of using it again.

However, enough of the sob stories. I'm not sure specifically how I'm going to design my setup, but in my head, I've pictured an L-shaped desk with a widescreen monitor as my main workstation and a smaller gaming monitor on the side. Of course, I don't want cables in sight so this is going to be my next main project. By the next issue, I shall be sitting comfortably at my desk enjoying my new setup. Phew!

↘ submit your How To project idea to: editor@maximumpc.com

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

Xbox Adaptive Controller

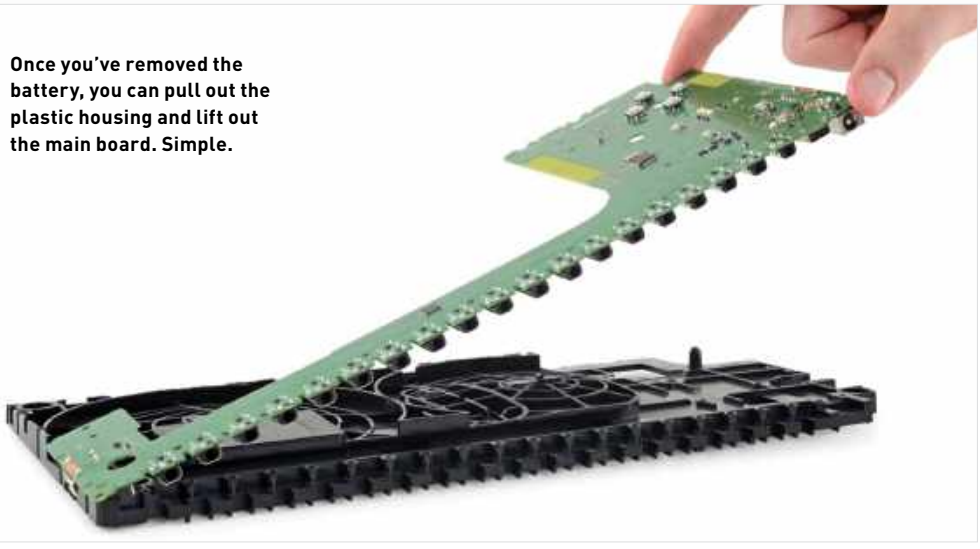

Once you've removed the battery, you can pull out the plastic housing and lift out the main board. Simple.

About iFixit

iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their own products. To learn more, visit www.ifixit.com

There may not be many buttons, but a total of 19 3.5mm connection ports allow for limitless number of control options.

The packaging has been designed to be as accessible as possible, with lovely big handles for lifting out the hardware.

Four rubber feet conceal the T9 Torx screws needed to open the controller up and find out what's inside.

This controller is huge, so the battery doesn't take up much space.

BACKGROUND

The Xbox Adaptive Controller has been built from the ground up to allow accessible gaming. It features programmable buttons and lots of ports for external switches, buttons, mounts, and joysticks.

MAJOR TECH SPECS

- Compatibility: Xbox One, Xbox Series S/X, Windows PC
- Dimensions (HWL): 23 x 130 x 292mm
- Wireless Protocols: Xbox Wireless Adapter, Bluetooth
- Mount Compatibility: 1/4-20 screw for AMPS, °-20 tripod screw
- Charging: 9ft USB-C cable, 7.79 Wh rechargeable Li-Po battery
- Ports: 19x 3.5mm, 2x USB 2.0, USB-C, 3.5mm stereo jack
- Weight: 552g
- Price: $100

KEY FINDINGS

- Microsoft's Xbox Adaptive Controller is the most customizable game controller ever made, and is designed to be as adaptable as possible, at least on the outside. The main unit is flat and can be mounted onto a tripod for those who may struggle to grip a traditional controller. The rear panel has an astonishing 19 3.5mm connection ports, representing each of the buttons, sticks, and triggers used in the Xbox control scheme. These can be connected to switches that can be controlled by your head, foot, or mouth, many of which are made by medical supply companies and can cost up to $200 apiece. It's a remarkable bit of kit for gamers who cannot use a traditional controller.
- Special praise must go to Microsoft's packaging team. The boxes all have ring-pull-style fasteners that don't require scissors and, in the box, big handles pull out the device, making it as easy as possible to get to the product. The unit is wider and flatter than the original Xbox controller, with two large buttons on top mapped as 'A' and 'B' by default. The side boasts a 3.5mm headphone socket, while the rear has a USB-C socket for the long 9ft cable—another thoughtful design. There's an internal battery, so no worrying about finding AA batteries.
- We want to find out whether the Xbox Adaptive Controller is as customizable on the inside. Flipping it over, we find four rubber feet that need to be peeled back to reveal T9 Torx Security screws. Not our favorites, but still an easy way to access the internals. Inside the back panel is a Microsoft easter egg—an engraving of a honey badger. Don't ask us what the relevance is, but it's a bonus for those unafraid to open up their products.
- The first thing to come out is the 7.79Wh Li-Po battery, which is about the same size as the two AA batteries in the standard Xbox controller. The battery needs to be removed so that the plastic covering the main board can come out. There's some encouraging news, in that the headphone port is modular and can easily be replaced. The other 20+ ports, however, are not.
- The plastic and the mainboard are tied together with 12 Torx security screws and one Phillips screw, which seems overkill. There's only one board, but so many ports are soldered down, so it will be difficult to service if any break. On the plus side, it's unlikely you'll use all 19 ports, and everything can be modified in software, so you can use another port and remap its function in the app. The buttons on the controller use plastic covers and spring mechanisms, so they won't be easy to replace.
- Repairability score: 7 out of 10 (10 is easiest to repair). We wish Microsoft had used more standard screws on this controller, but at least there's no adhesive in sight, which means disassembly is straightforward and non-destructive. The battery is the most likely component to need replacing, which can be done by removing four screws. However, high wear-and-tear components, such as the springy A and B button covers, will require special tools to repair if they break.

Streamline and speed up your PC

YOU'LL NEED THIS

BOOTRACER AND AUTORUNS
Wise Game Booster (for on-the-fly optimization)
Process Lasso (to rein in resource hogs)

DOES YOUR PC take longer to boot than it used to? In most cases, you can trace this to the growing number of apps on your system, thanks to their tendency to insert themselves into the start-up process even when they're not essential. What about other parts of your system? Does your PC suddenly struggle to play certain games, or have you started to notice things grinding to a temporary halt every now and then for no obvious reason?

If the answer to any of these frustrations is yes, it sounds as though Windows needs streamlining to eliminate the resource hogs and divert precious system resources to where they're most needed. In this issue, we'll focus on eliminating those bottlenecks and reveal ways in which you can keep unruly processes from disrupting your computing life.

We'll benchmark your current start-up times and reveal how to cut back on the programs that load with Windows (spoiler—only a few should start with your OS), then reveal the various ways in which you can keep your PC running smoothly, from freeing up system resources when you need an extra bit of oomph to ensuring errant processes are kept in check. –NICK PEERS


A

1 UNDERSTAND WINDOWS STARTUP

No, you're not imagining things—your PC really does take longer to start than it used to. As you install new apps, a surprisingly large number attach themselves to Windows' start-up routine, which is fine for essential software like your security package, but not so great when all the app is doing is shaving a few seconds from its own start-up time by reserving a section of RAM for itself. And there's the rub: not only do start-up apps extend your boot time but they also eat into your precious system resources, creating a drag on your PC's day-to-day performance.

›› Sometimes, though, a slow boot time is indicative of something else, particularly if your PC takes an age to reach the login screen. The box reveals the various phases of the start-up routine, and what to do if you experience delays even before reaching the login screen.

›› In most cases, however, your slow startup can be traced to too many apps loading with Windows. Step one to addressing the problem lies in benchmarking your startup. Using a stopwatch is fiddly, but the good news is that there are apps out there that can do the job for you.

›› For a quick and dirty benchmark of how long Windows take to load, followed by how long for all those apps to load themselves, try the portable version of Startup Timer (https://startuptimer.com/). You'll get two figures: 'a fully started', plus 'desktop first shown in' [**Image A**]. Beneath this is a list of detected start-up items, each with a percentage figure allocated to them—the bigger this is, the longer that app requires to load.

2 BENCHMARK YOUR STARTUP

Startup Timer is useful as a quick exercise, but for a more detailed look—and access to tools to help optimize your startup—head to https://greatis.com/bootracer/ to download the free version of BootRacer. After extracting BootRacerSetup.exe from boot_racer.zip, double-click the file and install the application. Once done, leave all four boxes ticked and click Finish.

›› When the main screen appears, you'll see two options, Full Boot Test and Startup Control, plus a large Start button. Click this to open the Quick Speedup tab of Startup Control and click 'Start Test' followed by Yes to reboot and perform a full benchmark of your system.

›› When you next log in, you'll see a series of messages above the Taskbar notification area revealing BootRacer is monitoring your system. Among these is a countdown clock declaring "Windows is almost ready". As start-up items load around you, the clock will keep resetting to 10. Eventually, it counts down to zero and you'll get your first report [**Image B**]. It's likely to say average or slow, indicating that there's room for improvement.

›› Underneath is a blue button that alternates between 'Auto close' and 'Know more?'. Click this to start our deep dive into optimizing your startup. If it closes, open BootRacer via the Start menu or search tool.

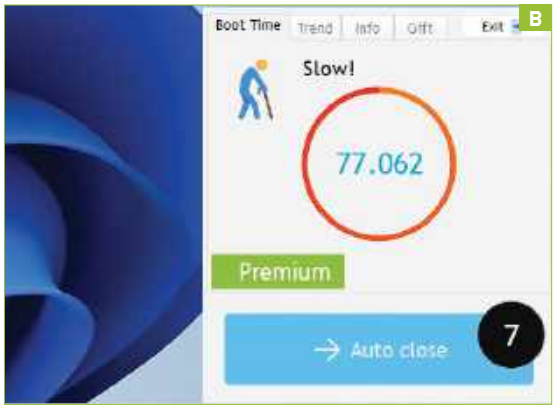

B

© STARTUPTIMER, BOOTRACER

MANAGE YOUR SYSTEM WITH PROCESS LASSO

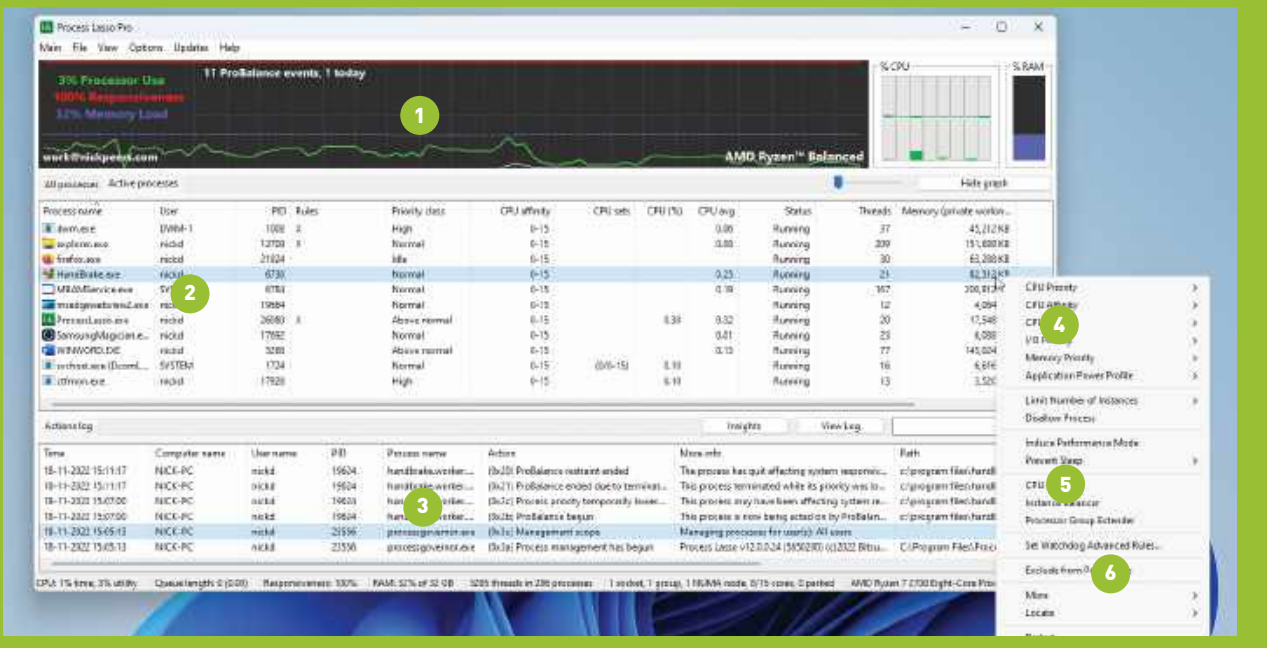

1. SYSTEM RESOURCE USAGE
This graph tracks your processor's CPU use and responsiveness, plus memory load. When a process threatens to reduce overall responsiveness, Process Lasso steps in.

2. PROCESS INFORMATION
Monitor how individual processes are performing. To focus on running apps, switch to the Active Processes tab to filter system processes.

3. ACTIONS LOG
This enables you to track how Process Lasso monitors—and manages—your processes. Look for references to ProBalance being enforced on resource-heavy apps.

4. MANUAL SETTINGS
In most cases, Process Lasso knows what it's doing, but if you need to override it for any reason, right-clicking the process allows you to do so.

5. CPU LIMITER
This setting allows you to restrict the number of processor cores available to the selected process when the system needs additional resources.

6. BYPASS PRO-BALANCE
There will be times when you want certain apps or services to have access to all the resources they need—select 'Exclude from ProBalance' to prevent Process Lasso from managing them.

3 EXAMINE RESULTS

A window opens with a detailed summary of the boot results. BootRacer splits the boot process into four stages: Pre-boot (your PC's initial hardware checks), Windows Boot, Password Timeout (the time while your PC is sitting at the login screen waiting for you to act), and Desktop. Each section is given a time in seconds, but only the Windows Boot and Desktop times count toward the final Boot Result calculation [Image C].

» Beneath this is a reference to 'Explorer Startup Delay: 10 sec=Not optimal'. Confusingly, this isn't a reference to the 'Windows is almost ready' clock you saw earlier but is actually a reference to an arbitrary 10-second delay Windows inserts into the post-logon process before start-up apps are processed.

» Leaving it in place ensures the desktop is ready to use more quickly, but if you needed to prioritize a specific start-up app or you're running a fast machine where the delay isn't as important then you can reduce the delay to just two seconds by clicking Optimize followed by 'Remove Delay: 10 Seconds'.

» BootRacer will tell you this is a premium feature, but by clicking the 'Use Free Bonus' button it will set the delay to just two seconds instead.

» You'll be prompted to 'Restart Computer' to apply these changes to see whether BootRacer measures any improvement. If you want to go further and disable the delay entirely, you can open Registry Editor (press Win + R, type regedit and hit Enter) and navigate to the following key:
Computer\HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\Serialize

» Double-click the Startupdelayinmsec value in the

FINER START-UP CONTROL

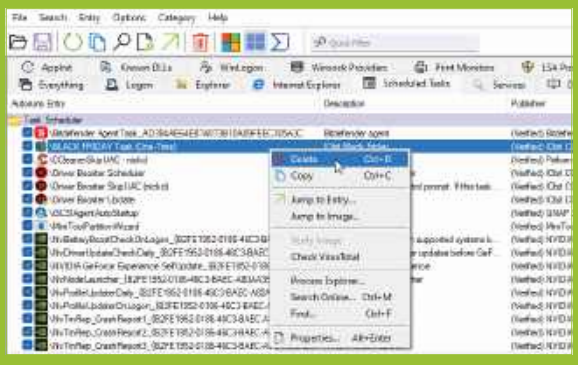

If you want more control over startup, including viewing hidden parts of the start-up process, download Autoruns from https://docs.microsoft.com/en-us/sysinternals/ You'll find it under Process Utilities. Unzip its contents and then right-click either Autoruns.exe (32-bit) or Autoruns64.exe (64-bit) and choose 'Run as administrator' to launch it.

You'll see a huge list of items appear under an all-encompassing Everything tab. The trick is to quickly navigate away from this tab to focus on specific start-up types. The Logon tab closely approximates what's shown by BootRacer but may well include additional entries. You'll also see that those items you have disabled or removed in BootRacer are shown as being enabled here—that's because BootRacer suppresses this list in favor of its own settings (uninstall BootRacer, and they will all flood back).

Some tabs are empty but focus on Scheduled Tasks as well as Explorer and Services tabs. Look for items marked in red or yellow—try disabling these first by unticking them. If there are no consequences, you can later right-click them to delete them. Another benefit of Autoruns is that it allows you to check start-up entries for possible malware—right-click an entry and choose 'Check VirusTotal' to upload the file in question for a quick online malware scan.

right-hand pane and change its value from 2,000 (milliseconds) to 0. Reboot and benchmark the change again.

›› Should you find the tweak adversely affects boot times—that's more common on older PCs and those running from slower hard drives—then simply click Optimize or Change again under Explorer Startup Delay, but this time click 'Restore Factory Settings', which simply deletes the Startupdelayinmsec value in Registry Editor to reinstate the default 10-second delay.

4 REVIEW APP START TIMES

Thankfully, the main bulk of BootRacer's Startup Optimizer tools are free to use. First, you need to enable the feature—look for the 'Enable Control' button in the carousel beneath the boot result (if it's not showing, click the first tab beneath the carousel to bring it back into focus). After clicking this, tick the box under 'Enable Startup Control', then switch to the 'Enable Measuring Startup Time' tab and click 'Restart your PC and Analyze Results' to reboot and benchmark your startup once again.

›› This time, when you log back in, you'll see BootRacer has instigated its own delay on launching start-up apps—once the timer counts down to zero, it starts loading them in sequence (and measuring their start-up times). Once finished, you should discover this simple tweak on BootRacer's part may have already

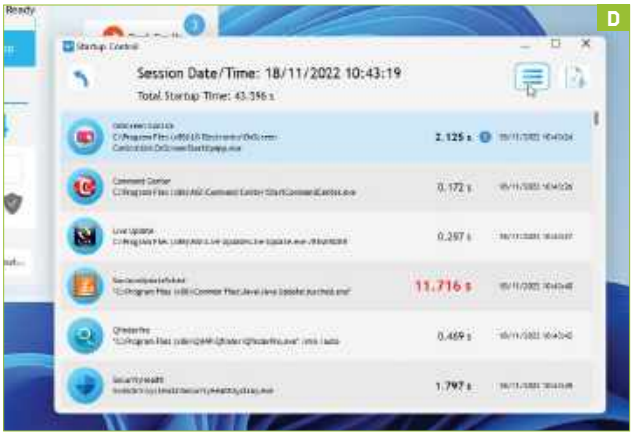

delivered noticeable savings on top of any recorded by disabling Windows' 10-second delay.

›› You'll also see exactly how many apps are set to start with Windows—by clicking the 'Check Results' button followed by 'Find Slowdown...' under 'Startup programs times:' you'll be shown a list of apps in the order they loaded along with the time it took for them to start [Image D]. Look out for items marked in red as particularly slow starters, one notable example is Java's update scheduler.

5 OPTIMIZE STARTUP

Once you've made a note, close the window and click the Startup Control button to start managing the list. Start by disabling those you don't need starting with Windows—they will either take longer to start, or you'll need to remember to check for updates manually. To do this, simply remove the tick next to the program and you'll see its status set to Disabled.

›› If you're convinced you don't need a start-up program, you can delete it permanently from the list—roll your mouse over it and click the bin icon [Image E], although we suggest first disabling it and monitoring startup over several days to ensure it's not actually needed for any reason.

›› As you trim the list, click 'Undo...' if you accidentally remove a start-up item and want it restored. The next thing to do is change the order in which items are loaded during startup. If there's an app you want to be available quickly, click the 'Set Order' button and use the arrows to promote it up the list.

›› You may notice an 'Anti-Malware' button in the top right-hand corner—this is a Premium feature, so ignore

it; your existing security software should already be offering protection against this, but you can also vet your start-up items for free using another start-up manager (see the box).

›› Once you've finished tweaking, reboot your PC and your start-up times should improve further. You may also notice, particularly on lower-powered machines, that your PC is faster in day-to-day use too. In future, you can disable BootRacer or leave it to monitor your start-up times to ensure you don't slip into bad habits. Don't panic if, occasionally, boot times are much slower; this can occur when installing Windows updates.

6 OPTIMIZE YOUR SERVICES

Start-up apps is just one area where you can trim the fat from system resources. Services—low-level programs used by both Windows and third parties such as your security tools—is another area where there may be opportunities to reclaim some precious resources.

›› There are two ways to tweak services—on an ad-hoc basis whenever you need additional system resources for certain tasks, from playing games to ripping media, or more permanently. The former option is best for most because it allows you to squeeze every last bit of performance out of your system by temporarily disabling both background tasks as well as services that can quickly be restored once you're done with a single click.

›› This type of app is commonly referred to as a game booster for obvious reasons, but can of course be used whenever you need that extra burst of speed. There are dozens to choose from, and they all do much the same thing. One simple-to-use tool is Smart Game Booster (https://gamebooster.itopvpn.com).

›› The Free version contains all the key functionality you need—just click Boost before playing a game. You'll be shown the performance uplift, typically around 10 percent, and when you're done, click Restore to bring things back to how they were.

7 MANUALLY TWEAK YOUR SYSTEM

One-click apps are all well and good, but if you want to exert more fine control over your system, we recommend trying another lightweight (and free) tool called Wise Game Booster (www.wisecleaner.com/wise-game-booster.html).

›› The program opens to the empty 'My Games' tab, where you can store shortcuts to any application you like. Beneath this will be a list of 'issues' the program has found, divided into three sections. You can choose to optimize them all from here, but it's a good idea to review what kind of optimizations to expect by exploring the System Optimizer, Process Optimizer, and Service Optimizer tabs in turn.

›› System Optimizer contains no fewer than 35 'recommended' optimizations, some of which may already be set by your system. These optimizations are best reviewed carefully and set as permanent tweaks to your system. Note, if you click 'Resume All' to undo them, all optimizations, not just those you've set in the program, are removed.

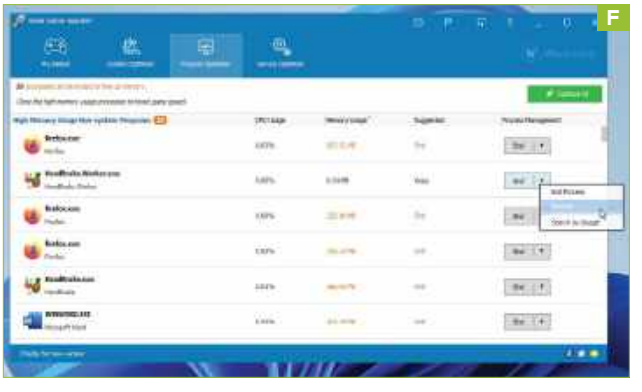
F

© PROCESS OPTIMIZER

›› Of more potential use for quickly clearing resources for gaming is Process Optimizer [**Image F**], which identifies memory-intensive processes that Wise Game Booster thinks are worth ending—clicking 'Optimize All' will end these promptly with no 'restore all' button, so review carefully before diving in.

›› Finally, Service Optimizer identifies what the program describes as 'irrelevant' services that can be safely ended to provide an extra boost—again an 'Optimize All' button exists to speed things up.

8 REIN IN RUNAWAY PROCESSES

Another way to allocate your PC's finite resources is by managing your processes' priorities. The higher the priority, the more demands a process can make on the system. By default, most processes are set to run at 'Normal', but there are options to give processes more (Above normal, High) or less (Below normal, Low) priority compared to others. You can set these manually via Task Manager's Details tab (right-click a process and choose 'Set priority'), but it's a fiddly process and there's potential to bring your PC grinding to a halt, and not simply by accidentally setting a process to Realtime, thereby giving it a higher priority than Windows.

›› Thankfully, some tools allow you to better manage the competing demands of processes. Though they are more suited to low-powered PCs that struggle whenever you open more than a couple of browser tabs, these tools can help higher-end PCs too.

›› Our favorite tool is Process Lasso (https://bitsum.com/). The annotation opposite provides you with a key run down of its features, while a Pro license gives you complete control of your system (prices start from $25 for a one-year, single-PC license, but $40 buys you a lifetime license), the free version has all the core functionality you need if you're happy for the app to do the heavy lifting for you.

FIX START-UP DELAYS

The boot process consists of several phases: first is Pre-boot when your hardware runs start-up checks and the boot manager waits to hand over to your selected operating system. Next is the Windows loading process (Windows Boot), then a pause at the login screen, followed by the post-logon process (or Desktop), as Windows loads your user profile along with apps configured to run at startup.

If your pre-boot time is longer than five seconds, this may indicate several things: you may not have enabled Fast Boot in your system UEFI—there may be a reason for this (dual-booting with Linux), otherwise, give it a go. There may also be a delay if you have a dual-boot menu that pauses for a lengthy period before booting the default OS.

Delays in the Windows Boot portion of the boot process may indicate errors on the drive (press Win+X, choose 'Terminal (Admin)' or 'Command Prompt (Admin)', then run 'chkdsk c: /f /r /x', selecting yes to perform the check and repair on your next boot). If you've not yet upgraded your boot drive to SATA or NVMe, you might also want to check to make sure it has recently been defragged (type defrag into the Search box).

How to Colorize B&W Photos in Photoshop

YOU'LL NEED THIS

PHOTOSHOP CC & A CREATIVE CLOUD SUBSCRIPTION
A suitable image to edit.
Multicore 64bit processor, 2GHz or faster
Windows 10 64-bit (version 1909) or later
8GB RAM (16GB recommended)
DirectX 12-compatible GPU less than seven years old
(4GB RAM recommended)

ONE OF THE BENEFITS of a Creative Cloud subscription is that you get almost constant updates to Adobe's creative apps. You can also access beta versions of parts of the software, which is exactly what we're doing here with Photoshop.

The famous photo-editing app has had neural filters, which use machine learning networks to achieve their aims, for a while now, but they're an incomplete set and some are still in beta. They're all interesting to some extent, but one in particular jumps out—the ability to restore and colorize old black-and-white photos with a few clicks. Long-time readers may remember that this is something we tackled in *Maximum PC* back in 2017 but, thankfully, it's now much easier.

It's not the only neural filter though. You can do dreadful things with portraits using Smart Portraits, from simply changing the direction of lighting to forcing a face to smile, or become angry, when it's not. The Makeup Transfer tool can take the makeup applied to one face and slap it on another, which is surely going to lead to an outbreak of online clowns and drag queens, while Super Zoom attempts to increase the resolution of a cropped image. There's a lot to go through here, so let's get stuck in. –IAN EVENDEN


A

1 FIND THE FILTERS

You don't need to download the beta version of Photoshop to find these, though you can certainly do that (see panel on page 71) if you want the latest test filters. Just make sure you've updated to the latest Windows version of Adobe's image-editing app (version 2023, which for us is 24.0.0).

» You'll find the neural filters, including some marked as beta, under the Filters menu. They open in a sidebar of their own [**Image A**], allowing you to flick them on and off, although we noticed that sometimes they wouldn't be selectable from the menu if the app had just started up. However, waiting 10 seconds or so usually cleared the issue.

» Before you can use the filters, however, you need to download them. Some of them can be rather large, with Photo Restoration coming in at over 800MB. Click the cloud-arrow icon to start the process.

» New filters are constantly being added to this section of the app, and there's even a wait list pane, in which you can express your interest in potential future neural filters. As we write, there are AI noise reduction, image relighting, AI face generation, and a filter for adding long-exposure effects to images of water in this section. Expressing your interest pops up a box in which you're asked to explain why you'd like to see a particular effect added to Photoshop.

» Note that some of the filters require cloud processing, and these are marked as such. If you're using something simple like Skin Smoothing, a 'Processing on device' message appears at the bottom of the screen while they're working. Switch on something more complicated, like Smart Portrait, and this changes to 'Processing in the cloud...' with a commensurate increase in the amount of time it takes to finish. An option at the top, behind a three-dot menu, allows you to switch GPU acceleration on and off.

2 OLD PHOTO RESTORATION

The new neural filters make restoring old photos a much easier process. It's not quite one-click but it's certainly much faster than identifying matching shades of gray and mapping new colors to them.

» There are two neural filters that will be able to help: Photo Restoration, which is in beta, and Colorize, which isn't. The latter handles the process of turning an old monochrome photo into a color one, while the former does noise reduction and scratch removal, and also adds sharpness and contrast to the image.

» If you've got old images you'd like to fix up in this way, you'll probably need to scan them. The rule with all 'AI' applications is that the more data you give the app to work on, the better the result, so if you're scanning an


B

C

image, make sure to do so at 300dpi or even higher if you're going to want to produce larger prints. Using small images destined for the web rarely gives an acceptable result, and you can always make a smaller file later on if you need to.

›› Get your image into Photoshop from your scanning software, and go to Filter → Neural filters. This opens the sidebar on the right. From here, you can choose Photo restoration, which at the time of writing, was at the bottom of the list and marked with a Beta symbol. Simply sliding the switch to the right will process your image on your PC, and the time it takes depends on the size of the file and the power of your machine. You can see the result of our first run in [**Image B**].

›› While the Neural filters interface has a before and after button, you can click to see the changes made to your image, there's no live divider as in Luminar, so it's hard to capture the differences in the picture in a screenshot. The filter has increased the contrast in the image and sharpened it up considerably. The faces of the people are clearer, but the app hasn't detected them as faces. If Photoshop thinks faces are present, a new slider marked Enhance Face appears, which does precisely what you'd expect it to, but it's not worth using on an already clearly defined face. Small, indistinct faces aren't detected, so it's only really usable on faces that occupy the middle ground—indistinct enough to be improved, but with enough facial features to be recognizable to the app.

3 REMOVE THE SCRATCHES

Scratches still appear on the image, so let's do something about them. There are two main sliders in the Photo restoration filter interface: Photo Enhancement and Scratch Reduction. Changing the position of either of them restarts the processing, meaning you'll need to wait to see the result.

›› The normal Photoshop Dust and Scratches filter can make an image soft, as it's applied indiscriminately across the whole surface, but the neural filter isolates just the scratches. Using the normal Healing tool will give you a similar effect, but takes much longer. Sliding the Scratch reduction slider to the right intensifies the effect, though you won't see it increasing in real time due to the processing breaks involved.

›› A large scratch in the upper right of our picture was removed by setting a value of around 50 on the slider, but smaller ones remained. On some photos, this filter can identify cables or other thin structures as scratches, wiping them from existence and attempting to plug the gap in a similar way to the Content-aware fill tool. Watch out for this happening, as there's currently no way to selectively apply the neural filters without making a selection first, though the result can be output to a new layer, allowing you to blend the original back in.

›› If you drop down the Adjustments section below the two main sliders, there are further options. You can reduce image noise, and tone down JPEG artifacts if your image has been resaved many times. More interesting is Halftone artifacts reduction, which can help if your image has been scanned from a newspaper. Again, all these fixes are applied simply by pushing the slider to the right, and waiting for the processing to complete.

D


E

4 COLORS

The Photo restoration filter can be combined with the Colorize filter to complete the modernization of an image. The quickest and easiest way to do this is to leave the Auto Color Image box checked, and let the local processing do its magic [Image C]. Dropping down the Adjustments section, which is located below this, allows you to choose a profile to guide the colorization process.

›› Leaving the profile selection set to None and the Auto box ticked gives us a rather muted colorization. The neural filter has picked up that the walls of the house are covered in moss and, therefore, should be greenish over the more yellow sandstone, but has missed a large part around one of the windows. The people in the center are naturalistically colored but still appear somewhat retro in nature.

›› Overall, it's a fairly suitable treatment for the subject matter and age of the photo. We can do better though. By choosing a profile, we can restart the processing with new parameters, which will produce a different result. The profiles present at the moment tend to lean in one direction: retro. Every option, apart from 'None', has retro in its title. It's reasonable to think Adobe will add more options over time, but for now, it's retro or nothing.

›› Retro high contrast completely ruins the image, introducing too much red into the colorization process [Image D] but Retro dark gives a much more convincing result. The profiles can be further altered by using the sliders beneath—these control everything from the strength with which they are applied to the color saturation and the precise color balance.

›› There are also sliders for the reduction of color artifacts and image noise, should they be present in your chosen photo.

5 GO MANUAL

Written under the small version of your original image at the top of the Neural filters panel are the words 'Manually color image'. Clicking in the image thumbnail opens a color picker, which allows you to place control points at the points you click, where the color you've chosen will be applied.

›› It isn't an intuitive process, but we managed to place points across the front of the house and attached a light pink color to make the stones less green while leaving the plant life its usual verdant hue.

›› Once you've created multiple points, selecting them and then clicking on the color swatch below allows you to change the color associated with that point, and a small slider allows you to increase or decrease the effect [Image E].

›› Even while you've got a control point selected, the color balance sliders below continue to affect the entire image, so if you find too much magenta creeping in, you can negate the effect by mixing in green.

6 PORTRAITS

Opening a black-and-white portrait, you can use the face detection algorithm to not only put the color back but to adjust the sitter's features too. [Image F] shows a neural filter colorized version of an image we'd made grayscale on the left and the original color image on the right. It's a pretty good match, except that the subject now looks like he's wearing lipstick and it hasn't got his skin tone quite right.

›› We then decided to use the Smart Portrait neural filter, which processes your data in the cloud, to mess with this face further. We slid the Happiness and Facial Age sliders around, and experimented with changing the head position, to create the portrait you see on the left in [Image G].


F


G

›› We found that, while effective, the Eye direction slider made the subject's eyes too fuzzy, detracting from the overall effect. We put them back where we found them, looking over to the right.

7 MAKEUP

As our portrait now looks like it's had a little help in the lipstick department already, we thought we'd try the Makeup Transfer neural filter to see if we could make things worse. Or better. So we opened an image of a young lady in full Dia de los Muertos-style makeup to use as a source, and let it rip. The neural filter produced the image you can see [**Image H**] and doesn't contain any sliders with which to customize the result. The filter has concentrated on the lips and eyes—common enough places to find makeup—but has all but ignored the makeup on the rest of the face, apart from a purple smudge near the right-hand temple, which might have come from a flower in the source image instead.

›› Opening a picture of a more subtly made-up model and using it as the source was even less successful, with almost none of the predominantly white facepaint carrying over.

8 DEPTH BLUR

One final thing you can do with the Neural filters is to blur the background, much like a smartphone camera's portrait mode. We found it to be less effective, but more natural-looking than the phone, with the filter able to automatically recognize the subject. There's also a manual setting in which you can place control points to define the bits you want to be in the foreground and sliders to control the amount of blur and to add haze, as well as color controls for saturation and balance.

›› The result is [**Image I**], with the original version on the right. The effect isn't extreme, and we had to pump the Blur slider all the way to the right. The filter may have realized that the car isn't far behind the subject, so wouldn't be blurred with a large-aperture camera lens.

›› Neural filters are hit-and-miss at the moment. They aren't a one-click fix for restoring old photographs, and can't change a portrait as if the sitter had been in a different position. They are best used sparingly, to touch up images that are close to being right already.

DOWNLOADING BETA VERSIONS

For the most cutting-edge version of Photoshop, you need to get hold of the beta. Luckily, this is easily available through Creative Cloud, with a Beta versions section on the left-hand side of the central hub app.

Once you've found it, installation follows the same process as any other Creative Cloud app—you won't be able to install beta versions of applications you aren't already subscribed to. They're essentially pre-release versions of the app suite, used to preview new features and for Adobe to gather feedback on its ideas before setting them loose on the public.

Neural filters were once part of the beta program, but have now moved over to the stable release version of the app, even though some of the filters themselves are still in beta.

The beta version opens with a different splash screen from the normal release, but getting it to open can be tricky if you already have Photoshop running. We selected it from the Start menu and Photoshop 2023 itself popped up in its place. It wasn't until we'd quit that app that we were able to get Photoshop Beta to run. You can confirm which version you're running by dropping down the Help menu and choosing About Photoshop, as the beta app looks similar to the mainstream one once the splash screen is out of the way. The beta app will be called Adobe Photoshop (Beta) and will have a release number slightly higher than the 'normal' version.

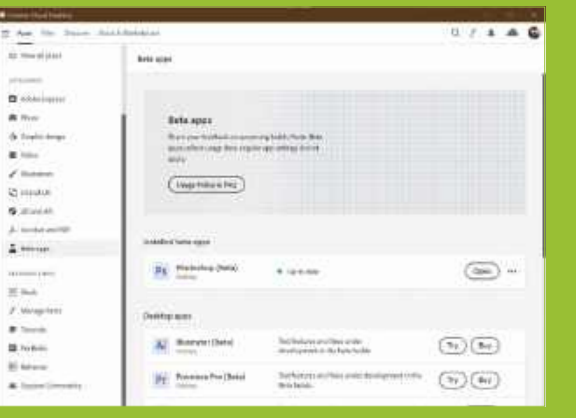

The beta also has a feedback button on its interface, which looks rather like a chemical flask that's half full (or half empty, if you're a pessimist). This opens a dialog box in which you can communicate your thoughts on the beta app and its new features to Adobe.

You also get a Beta Overview screen that tells you which features are currently being tested—Photoshop Beta receives updates just like any other Creative Cloud app—and what they do. You can click on the feature names to learn more about them.

Feedback can be given in a few forms, but the easiest is the happy/sad faces you'll notice on the neural filter screenshots—clicking on the sad face opens a window in which you can add a comment to your disappointment, and even send the image you were working on back to Adobe so they can see how terrible their work is.

LAB NOTES

GUY COCKER, EDITOR-IN-CHIEF

Nvidia nonchalance

A reader is perplexed by the 4000-series and so is the market

THIS MONTH, WE RECEIVED a letter from a reader, BJ Koho (see page 94), who didn't understand the fuss we've made over Nvidia's RTX 4000-series. He says that although the new cards are fast, that's the least we should be expecting—his issue is that this speed comes with such hefty financial, thermal, and power demands, that he doesn't understand the point of them in the current economic climate.

True, we were impressed with the RTX 4090 in our review in the December 2022 issue, and I was blown away by DLSS 3 4K gaming in my build in the Holiday 2022 issue, I immediately went out and bought a card for myself. A luxury, for sure, but how else is a PC gaming nerd meant to keep himself warm this winter!?

Joking aside, we've been rather less effusive about the RTX 4080. Jeremy Laird wrote that the card was straight up "bad" in his column in the December issue, while our review in this issue (see page 74) finds the card difficult to recommend overall. The price-to-performance ratio is so poor compared to the 4090, and it's effectively priced so close to the flagship card, it just doesn't offer good value for money.

Nvidia's strategy feels copied directly from the Apple playbook. Like Nvidia, Apple is a market leader with fantastic flagship products that are easy to recommend for those with the money. But Apple hobbles its non-flagship products and prices them so close to its flagships that you're pushed into buying the pricier model. Take the iPhone 14, for example—a fine phone, but lacking the Always-On Display, Dynamic Island and 120Hz refresh rate of the Pro model, which costs $200 more. Apple knows people want those features, so they will pay the extra. Same with the RTX 4080 and 4090.

Nvidia's latest cards seem to be in ready supply.

What's interesting is that, like our reader BJ Koho, the market seems to have had a lukewarm response to these new cards. As I write this, one week on from the 4080 launch, I can find an abundance of both 4080 and 4090 partner cards available on online retailers. Sure, Founder's Edition cards are sold out on Nvidia's store, but it feels as though Nvidia has discovered pretty quickly that there's a limit to the number of people who'll spend $1,200+ on a GPU right now.

SAM LEWIS

Staff Writer

When you move home, one of the most alarming things you notice is just how much unnecessary garbage you acquire through the years. Some of which hasn't seen the light of day for years. It was like unboxing time capsules from my childhood.

Anyway, amidst all the bits and bobs, I stumbled upon some hidden gems. One of which was my original Nintendo DS, which thankfully still had *Mario Kart* inside. It also still had the original charger so I didn't have to hunt around for half an hour to get a quick hit of nostalgia.

The same couldn't be said for the next gem I pulled out, an old IBM ThinkPad T42. If you're unfamiliar with this brick of a unit, it's essentially the Nokia of the laptop world. The problem, though, is I couldn't find the charger, so that's on the to-do list. I'm sure there's a lot of my old school work from back in the day still on there.

Last but not least, I found my beloved iPod Nano 3rd gen. This was, and still is to this day, the best gift I've ever received. Ever since my dad got the original iPod Classic, I wanted one, so to get one of these when I was young blew me away. That statement alone should show you my age! This tiny, quirky device is the reason I'm such a music nut. So if you have a load of old boxes in storage somewhere, take a look through the archives, you might just find some hidden gems.

Even with 1,152 zones, AOC's latest miniLED monitor is pretty clunky.

Editor's Pick:
AOC Agon Pro AG344UXM

Time for miniLED?

SIX MONTHS ago in this very section of the magazine, I reported that I was finally ready for miniLED to take my money. At last, I'd seen a great display with miniLED tech. The catch was that it was a laptop screen, not a desktop panel, specifically the Asus ROG Flow X16.

I've recently tested a few more cutting-edge miniLED desktop monitors and they are all, without exception, a bit rubbish. One of them will be getting the full review treatment next issue, so I won't spoil the anticipation even if I'm sure the *Maximum PC* readership will spot it instantly.

Instead, I give you the AOC Agon Pro AG344UXM, a 34-inch ultrawide gaming panel with some spectacular specs. This 21:9-aspect panel packs a serious punch. It's DisplayHDR 1000 certified, hits 1,000 nits peak brightness, packs a miniLED backlight with 1,152 dimming zones, has a 1ms pixel response, and 170Hz refresh. But like every other miniLED monitor I've seen, it's just a bit borked.

With local dimming enabled, it does some odd things. For instance, you can see dark borders on the inside of bright objects. More generally, you get weird inconsistencies in backlight power wherever there's a transition between bright and dark objects.

It's not at all hard to understand why. Even with 1,152 zones, you're looking at roughly 52 by 22 zones (can someone out there do the math, try as I might I can't come up with a 21:9-aspect integer resolution that works out to exactly 1,152 zones) and each zone measuring something in the region of three-quarters of an inch wide and tall. That's still quite a large block for each zone and there is simply no way of mapping that to even align perfectly with the edges of larger objects, much less precisely light up something small and detailed correctly, like a star field.

So, compromises are unavoidable. In broad terms, either you crank the relevant backlight zones up enough to do justice to bright objects and put up with some blooming and light bleed in darker areas. Or you prioritize black levels and leave brighter objects looking a bit dull. Of course, in practice, you can aim anywhere between those two extremes, but you're never going to get it exactly right.

What's more, as soon as you have a backlight running its own algorithms, you then immediately also run into response and syncing issues. An LED backlight has a different response profile than an LCD panel. So, again, it's never going to be possible to perfectly align the backlight response to the LCD panel response. Cue more compromises.

Admittedly, the drawbacks of miniLED tech are less obvious playing games or watching movies than on the Windows desktop. But even then, the experience never quite adds up to the sum of the typically eye-popping specs of miniLED monitors. Meanwhile, on the desktop I find the backlight algorithms are invariably clunky enough that I want to disable local dimming, leaving me wondering why I spent $1,000-plus.

At least, I would be wondering where my money had gone had I bought a miniLED monitor. But I've seen enough now to make that extremely unlikely. The reality is that there's no substitute for proper per-pixel lighting when it comes to HDR performance. It's go OLED or go home, at least until micro-LED becomes more affordable.**–JL**

$1,799, us.aoc.com

Reviewed...

74 GeForce RTX 4080 Founders Edition

76 Velocity Micro Raptor Z95

78 Sapphire Radeon RX 6700

80 Asus ROG Strix Z790-A Gaming WiFi D4

82 Asus ZenWiFi XT9

84 Samsung Odyssey Ark

87 SteelSeries Aerox 5 Wireless

88 SureFire Harrier 360

89 HyperX Pulsefire Haste Wireless

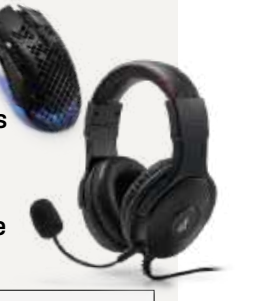

90 Call of Duty: Modern Warfare II (2022)

92 Authy vs. Google Authenticator

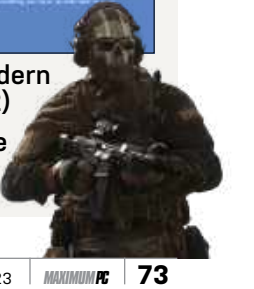

The Nvidia RTX 4080 Founders Edition has the same cooler as the 4090, but a lower power draw.

GeForce RTX 4080 Founders Edition

Does it measure up to the RTX 4090?

WE WERE INCREDIBLY impressed with Nvidia's Ada Lovelace architecture and the first graphics card to use it, the RTX 4090 (see December 2022, page 74). The 4090 offered massive generational performance gains with only a slight price increase relative to its 3090 predecessor. For the second outing, the RTX 4080 delivers a sizeable increase in performance over the old RTX 3080, but it inherits the substantially higher price point of the RTX 3080 Ti.

It's not that the 4080 is a bad card, but previous generations have kept the 80-class models in the $500–$800 price bracket. The 3080 Ti got away with its high price because it launched during the worst graphics card shortage we've ever witnessed thanks to chip shortages and cryptocurrency mining. After the painful availability and value of GPUs all last year, the 4080 launch price feels like a slap in the face to gamers. It's very much out of reach for mere mortal PCs.

That's unfortunate, but it's also hardly surprising. Early indications are that Nvidia and its partners sold every 4080 card priced under $1,500 at launch, although in the weeks after, plenty of partner boards have been available. As a result, it appears as though demand won't remain high and the 4080's price may plummet in the coming weeks—a repeat of what happened with the RTX 3090 Ti.

GARNERING GOODWILL

And the price may need adjustment, depending on how AMD's RX 7900-series parts stack up (see our RDNA3 feature on page 32). Nvidia might have the technological advantage with DLSS 3 and superior ray-tracing hardware, but if AMD can take the lead in traditional gaming performance while staying reasonably close in ray-tracing, its lower prices could garner a lot of goodwill and potential customers.

Looking at the performance, the RTX 4080 beats the RTX 3090 Ti by 10–40 percent, depending on the game, with a 19 percent lead at 4K on average—that was the fastest graphics card available just two months ago! But the 3090 Ti was also selling for $1,099, so you have to pay more for the upgrade. Also at 4K, the RTX 4080 leads the 3080 Ti by 36 percent, trounces the 3080 by 50 percent, and churns out 32 percent more performance than AMD's current best, the RX 6950 XT.

The 4080 delivers on the efficiency front at least. While the Total Board Power (TBP) is 320W, across our test suite, it averaged 290W at 4K, dropping to 250W at 1440p and 220W at 1080p. Even with manual overclocking, which improved performance by seven percent, power use averaged 310W and only a few games broke the 320W figure.

But the problem is the huge 71 percent jump in generational pricing and its proximity to the RTX 4090. The 4090 costs 33 percent more than the 4080, but it's also 35 percent faster at 4K. If you're willing to spend over a grand in pursuit of gaming performance, why not go the whole hog and buy a 4090? You can even argue that the 4090 represents better value overall. That's especially true for professional workloads such as Blender, Octane, and V-Ray, where the 4090 is 30–45 percent faster and the extra VRAM can prove beneficial.

It's an interesting change in tactic for Nvidia. It went big with AD102 and the 4090, leaving a rather large gap compared to the 4080. The 3080 was much better value than the 3090, and the same goes for the 2080 over a 2080 Ti. Titan cards, meanwhile, offered a few professional extras, double the VRAM, and a slight performance boost along with a massive price premium. Now the 4080 costs as much as a Titan Xp, the 4090 delivers similar bang for your buck, and we can't help but wonder where the eventual RTX 4070 will land. **–JARRED WALTON**

VERDICT 7

GeForce RTX 4080 Founders Edition

+ ONE GIANT LEAP Second-fastest GPU; impressive efficiency; DLSS 3 and RT.

– ONE SMALL STEP Huge price increase; big performance gap; RDNA 3 looks promising.

$1,199, www.nvidia.com

BENCHMARKS

	RTX 4080	RTX 4090	RTX 3090 Ti	RX 6950 XT
10 Game Average	124 / 83	**144 / 107**	108 / 70	104 / 63
Borderlands 3	166 / 89	**216 / 130**	138 / 80	144 / 78
Control (DXR)	94 / 46	**135 / 67**	79 / 40	53 / 25
Cyberpunk 2077 (DXR)	63 / 30	**86 / 44**	50 / 25	29 / 13
Far Cry 6	150 / 109	**153 / 133**	140 / 91	150 / 94
Flight Simulator	77 / 76	82 / **81**	**87** / 70	82 / 55
Forza Horizon 5	151 / 126	**172 / 149**	115 / 88	136 / 98
Horizon Zero Dawn	176 / 114	**177 / 156**	162 / 102	170 / 93
Minecraft (DXR)	83 / 44	**106 / 62**	69 / 35	37 / 18
Red Dead Redemption 2	157 / 104	**179 / 138**	133 / 94	122 / 84
Watch Dogs Legion	125 / 88	**132 / 113**	109 / 71	121 / 69

Best scores are in bold. All testing conducted with a Core i9-12900K, MSI Pro Z690-A WiFi DDR4, 2x16GB DDR4-3600 CL16, 2TB Crucial P5 Plus M.2 SSD, Cooler Master MWE 1250 Gold V2. Scores are average framerates at 2560x1440 / 3840x2160 ultra, with ray tracing enabled in Control, Cyberpunk, and Minecraft.

SPECIFICATIONS

Architecture	AD103
Lithography	TSMC 4N
Boost clock	2505MHz
GPU cores	16384
Memory	16GB GDDR6X
TFLOPS FP32	48.8
Bandwidth	717GB/s
TBP	320W
Connectors	HDMI 2.1, 3x DisplayPort 1.4

Velocity Micro Raptor Z95

Top-end components, overclocked for your pleasure

WE REVIEWED a Velocity Micro PC in our June 2022 issue–the Raptor Z55, which at the time was a performance powerhouse with a Core i9-12900K CPU and RTX 3080Ti GPU. Six months later, the cutting edge has advanced, thanks to the arrival of Intel's Raptor Lake processors and Nvidia's Ada Lovelace graphics cards. Velocity Micro has issued a new Raptor, and the Z95 is the first pre-built system we've reviewed using this latest generation of components. The eagle-eyed among you will spot it has a similar spec to the machine we built last month (Holiday 2022 issue), incorporating an Intel Core i9-13900K, RTX 4090, and 32GB DDR5 RAM. Given that was a machine we loved for 4K gaming and creativity, is this pre-built the way to go for those who don't want to go about making their own?

Velocity Micro's Raptor line specializes in offering high-end components encased in a roomy chassis with overclocks pre-applied. In our case, that's a boost up to 6GHz on the Intel Core i9 processor, which resulted in a small performance gain on the single-core test (see below) but a 12 percent boost to the multi-core result. As you'd expect from a pre-built, Velocity Micro also does a good job of syncing the RGB lighting across components and keeping things quiet (although the GPU fans do sometimes ramp up under load).

Let's start with the good. This is the fastest PC you can buy right now thanks to that killer combo of the Intel Core i9-13900K and RTX 4090. We're talking 4K Ultra settings at 60 frames per second in all but the most demanding games such as *Cyberpunk 2077*, while enabling DLSS 3 in that title takes it up to a 110fps average. It also comes with a 1,000W PSU, giving the machine the power needed for those overclocks. Basically, if you want the fastest gaming PC available in 2023, don't want to build it or overclock it yourself, and have a generous budget, you'll love the Velocity Micro Raptor Z95.

The issue for anyone comfortable doing the building and tinkering is that you're paying quite a premium for this machine as configured at $5,699. The similar machine we built for the Holiday 2022 issue was nearly $2,000 cheaper. Sam's budget RTX 4090 build in this issue (see page 16) offers similar gaming performance for under $3,000—nearly half the price of this Raptor Z95. As lovely as the Raptor Z95 is, that's hard to justify.

Given the price, buyers also have a right to expect components to be more future-proofed than they are. The Asus ROG Maximus Z690 HERO is a great board, but the Z790 version ($100 more at retail) has an extra USB 3.2 port, better PCI express lane configurations, and higher DDR5 clock speeds. Likewise, 32GB of DDR5-6000 RAM will be fast and capacious enough for most users, but 64GB could be expected on a machine at this price. The Raptor Z95 also comes out just ahead of the expected release of PCIe 5.0 SSDs, but the Samsung 980 Pro used here feels a little outdated with the faster 990 Pro now out in the wild.

From the perspective of someone who regularly builds PCs, this is disappointing. If you want a pre-built, then Velocity Micro has crafted a beautiful machine in the Raptor Z95 and has done an excellent job maximizing its performance. Be aware you're paying a premium over building it yourself, but no one said being on the cutting-edge was cheap. **–GUY COCKER**

VERDICT 7

Velocity Micro Raptor Z95

+ HIGH VELOCITY Incredible performance; clean design; 1000W PSU.

– DRAG CITY High price; no Z790 mobo or 64GB RAM; faster SSDs available.

$5,699, www.velocitymicro.com

BENCHMARKS

	ZERO-POINT	
Cinebench R23 Single-Core (Index)	2,247	2,304 (3%)
Cinebench R23 Multi-Core (Index)	35,457	39,743 (12%)
CrystalDisk QD32 Sequential Read (MB/s)	7,057	6,874 (-3%)
CrystalDisk QD32 Sequential Write (MB/s)	5,303	4,953 (-7%)
3DMark Fire Strike Ultra (Index)	24,528	24,432 (0%)
Cyberpunk 2077 (fps)	78	84 (8%)
Cyberpunk 2077 RTX (fps)	43	43 (0%)
Metro Exodus (fps)	140	139 (-1%)
Metro Exodus RTX (fps)	105	107 (2%)
Total War: Three Kingdoms (fps)	100	104 (4%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of our cover build from our Holiday 2022 issue, featuring an Intel Core i9-13900K, MSI GeForce RTX 4090 Gaming X Trio, ASUS ROG Maximus Z790 Hero. Corsair Vengeance 32 GB (2 x 16 GB) DDR5-5200 CL40, PNY XLR8 CS3140 1TB and NZXT C850 80+ Gold. All game tests were performed at 3,840 x 2,160 at the highest graphical profile.

SPECIFICATIONS

CPU	Intel Core i9-13900K @ 6GHz
GPU	PNY GeForce RTX 4090
RAM	32GB (2x 16GB) DDR5-6000
Motherboard	Asus ROG Maximus Z690 HERO
Storage	2TB Samsung 980 Pro NVMe SSD
Cooler	Velocity Micro 360mm Liquid CPU Cooler
Case	Velocity Micro GX6 Plus ATX
Front I/O	3.5mm audio, 3.5mm mic, 2x USB 3.2 Gen 1, USB 3.2 Gen 2 Type-C
Rear I/O	2x Thunderbolt 4 USB Type-C ports, 7x USB 3.2 Gen 2 ports (6x Type-A + USB Type-C), 2x USB 2.0 ports (2x Type-A), HDMI port, Wi-Fi Module, Intel 2.5Gb Ethernet port, 5x Gold-plated audio jacks, Optical S/PDIF out port, BIOS FlashBack button, Clear CMOS button
PSU	EVGA SuperNOVA 1000W Gold
OS	Windows 11 Home
Dimensions	20.25 x 20.5 x 8.25-inches
Weight	30lbs

The Raptor Z95 is a quiet and attractive build, but the RGB fans are hidden.

Sapphire Radeon RX 6700

Strong competition for Intel's Arc

AMD ORIGINALLY launched its Navi 22 GPU with the RX 6700 XT in March 2021. We expected there would be an RX 6700 non-XT arriving shortly after, using harvested chips that couldn't meet the requirements of the higher tier part. Then months passed with no such part in sight, and AMD went so far as to refresh its desktop parts with a higher binned and higher priced RX 6750 XT (see review, August 2022 issue).

Now that GPU cryptocurrency mining is dead and graphics cards aren't immediately selling out, AMD and its partners are finally ready to release the long-awaited 6700 10GB. Why didn't we get this GPU last year? Probably because AMD was using the same GPU in the mobile RX 6700M. That's too bad, though it might not be too late.

This is exactly the sort of card we expected to see, with 36 Compute Units out of Navi 22's maximum configuration of 40 CUs. One of the six 32-bit GDDR6 controllers also gets fused off, leaving a 160-bit interface. Along with the missing memory controller goes 16MB of Infinity Cache, leaving 80MB—still plenty for 1080p and 1440p gaming, which is where the 6700 10GB shines.

Of course, this is also the same old RDNA 2 architecture from over two years ago, which has certainly lost some of its luster now that we're ready for RDNA 3 GPUs to arrive (see page 32). Then again, at roughly triple the cost, plenty of people are still looking for something to fit a midrange budget rather than blowing a huge chunk of cash on the latest and greatest hardware. But there will be compromises, like the mediocre ray-tracing performance that's common for RX 6000-series cards, as well as the lower-quality video encoding hardware.

Still, as far as bang for your buck is concerned, the 6700 10GB sits in the upper tier of current GPUs—only AMD's RX 6600-series cards clearly beat it, though the 6700 XT also edges past it for those who might want a bit more muscle. It also depends on current prices, which are in a constant state of flux at present. The RX 6650 XT launch MSRP was $399, a product of the crypto mining days, but you can now find those cards starting for as little as $259. The RX 6700-series prices have been dropping as well, so it's important to check around to see where things stand before making a decision.

Performance from the 6700 10GB across our test suite came out ten percent faster than the RX 6650 XT, 14 percent ahead of the RTX 3060, and five percent slower than the Arc A770. That does include three-ray tracing games; if you drop those from the list, the 6700 10GB ends up being 25 percent faster than the RTX 3060 and 11 percent ahead of the Arc A770. Intel's drivers continue to improve, but they're still several years behind AMD and Nvidia in terms of maturity.

Power and efficiency are strong points for the card, and while Sapphire officially lists a TBP (Total Board Power) of 220W, in testing we saw a maximum power use of less than 200W while gaming, and an average power use of 186W. That's about the same as we measured with the RX 6650 XT, which runs at higher clocks and thus loses some efficiency.

Ultimately, Sapphire's prospects for the card come down to price and availability. Getting 25 percent more memory than the 8GB Navi 23 cards can help in a few games. Meanwhile, bandwidth and compute performance aren't significantly higher than the RX 6650 XT, and the 6700 XT adds another 2GB of memory and may only cost $25 more. The same logic that might lead someone to choose the 6700 10GB over the 6650 XT would also suggest going for the 6700 XT. And that might be precisely why it's late to the party. **–JARRED WALTON**

BENCHMARKS

	RX 6700 10GB	Arc A770 16GB	RTX 3060	RX 6650 XT
10 Game Average	81 / 57	**85 / 61**	71 / 50	74 / 50
Borderlands 3	**104 / 70**	94 / 67	73 / 52	103 / 67
Control (DXR)	67 / 43	**86 / 57**	75 / 46	59 / 37
Cyberpunk 2077 (DXR)	26 / 23	**50 / 30**	35 / 30	23 / 18
Far Cry 6	**123 / 87**	104 / 81	97/ 71	119 / 82
Flight Simulator	69 / 47	**72 / 51**	69 / 49	67 / 46
Forza Horizon 5	**98 / 81**	86 / 74	72 / 61	88 / 69
Horizon Zero Dawn	**121 / 92**	93 / 83	100 / 78	115 / 83
Metro Exodus (DXR)	92 / 38	**120 / 55**	83 / 34	79 / 36
Red Dead Redemption 2	82 / 67	**93 / 74**	69 / 56	72 / 59
Watch Dogs Legion	**89 / 63**	78 / 60	64 / 48	75 / 53

Best scores are in bold. All testing conducted with a Core i9-12900K, MSI Pro Z690-A WiFi DDR4, 2x16GB DDR4-3600 CL16, 2TB Crucial P5 Plus M.2 SSD, Cooler Master MWE 1250 Gold V2. Scores are average framerates at 1920x1080 / 2560x1440 ultra, or 1920x1080 medium / ultra in DXR enabled games.

VERDICT **7**

Sapphire Radeon RX 6700

+ GRAND ENTRANCE 10GB memory; decent value; proven drivers

– LAST CALL Limited options; late to market; 6700 XT.

$329, www.sapphiretech.com

SPECIFICATIONS

Architecture	Navi 22
Lithography	TSMC N7
Boost clock	2495MHz
GPU cores	2304
Memory	10GB GDDR6
TFLOPS FP32	11.5
Bandwidth	320GB/s
TBP	220W
Connectors	HDMI 2.0b, 3x DisplayPort 1.4a

The Sapphire Radeon RX 6700 card is a more compelling option than the 6650 XT that launched in early 2022.

Mixing white, black, and silver colors doesn't quite work, but we like the Pac-Man iconography.

Asus ROG Strix Z790-A Gaming WiFi D4

Good value Raptor Lake motherboard

LAST MONTH, we reviewed the Asus ROG Maximus Z790 Hero ($630), a top-of-the-line board for the latest generation of Intel CPUs. It had loads of features including support for PCIe 5.0 M.2 SSDs, DDR5 RAM, and a shedload of the latest I/O ports. However, in our cover build, we pitted it against this Asus ROG Strix Z790-A Gaming WiFi D4 ($380) and found that, for gamers at least, there was no performance difference between the two boards. With the RTX 4090 ($1,600) set to drain the budget of many a PC gamer, could this Strix model be a potent pairing for the cheaper Intel Core i5-13600K?

This is a striking-looking mobo with lots of white and silver detailing on the various heatsinks. Design aside though, this board has all the important ports and sockets that gamers demand, including two sets of ARGB/RGB headers for AURA lighting, four PCIe 4.0 M.2 slots, and a PCIe 5.0 x16 slot for any future graphics cards that use the newer standard.

The four SATA ports might not be enough for those still using 3.5-inch or 2.5-inch hard drives, though, while the decision to put the AIO pump headers between the top M.2 SSD slot and the CPU socket makes things pretty busy if you're using a water cooler. There are thoughtful touches, though, including an accessible PCI Express release latch so you can easily get a GPU in and out.

I/O options are also plentiful, and we particularly appreciate the DisplayPort and HDMI outputs. There are eight USB-A ports, two USB-C ports, and a 2.5Gb Ethernet port. The board also supports WiFi 6E, has digital audio out, and a USB port for performing BIOS updates, plus a BIOS flashback button. The board can also provide 30W PD for cases with compatible USB-C ports, and the I/O panel also has Pac-Man chasing a ghost, a nice little touch for its gaming audience.

As we covered in our previous issue, benchmark performance from the board is excellent, and there isn't much to distinguish between this and a DDR5 board in most tests. The Cinebench R23 Multi-Core test showed a nine percent performance hit over a DDR5 board when using the Intel Core i9-13900K processor but for gamers, our tests showed no difference that was outside a margin of error. That's good news for anyone looking to save money, although future games may take advantage of DDR5 RAM and PCIe 5.0 storage, neither of which this board supports.

If you want to max out the overclocking potential of your newest Intel CPU, this isn't the board for you, as it doesn't have the power stages needed. On the flip side, if you're happy for the BIOS to do the work, then Asus's interface will allow you to push the CPU and RAM safely. What's remarkable is that with a decent DDR4 kit, like the PNY XLR8, it can keep up with the lower bandwidth offered by DDR5 in real-world tests, which is a great way to save money. It may not be future-proofed, but for performance on a budget, the Asus ROG Strix Z790-A Gaming WiFi D4 is a strong recommendation. **–GUY COCKER**

VERDICT 8

Asus ROG Strix Z790-A Gaming WiFi D4

+ A-GRADE Excellent performance; great I/O; GPU easy-release button.

– E-GRADE No DDR5/PCIe 5.0; non-uniform styling; fiddly AIO header location.

$380, www.asus.com

BENCHMARKS

	ZERO-POINT	
Cinebench R23 Single-Core (Index)	2,247	2,171 (-3%)
Cinebench R23 Multi-Core (Index)	35,457	32,164 (-9%)
CrystalDisk QD32 Sequential Read (MB/s)	7,057	7,022 (-0.5%)
CrystalDisk QD32 Sequential Write (MB/s)	5,303	5,295 (-0.2%)
3DMark Fire Strike Ultra (Index)	24,528	24,519 (0%)
Cyberpunk 2077 (fps)	78	77 (-1%)
Cyberpunk 2077 RTX (fps)	43	43 (0%)
Metro Exodus (fps)	140	140 (0%)
Metro Exodus RTX (fps)	105	104 (-1%)
Total War: Three Kingdoms (fps)	100	101 (1%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our Zero Point system consists of our Build It machine from the December 2022 issue (page 16), consisting of an Asus ROG Maximus Z790 Hero motherboard, Intel Core i9-13900K CPU, MSI GeForce RTX 4090 Gaming X Trio GPU, Corsair Vengeance 32 GB DDR5-5200 RAM, and PNY XLR8 XS3140 PCIe 4.0 NVMe 1TB SSD. The ASUS ROG Strix Z790-A Gaming WiFi D4 system used the exact same components, apart from PNY XLR8 Gaming Rev RGB 32GB (2 x 16GB) 3200MHz DDR4 RAM. XMP was enabled on motherboard BIOS but no other overclocking performed.

SPECIFICATIONS

CPU support	Intel Socket LGA1700 for 13th Gen and 12th Gen Core Processors
Chipset	Intel Z790 Chipset
Memory	4x DDR4 DIMM, Max 128GB
Display ports	DisplayPort, HDMI port
Storage	4x M.2 slots, 4x SATA 6Gb/s ports
Ethernet	Intel 2.5Gb Ethernet
Wireless	Wi-Fi 6E, Bluetooth 5.3
USB	USB-C 3.2 Gen 2x2, 3x USB 3.2 Gen 2 (2x Type-A, 1 x Type-C), 4x USB Type A 3.2 Gen 1, 2x USB Type A 2.0 ports
Audio	ROG SupremeFX 7.1 Surround Sound High Definition
RGB	Aura Sync
BIOS	256MB Flash ROM
Form factor	ATX
Size	12in x 9.6in (30.5cm x 24.4cm)

Asus ZenWiFi XT9

Affordable mesh system that doesn't disappoint

IF YOU'VE STRUGGLED to extend your Wi-Fi signal to all corners of your home, you may have seen the appeal of a mesh network. This tech combines two or more devices to improve Wi-Fi coverage. Traditionally, however, there's a trade-off between price and functionality.

Choose an entry-level system like Google Wifi or Amazon eero, and you're restricted to managing your network via an app while giving up advanced tools like VPN or server capabilities. At the other end of the scale are mesh systems with all the bells and whistles but huge price tags, such as Netgear's Orbi AXE11000 system that we reviewed in May 2022.

Asus's ZenWiFi XT9 places itself in the middle, with this two-pack setup costing under $500. The sacrifice in top-line speed comes from opting for WiFi 6 rather than WiFi 6E (with an additional 160MHz channel in the 6GHz spectrum) but, elsewhere, Asus doesn't cut corners. Like all ZenWiFi systems, you get a pair of powerful, feature-packed routers.

Setup is best done through the app, but while this boasts plenty of advanced tools for on-the-go tweaking, you'll want to explore the router's browser interface for maximum control. Here, you'll find all the bells and whistles you need, including adaptive QoS, VPN (both server and client support), and built-in security and parental controls, both of which are free.

The XT9 also supports both wireless and wired backhaul connections for managing the mesh network. However, unless both devices are within the length of an Ethernet cable of each other, you'll be using the wireless backhaul, which eats up over half the available bandwidth via its own dedicated 5GHz channel. This leaves a theoretical maximum of around 3,000Mbps for regular traffic, which soon drops away as you move further from the main router. In our tests, bandwidth on the 2.4GHz channel halved between router and satellite alone.

We set up the XT9 in a two-story duplex, installing the main router at the top of the stairs and the satellite in the room below. Our subjects were an iPhone (WiFi 5) running the WiFi SweetSpots app, and our desktop PC with a WiFi 6 card situated away from the satellite. We used WiFi Analyzer and inSSIDer to test the XT9's capabilities, backed up with a 1GB read/write test using LAN Speed Test.

We found that while the XT9 falls short of Netgear's top-of-the-line Orbi system, it held its own with similarly priced devices. We saw read/write speeds of over 100Mbps over the local network on the desktop, while the iPhone recorded speeds of around 450Mbps when standing next to the main router, and 350Mbps in the same room as the satellite. Though speeds dropped as we moved away from the router, we enjoyed a 30Mbps connection in our backyard. Moving the satellite to another downstairs room doubled this to 60Mbps, a reminder that satellite placement is critical when configuring your mesh system.

If we've any criticisms of the XT9, it's the setup process, which isn't as simple as some of its rivals. We also struggled to set up our PPPoE internet connection, although a search online suggests this problem affects only a minority of ISPs. Once things were up and running, though, the connection was fast and stable.

While it can't match the Orbi AXE11000 in terms of raw performance, the ZenWiFi XT9 is still a major step up from older WiFi 5 systems, and a great choice for those looking for a competitively priced mesh system that doesn't compromise on advanced features. **–NICK PEERS**

VERDICT **8**

Asus ZenWiFi XT9 mesh system

+ STREAMING HEAVEN Excellent performance; advanced features; competitively priced.

– BUFFERING HELL Not the easiest to set up; PPPoE connection issues.

$499, www.asus.com

SPECIFICATIONS

Wireless	574 + 2402 + 4804, 2.4 GHz, 5GHz-1 & 5GHz-2
Security	WPA3 (with WPA2 fallback)
Key specs	Quad-core 1.7GHz CPU, 512MB RAM, 256MB flash storage
Ports	2.5Gbps (WAN), 3x 1Gbps (LAN), USB 3.0 per device
Extras	Alexa compatibility, VPN server/client, BitTorrent client, security, parental controls
Dimensions	6.29 x 2.95 x 6.35-in (per device)
Weight	1.64lb (per device)

The ZenWiFi XT9 will boost your Wi-Fi network's range, capacity, and speed.

The Samsung Odyssey Ark offers a truly immersive gaming experience.

Samsung Odyssey Ark

A magnificent monitor for mad money

© UNSPLASH

IF THERE'S A COST of living crisis going on, Samsung and Nvidia haven't gotten the memo. This 4K Odyssey Ark costs $3,500, and the RTX 4080 or 4090 realistically needed to drive it are over $1,200. But if you want to experience the best of PC gaming, they're perfect bedfellows. The Ark is packed with features, it transforms from landscape to portrait, displays three different sources at once, and boasts TV-like smart apps. But none of those features is a reason to buy it—the reason is that it's just a great gaming monitor.

The Odyssey Ark is a colossal 55-inch display, with a 4K panel and 1000R curvature—a pretty deep bend on a monitor this size. It takes at least two people to unbox it, screw it into the stand and lift it onto a desk, but it's pretty easy to put together overall. It glides between landscape and portrait (or 'Cockpit') modes with little effort, although we'd have liked it to be automated at this price.

Like Samsung's premium TV offerings, connectivity is managed through a separate One Connect box, with power, networking, and AV all running through one cable to the TV. As someone who has reviewed a fair share of TVs and monitors over the years, I love this feature. It helps that the One Connect box has four HDMI 2.1 inputs to support loads of devices, although DisplayPort is an odd omission for a PC monitor.

Using the monitor in its standard widescreen configuration for work or gaming, you'll find it simultaneously ridiculous and awesome. Ridiculous because it offers basically zero privacy for work-related tasks, even if you use Split Screen in Windows 11. It's not a professionally calibrated monitor, but it does offer a respectable level of color accuracy for video and photo editors, but the thought of doing such work on it is almost laughable, as you'll have to move your entire head when switching your gaze between the timeline and the preview window. For gaming, it's fantastic, with great colors and contrast, particularly running in HDR. The 165Hz refresh rate is welcome, although the lack of official G-SYNC is disappointing at this price (AMD FreeSync is supported).

Otherwise, the Odyssey Ark is packed with features, most of which we found a distraction. Take Cockpit mode—you can run three sources at once, but you have to arch your neck to see anything at the top, and you lose HDR support when running multiple sources. Samsung's smart apps are great, but who wants to use Xbox Game Streaming on what is clearly a PC monitor? We also found the huge Ark Dial remote redundant, in that it makes controlling the display more confusing than the standard remote control.

The Samsung Odyssey Ark is the ultimate display for truly immersive 4K gaming thanks to its immense size and visual fidelity. It has its flaws, with many major features including Cockpit mode feeling redundant. But when you've played a game in landscape on the Ark, you'll never want to go back. **–GUY COCKER**

VERDICT **7**

Samsung Odyssey Ark

ARK-ANGEL Great color and contrast; 4K @ 165Hz; excellent HDR.

ARK-DEMON No G-Sync; redundant Cockpit mode; awful Ark Dial remote.

$3,499, www.samsung.com

SPECIFICATIONS

Screen size	55-inch
Screen curvature	1000R
Aspect ratio	16:9
Panel type	VA
Brightness	600cd/m2 (Typical)
Contrast ratio	1,000,000:1 (Static)
HDR	Quantum HDR 2000
HDR format	HDR10+
Resolution	3840x2160
Response time	1ms
Viewing angle	178°(H)/178°(V)
Color gamut	95% DCI Coverage
Refresh rate	165Hz
Mini LED local dimming	1,056 zones
Variable refresh rate	AMD FreeSync Premium Pro
Inputs	4x HDMI 2.1, 2x USB A, Ethernet
Power consumption	190W (Typical)
Dimensions (with stand)	117.5 x 110 x 37.9cm (WxHxD)
Weight (with stand)	41.5kg

SUBSCRIBE TODAY

Find your next great read

On iOS & Android!

Subscribe today and get instant access on your iPad, iPhone or Android device. Additional print options available.

www.magazinesdirect.com/B2XMPC

SteelSeries Aerox 5 Wireless

Does this ultra-light mouse live up to its price?

WHEN WE SEE A MOUSE costing $140, it can be difficult not to have incredibly high expectations. It doesn't help when that mouse comes from a reputable company such as SteelSeries, which already has an arsenal of great peripherals under its belt. So will the SteelSeries Aerox 5 Wireless prove itself worthy or fail to live up to its lofty asking price?

The first thing we note is the consistently impressive unboxing experience SteelSeries offers. We are greeted by a lightweight braided USB Type-A to USB Type-C cable, a USB Type-C dongle, and a USB Type-C extension adapter, meaning we have plenty of great connectivity options. Of course, the most popular choice with this mouse should be its wireless dongle, but it's great to have Bluetooth and detachable cable options just in case.

The Aerox 5 mouse uses SteelSeries Quantum 2.0, a new wireless technology containing dual channel optimization, which results in no packet loss. It is ultra-low latency and feels extremely snappy and accurate, much like using a wired mouse, only without the common cable annoyances. Not only does this perform well, but it also optimizes battery performance, with SteelSeries boasting that this mouse can achieve up to 180 hours before it needs a charge. We weren't able to test this mouse for 7.5 days without a break—that would be quite a push—but in over two weeks of typical usage, we didn't have to charge it once. That said, charging isn't a problem, as 15 minutes of charging gets you over 40 hours of battery life.

This Aerox 5 wireless mouse falls in between the all-out Aerox 9 wireless and the Aerox 3 wireless. Design-wise, they are fairly similar, sporting an aerated honeycomb shell design that maintains its minimal weight and also reduces the occurrence of sweaty palms. Although the mesh stylization has some benefits, one major drawback of the design is that you can see right into the internals of the mouse and the RGB lighting hasn't got the best spread due to this design. You can dive into the SteelSeries Engine application to tweak and adjust this but it isn't as smooth as we'd like.

MEDIUM-TO-LARGE MOUSE

The Aerox 5 wireless is on the medium-to-large size and therefore suits both claw and palm-style grips. It's a comfortable mouse and the build quality is great, with no rattling or creaking going on. That's impressive seeing as it only weighs 0.16lbs, as such a lightweight structure can often result in a cheap feel.

What's even more impressive is that the mouse is so light, yet features nine programmable buttons. On the top are the left and right click, smooth scroll wheel and a profile cycle button for different presets (adaptable in the SteelSeries engine). On the left side is a five-button side panel, which features a custom-built up/down flick switch. This provides great versatility and can be mapped to suit your in-game needs. Not only could you use this to your advantage when gaming, but the programmable buttons are great for creative work, such as graphic design or video editing.

But the most important factor is how does this mouse perform? We already mentioned how reactive the Quantum 2.0 wireless functionality is, and that wouldn't be noticeable without the zippy SteelSeries TrueMove Air sensor on board. Co-developed by PixArt, this optical gaming sensor features tilt tracking, 18,000cpi, 400 IPS, and 40G acceleration. It's safe to say that it is a quick mouse perfect for gaming and FPS titles. Paired with its PTFE feet, it makes for a snappy mouse with a smooth glide.

There's no denying that it's an impressive mouse. It's lightweight, has a great build quality, stats that speak for themselves, a long battery life, plenty of customization within the SteelSeries Engine app, and it even features IP54-rated AquaBarrier that protects the mouse from dust, dirt, and oil. That's already a healthy list of positives for this mouse, but its one major drawback is the price. For $140, this mouse offers everything you want, but it's still expensive next to competition from HyperX. If it was just slightly cheaper, it would be a lot more exciting. –SAM LEWIS

VERDICT **8**

SteelSeries Aerox 5 Wireless

PLUS 5 Fantastic performance; lightweight; long battery life; nine programmable buttons.

MINUS 5 Expensive; below average RGB.

$140, www.steelseries.com

SPECIFICATIONS

Sensor	SteelSeries TrueMove Air
Sensitivity	18,000cpi
Polling rate	1,000Hz (adjustable)
Mouse switches	SteelSeries Golden Micro IP54 switches
Connectivity	Quantum 2.0 Wireless, Bluetooth, USB Type-A to USB Type-C
Weight	0.16lbs

SureFire Harrier 360

Bang for your buck or just thumping bass?

YOU SHOULDN'T always judge a product by its price tag. After all, there are plenty of good budget headsets around the $50 mark. The SteelSeries Nova 3 and Corsair's HS55 Stereo Wired are just two examples of well-rounded affordable gaming headsets at this price. But when a headset arrives promising high-performance 7.1 virtual surround sound, comfortable earpads, high-quality sound, and a highly-sensitive detachable microphone, we tend to be skeptical. So what's the story here?

Our first answer to this question comes when we pick up the headset. It's fairly light, which should help when wearing for long periods, but the choice of materials is poor and there's a sharp finish on the edges that, even for $50, is below what we'd expect. This is especially noticeable when you adjust the headset size on the forks. The quality is better on the headband and the earcups, as these contain a plush faux leather material. As soft as they are, the headset suffers from the fact that the clamping pressure is too tight. Even expanded to the max, the top of the earcups press too hard on your head, to the point where prolonged use can cause some discomfort.

The design isn't too outlandish, as long as you don't turn on the RGB functionality. It's an all-black plastic construction with the only splash of color (other than the lighting) being some subtle red stitching on the headband. Each earcup features a black mesh grille with SureFire's flame logo in the middle. Behind the grille is the RGB lighting, which has an even glow. Though SureFire is trying to adhere to the traditional gamer aesthetic, we'd have preferred less RGB and a better overall fit and finish on the headset.

The cable leading out from the left earcup allows you to control volume via a wheel, plus there's a mute switch for the detachable microphone, and a button to toggle RGB lighting. There are three options available; off (our favorite), static, and a soft cycle. Unfortunately, this controller has a rather cheap feeling to it as well.

There are no buttons or controls on the headphones themselves, which is good as you don't end up accidentally changing any settings. The cable is a 2m braided USB Type-A that isn't detachable from the earcup. We don't think you should expect this cable to last for years, but it feels thick and fairly durable overall.

GOOD CHARACTERISTICS

The SureFire Harrier 360 isn't off to a great start, but what really matters is the sound quality. The manufacturer describes this headset's audio as high-quality and, although it isn't exactly up there with the aforementioned SteelSeries and Corsair offerings from the same price bracket, the SureFire Harrier has some good audio characteristics.

Its virtual 7.1 surround sound works well for directional audio, which is important for competitive gaming. The high ends are fairly detailed and thankfully not painfully sharp with the volume cranked up. Midrange suffers with a resulting flat soundstage and unfortunately, this doesn't get much better with the bass. It's punchy, yet seems to overpower much of the detail provided in the top end. The headset works best within story-driven games but we recommend turning bass down in your game's settings if you can.

The microphone is detachable and the recording quality isn't too sharp, provided you aren't streaming and don't have it too close to your mouth. It does pick up a fair amount of background noise so bear that in mind if you are likely to use this headset for work meetings or with a mechanical keyboard.

The SureFire Harrier 360 is decent without being a bargain. It has an overall cheap feeling, and also lacks comfort and style. The problem here is that for a little bit more, you can get a more reputable and reliable headset that doesn't have the same quality control issues and we believe this would be a much better choice for most gamers. Although the sound quality isn't awful, its low to mid-range is muddy, with its saving grace being the top-end which provides good directional sound for gaming. Unfortunately, though, it's hard to recommend this considering the market surrounding it. –SAM LEWIS

VERDICT **5**

SureFire Harrier 360 Gaming Headset

JUMP JET Soft ear cups; effective virtual 7.1 surround sound.

CRASH LANDING Cheap materials; sharp edges; muddy bass; garish RGB.

$50, www.surefire-gaming.com

SPECIFICATIONS

Frequency response	20Hz-20,000Hz
Impedance	24 ohms
Connectivity	Wired USB Type-A
Design style	Closed back
Weight	0.81lbs

HyperX Pulsefire Haste Wireless

Marvelous eSports mouse at a decent price

HYPERX HAS made a name for itself making really good products that sell for reasonable prices, and even though ownership of the brand has moved from Kingston to HP recently, nothing much seems to have changed. The standard Pulsefire Haste (wired) mouse was released to appeal to the eSports crowd who wanted a fast and accurate mouse that was comfortable for long play sessions. Of course, a wire-free version of that mouse would seem to be an oxymoron, what with wireless introducing lag, dropout, and recharging concerns, but there are undoubtedly plenty of *Fortnite* and *Apex Legends* players out there who value the convenience of wireless over the small performance advantage offered by wired.

The Pulsefire Haste Wireless is distinguished by its honeycomb design, which allows air to flow through the mouse and to your hand. There's no fan in the mouse, so you don't notice the ventilation but the overall effect is that it makes the mouse extremely light, weighing in at just 61g, including the built-in rechargeable battery.

The mouse also features a symmetrical design, which makes it good for left- and right-handed users, although the two side buttons favor right-handers as they sit where the thumb rests. There are slight finger-shaped indentations on the primary mouse buttons, which feels comfortable, as well as a clicky-style mouse wheel and a button to switch between 400, 800, 1600, and 3200dpi so you can adjust sensitivity without having to go into Windows or HyperX NGENUITY software settings.

Speaking of software, HyperX's app can be downloaded from the Windows store and allows you to update the firmware, customize the buttons, and change the lighting effects on the scroll wheel. There are the usual options of solid, cycle, and breathing style effects, as well as the usual color options that can be synced to your other HyperX gear if you own any. We also like that you can save your lighting settings to the mouse so you can move between machines and keep your preferences—this is a pretty portable mouse after all. It's also worth mentioning the mouse is also available in white, if preferred.

The Pulsefire Haste Wireless comes with a 2.4GHz dongle, which is tiny and slots neatly into the underside of the mouse itself. We have to deduct points for the lack of Bluetooth compatibility, though, which would allow you to just grab the mouse and use it with most laptops, no dongle required. You could also just use a USB-C cable, one of which is included in the box, including a neat adaptor that adds a USB-A port for the wireless adaptor. And charging isn't something you should have to do a lot of, as this mouse lasts for 100 hours when not using the lighting. The wireless dongle was always extremely dependable during our review time, and with a 1,000Hz (1ms) polling rate, no normal human is going to notice any lag in their favorite first-person shooter or racing game.

There are also some great extras included in the box, including grip tape so you can make the mouse easier to hold depending on your preference, and extra PTFE skates, in case one of the four installed on the base of the mouse comes off or wears out over time. The packaging materials are also card and plastic, and easy to break apart and recycle once you've unboxed the mouse.

The Pulsefire Haste Wireless is another win for HyperX and continues its tradition of making great products for reasonable prices. In fact, at the time of review, we were seeing this mouse closer to a $40 price point, which represents great value. It's a really comfortable mouse to use, has a highly accurate 16,000dpi sensor, and comes with some excellent accessories.

We wish it had Bluetooth, and the lighting options are basic, but if you're an eSports gamer wanting a great value wireless mouse, this is it. –GUY COCKER

VERDICT **8**

HyperX Pulsefire Haste Wireless

MORE SPEED Comfortable design; accurate 16,000dpi sensor; great accessories included.

LESS HASTE No Bluetooth; basic lighting options.

$65, www.hyperx.com

SPECIFICATIONS

Included accessories	USB Wireless Receiver, 1.8m USB Cable, Grip Tape, Extra PTFE Skates
Buttons	6
Connection	Wired/2.4GHz Wireless
DPI presets	400, 800, 1600, 3200
Button durability	80 million clicks
Polling rate	1000Hz
Optical sensor	Pixart PAW3335
Warranty	2 years
Battery life	Up to 100 hours
Battery type	370mAh Li-ion polymer battery
Weight	61g

Try on the boots of Ghost and join the hunt for terrorists.

FIRST-PERSON SHOOTER

Call of Duty: Modern Warfare II (2022)

Task Force 141 is back for another global terrorist hunt

YOU CAN TELL a lot about a game by the things it makes you sit through before you actually get to play it. Launch the 2022 iteration of *Modern Warfare II* for the first time, and you have to scroll through endless license agreements and privacy policies, remember your account password to log in, then wait for the game to restart if you choose to play the campaign rather than the multiplayer.

There are even multiple developer and middleware logos to enjoy—it's an Infinity Ward game, but the studio had help—and you are left looking at a desert cliff while the shaders pre-compile on the first run. Birds circle, but you're there long enough to notice they're on a loop, fading out of existence before appearing again in the sky. The message here is that this game took a lot of work to produce.

Once you start playing, it's immediately noticeable how sharp the environments are and how, as a result, character models and vehicles can look artificial in comparison. Glenn Morshower from *24* and the *Transformers* movies makes a great General Shepherd, although Captain Price's beard is annoyingly distracting. Performance is generally excellent, recommended specs are low, and it will even run nicely on a Steam Deck as long as you're running Windows on it (because the anti-cheat system doesn't like Linux).

You're straight into the dusty boots of Ghost, a well-adjusted individual who wears a skull mask and gloves complete with skeletal outline, and for once the campaign mode isn't just practice for multiplayer. While it's still the kind of globe-trotting terrorist hunt we're used to, this time it all feels slower and more deliberate—your character needs to sprint more often to keep up with his squad. Some of this may be because guns are now faster, the latency between firing and hitting barely noticeable up-close, but still enough of a ballistic simulation to require thought and planning when firing from a distance.

Over 17 missions taking seven hours, you move from country to country and spectacular set-piece to spectacular set-piece. There's no such thing as a routine patrol, and you'll find yourself controlling a guided missile, swimming through a harbor, jumping in and out of moving trucks, creeping through the undergrowth in a ghillie suit, guiding a companion through an infiltration by watching him on CCTV, and of course raining hot explosive death on enemies from a circling airplane, all while keeping your allies safe.

It's fun, developing old ideas rather than adding many of its own, though the swimming level leaves you frustrated with the way it makes getting out of the water unnecessarily hard, your character bouncing up and down like an excited labrador by the side of the dock before he has the sense to grab onto the edge. By that time, the AI guards have

Gunsmith is the only way to level up your weapons.

Warzone 2 comes with a huge new desert map, Al Mazra.

Ghost, a well-adjusted individual in a skull mask.

Guide a companion through the mission.

A battle royale mode pitches players against each other.

begun lobbing grenades accordingly.

And as for that multiplayer, there are several differences and new features this time around. Reload canceling is gone, bunnyhopping has bitten the dust, and slide canceling is much harder to pull off, while changes to the cover system mean you'll still be able to shoot opponents behind walls if you're using the right gun. There's a new Apex Legends-style ping system for team comms, and a rebalance between killstreaks and scorestreaks.

The game launched with 13 multiplayer maps, but some are only usable in certain game types, some are used for 6v6 battles, while others are for the 32v32 'Ground War' games. Some expand in size to accommodate more players. Things will change pretty quickly after launch, with every new patch bringing in tweaks to balance weapons and fix issues.

Crossplay is mandatory for PC players, and with console owners able to turn it off, your mouse and keyboard skills will almost certainly be going up against gamepad-equipped troops. This has been controversial, as console players get aim assist options that PC players don't, though this negates the mouse user's innate advantage. The amount of aim assist is a topic of debate, with the general consensus being that it's too strong for everyone. Even worse is that split-screen multiplayer is off the table for PC users. Though it's rarer to see couch-based one-on-one shooters on PC, like the days of *Portal 2* co-op, it's nice to have.

The big news, however, is *Warzone 2.0*, the battle royale mode that pitches squads of players against one another in one enormous map. This time it's Al Mazrah, a desert environment with roads, rivers and mountains, several towns, and an airport. Other multiplayer maps touch upon its locations, but to get the full experience you have to play *Warzone*. Gunsmith is now the only way to level up your weapons, working your way through the upgrade tree for each type to unlock.

WEAPONS OF STRUCTURE

This weapon customization service has seen changes too. Only five mods are now allowed per gun, down from 10 in Vanguard, and you can share attachments across a weapons family—each pea-shooter has a group of related weapons you can unlock and share magazines or barrels, while other mods such as optics and ammo can be used on any weapon. The family tree structure means it now takes less time to unlock a particular weapon or attachment.

There's also DMZ, a new mode that's more PvE than PvP, and takes inspiration from games such as *Escape From Tarkov*. The idea is to get in, do whatever it is you came to do, and get out again while avoiding other players (who will need to watch out for bounties being placed on their heads if they attack you) and taking out the bots, who patrol exactly the areas you're trying to move through. It's a stealth co-op experience using the same huge map from *Warzone 2*, which players will get to know inside out.

Buying into a *CoD* game at launch is an expensive undertaking, with the standard game going for $70 and the Vault Edition, which offers the most additional content, including a complete season one battle pass, 50 tier skips, and a pack of skull-faced operators, lightening your wallet by one cent under $100. You would perhaps be wise to reconsider getting the latest *CoD* for the campaign alone, but it's much better value if you're intending to use the multiplayer mode, offering two years of intense content and battles.

It's a game of two halves, the 2022 edition of *Modern Warfare II*—it follows on from 2019's *Modern Warfare*, but has little to do with the 2009 game of the same name—is a multiplayer dream welded to an interesting, though somewhat familiar, campaign. –IAN EVENDEN

VERDICT **9**

Call of Duty: Modern Warfare II (2022)

SOAP Sharply rendered; gun-obsessed war against the usual suspects, or against your friends.

SOUP Missions are a bit samey; environments look better than character models; it's all very loud.

RECOMMENDED SPECS CPU, i7-4770/ Ryzen 5 1400. RAM, 12GB. GPU, GTX 1060/ RX 580.

$70, callofduty.com/modernwarfare2, M-rated

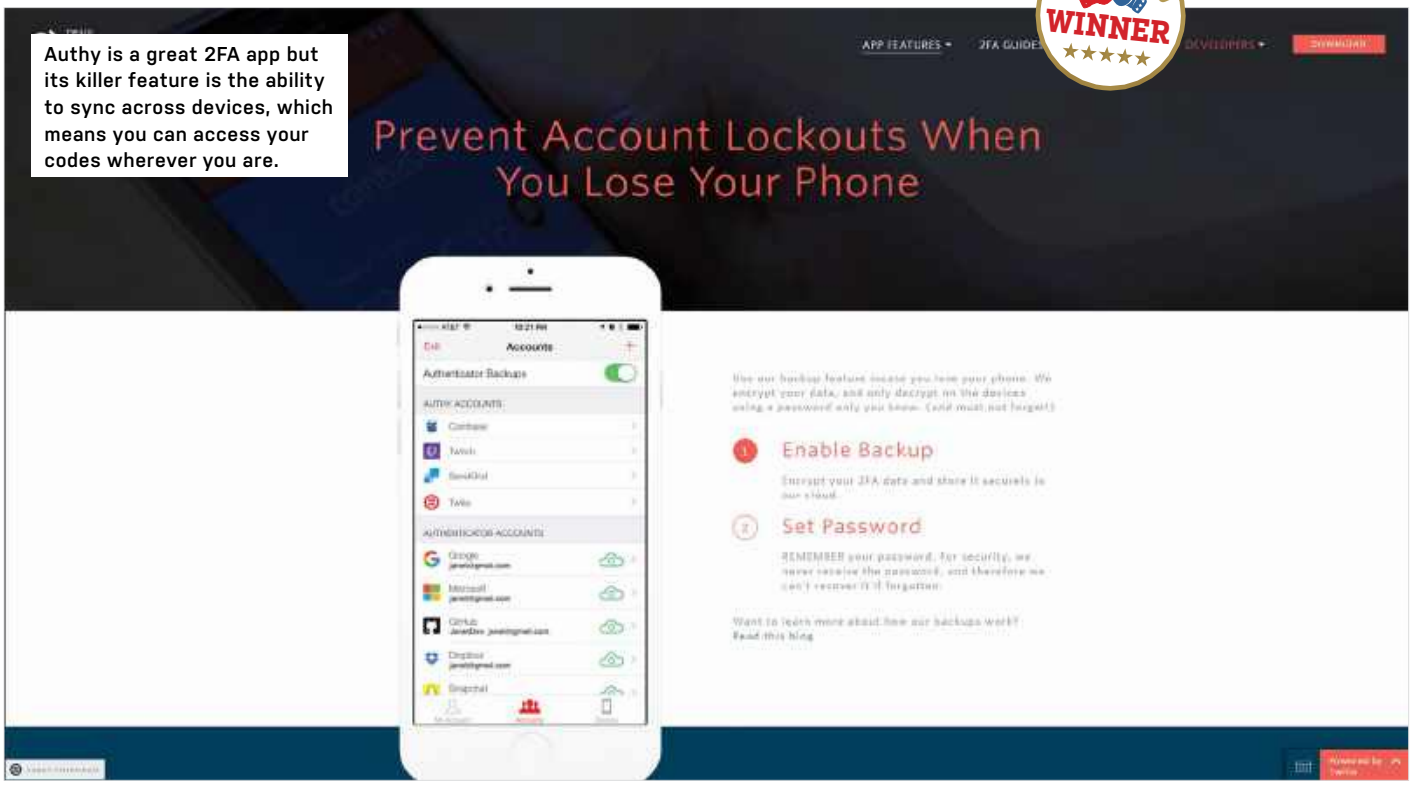

Authy is a great 2FA app but its killer feature is the ability to sync across devices, which means you can access your codes wherever you are.

Authy vs. Google Authenticator

Which is the better two-factor authentication app?

LOGGING IN to a service can be a real pain these days, in between being asked to change your password every six months, having to remember which street you grew up on as a child, or requiring you to check your phone for a text message code before you can gain access. So a request to set up two-factor authentication (also known as 2FA) might seem like yet another annoying layer between you and whatever account you're trying to access.

But it is really important to enable 2FA on any services that offer it, because it means that, even if your password is compromised, any malicious users won't be able to log in to your account without that code, which only you should have access to. This is especially important on any services where your financial information could be accessed or compromised, as 2FA acts as a barrier between your money and a nefarious user.

Hopefully, within the not-too-distant future, we'll be able to do away with passwords and 2FA altogether with biometric-based security. Microsoft has already enabled this for its own services via its Authenticator app. But for now, 2FA is an important security option for those services that offer it.

There's a de facto standard when it comes to 2FA, and that is Google Authenticator. Indeed, many services will explicitly tell you to use Google's app to scan the QR code that's displayed during the setup process, probably because it's assumed that over 50 percent of people will be using an Android smartphone, and if not, they will still have a Google account for which they can remember the password. But there are lots of different authentication apps, including Microsoft's aforementioned effort, LastPass Authenticator, and Zoho OneAuth, all of which do the job of providing you with a unique code when a 2FA-enabled website, app, or service requests one.

All authentication apps work in the same way—when you enable it on your chosen service (Facebook, for example), the site will create a QR code that you scan using your chosen app. The app will then display a six-digit code that you input into the service to confirm, and that's it. From then on, the app creates a new code every 30 seconds, and when the service requests it, you simply open up your app and input the code.

One service might ask you for a code every time, whereas others might only do it when it suspects something has changed—you might not have logged in for 30 days, for example, or you have tried logging in from an IP address that the service doesn't recognize. In this head-to-head, we'll be looking at the two most used and feature-rich options currently out there, Google Authenticator and Twilio's Authy.

Google Authenticator was the app that started it all and set the standard for two-factor authentication, and as such is a solid and dependable choice. The app is available on both Android and iOS platforms, meaning you'll need a phone or tablet nearby whenever you're asked for a 2FA code. Your codes can be moved over to a new device by exporting via a QR code, meaning if you ever move devices in the future, you can still access all of the services you need.

That's the theory, anyway—in practice, Google has sometimes been slow to update its app for new mobile operating systems, which has meant users have been unable to open the Authenticator app at all. That's unforgivable in this always-connected world, especially as Google offers no way to get your 2FA codes online or through a browser.

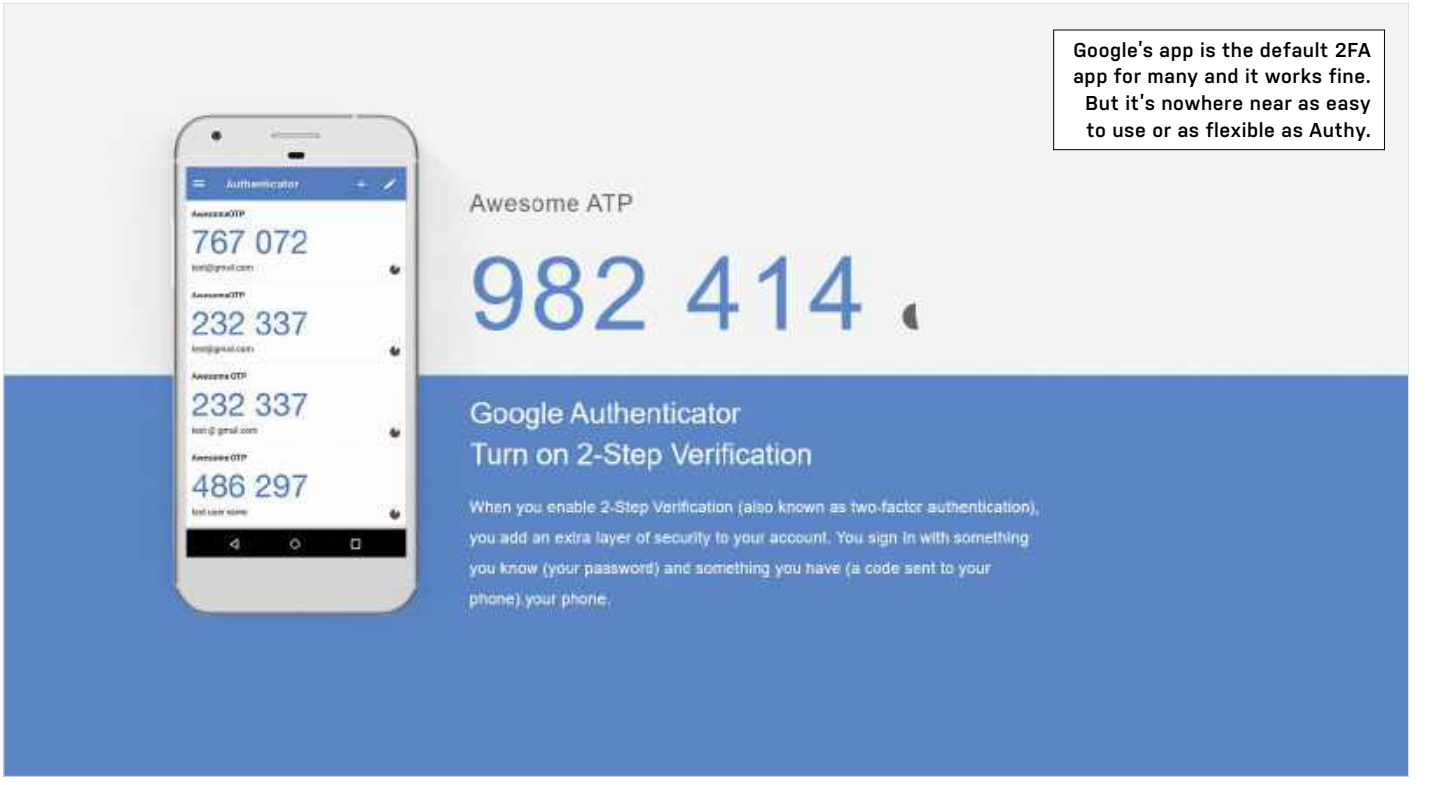

Google's app is the default 2FA app for many and it works fine. But it's nowhere near as easy to use or as flexible as Authy.

Google's interface could also do with a makeover, because as you build up a library of services using 2FA, it can be time-consuming to scroll through the list in the app. This is compounded by the fact that Google's app doesn't create icons for you, which would make it easier to distinguish between Twitch and Facebook at a glance, for example.

THE GREEN STUFF

Authy by Twilio does the same basic job as Google Authenticator, but it has developed over recent years to offer a lot more functionality compared to its big-name competitor. It still lists all of the services you add to it, but it will usually try and automatically add a logo when you scan the QR code, making it much easier to find a service when scrolling through the list on your phone. And if there isn't a logo available, you can use different colors of the default key logo to color code your services. I personally use green for my financial accounts—because of all the greenbacks in my possession, obviously.

That's not all when it comes to features with Authy, as it goes way above and beyond what's offered by Google Authenticator. Your accounts can be backed up online, meaning that if you lose your device, you can still log into Authy and access all of your 2FA codes. This also applies to running Authy on multiple devices at once—I run it on both my Android phone and Android tablet, which is handy for when I sometimes leave my phone in a different room. Authy also has a smart feature where you can simply hit 'approve' from a pre-existing version of the app when you're setting it up on a new device—you don't have to enter your password each time.

Authy feels more secure than Google Authenticator, as you can set it up with its own security settings. For example, you can set a PIN in order to open the app, or use your device's biometric security such as fingerprint or FaceID every time you want to access it. Not everything is great about the mobile app, however—there's no easy way to automatically sort all of your services in alphabetical order, so you have to do it manually. You sometimes have to go back and manually edit the names given to your services, as you'll find services that are just named after your email address or username with no other context provided.

Authy's killer feature, though, is that it runs not just on phones, but pretty much every device going, including Windows, Linux and macOS. That means you can access your codes directly from whatever device you're using, as well as copy and paste without having to use another device and type it out manually. While some people will feel less secure if they have Authy running on lots of different devices, you can always log out any other device from the app's settings, as well as see when they were last accessed.

Usually these software head-to-heads are relatively close, but in this instance there really is no contest between these two—Authy is the far better 2FA service. The most important thing to take away from this article is to enable 2FA where available, and if you prefer to use Google Authenticator or any other app to do this, then that's better than nothing. But really, there's no reason not to use Authy as your two-factor security option, because its features and flexibility make it as easy as possible to make every account you use as secure as it can be. –GUY COCKER

VERDICT 9 KICK ASS!

Authy

MASTER CODE Works across devices; auto online backup in case of device loss; logos make it easy to navigate.

CODE RED No way to sort alphabetically on mobile; basic interface; some services require manual renaming to make sense.

$free, www.authy.com

VERDICT 7

Google Authenticator

SAFE AS HOUSES Ubiquitous; simple to set up; easy transfer between devices.

ROCKY ROAD No online backup or sync; only works on mobile and tablet; no logos to distinguish between services.

$free, www.googleauthenticator.net

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> AIO Assistance
> Nvidia Hype Train
> 2 GPUs, 1 PC

PST Disorder

I've been a long-time fan of your excellent magazine, and I wanted to let you know I very much appreciated your Gmail vs. Outlook article (December 2022 issue).

It reminded me that there's one problem with Outlook that nobody really addresses, though—its infamous PST data files. What's particularly annoying is the fact that it doesn't matter how many emails you delete from how many folders, even permanently (meaning deleting emails from the deleted folder).

Even if you make an effort to list and delete emails that have a big attachment file, the PST file size doesn't get smaller. I've used the Compact Now feature in the Outlook Data Files Account Settings and it doesn't do anything.

The problem becomes crippling if you don't have a good internet connection. I frequently stay in small hotels outside the USA, and if you need to upload a 6-8GB file, it can take an eternity until you're able to close Outlook properly at the end of its uploading procedure. I want Microsoft to help us find a way to effectively reduce the size of Outlook's Data .pst file in a real and proportionate manner with the quantity and size of the emails that Outlook has. **–S. Nash**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: Believe me, I've spent many days of my life trying to solve the email conundrum. I don't delete emails, though perhaps I should, but I like to be able to quickly find press images and contacts from my old emails when I'm working.

I've tried setting up my own email server on my Synology NAS, which proved next to impossible—my plan now is to migrate everything to Yahoo!'s massive free email storage, but I can never quite bring myself to spend a weekend actually doing it.

Anyway, I digress from your problem, which is that your PST files are becoming unwieldy. Outlook stores email locally by default and syncs online, meaning if you delete an email locally, it also deletes it from the server. You could try reducing the amount of time Outlook goes back to download messages from—perhaps the last week or month would make things more manageable.

My next question would be why do you need your emails stored locally? It sounds like you do some traveling, so being able to answer emails when you don't have a connection, like on a plane or a remote hotel, makes total sense.

However, when you have an internet connection, you can still use your Outlook account online through a web browser, which will reduce the amount of local storage needed on your laptop. Or you could try a different app, such as Thunderbird or the Mail app that comes with Windows.

My experience with PST files is using them as a way of backing up my email into one convenient file for restoring later or keeping it as a failsafe. As a last resort, you could try uninstalling Outlook and reinstalling it—that might help rebuild a fresher (and hopefully smaller) version of the PST file.

Water cooling and cable tidying

Greetings from the deep south! It might not be practical, but I thought I'd send you a list of articles I'd love to see in future issues of *Maximum PC*.

'How to choose an AIO radiator size relative to your CPU'—currently, I'm using a Vetroo V5 to cool my Ryzen 3700x, but it's only getting down to 42°C at idle with an undervolt applied. If I want to venture into AIO cooling, what would be sufficient without spending $200 or more?

'Cable management tips when using a less-than-accommodating case'—I'm not gifted at doing this and have a Corsair Carbide Spec 01 case that drives me crazy every time I do a new build in it. I bought a Lian-Li LANCOOL 215 Midi-Tower for my current build, and it's much easier because it's so huge. Surely, there have to be some rules when dealing with cases that are less forgiving in their design?

Finally, an 'Overclocking/Undervolting for dummies' article. Thanks for doing such a great job on the magazine! **–T. Neumann**

↘ submit your questions to: editor@maximumpc.com

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: Thank you for your kind words on the magazine, and for all of the article ideas—as I keep saying every issue, it really makes my job a lot easier to have such generous readers sending great content suggestions!

All of these would make excellent 'How to...' or feature articles and I'm sure you're not the only reader to want to know about choosing your first AIO, cable management, or undervolting. However, in the meantime, I can answer some of the questions you raised in your letter.

On choosing an AIO radiator size—I'd say first of all that you don't need one for a 3700x, which is a good processor but definitely at the budget end of the market. Both Sam and I have used air coolers in recent builds, and a good one such as your Vetroo should be just fine.

If you really want water cooling though, I'd start with a 240mm one to save money. Look for budget brands, such as Arctic, Deep Cool, and Lian-Li—the latter you're already familiar with from your case. You can spend more on RGB fans or coolers if you want, but again, it sounds like a budget system, so why spend more than you have to?

Cable management is a tough one—there are times when I'm doing a cover build and I have to ask our photographer to help tidy up my mess, as he's much more organized than I am. Use the cable ties as much as possible, take advantage of space at the bottom of the case to hide excess cabling, and if everything's working, there's nothing wrong with slapping the back panel on and not worrying about it too much.

And finally, when it comes to overclocking, we've recommended a few times recently to take advantage of the automatic BIOS XMP and overclocking settings, which you should do if you haven't already. As for a deeper dive, I'm sure it's something that would go down well. Watch this space!

Nvidia nonsense

Quite frankly, I'm confused by all the positive press that Nvidia's RTX 4000 series is getting. Sure, the cards are faster—what would be the point if they weren't? But at what cost?

The MSRP of the two cards that have been released is high because Nvidia knows it can get away with charging ridiculous amounts (and more, thanks to eBay scalpers), while the high electricity bill of using one adds yet more to the cost of ownership. The thermals these things put out could heat a moderate-sized home in the winter, the power requirements have made 850W power supplies obsolete, and the fire hazards from overheating power connectors have been well documented.

This comes in the midst of a recession with raging inflation, high energy prices, and a climate crisis where we're being told to be responsible and reduce power usage. This series of cards seems like a poor 'just make it bigger' engineering choice when they should have been trying to find more performance, lower costs, and achieve better efficiency. **–BJ Koho**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: Your letter makes some really good points and I understand the frustration from many people over the lack of a new Nvidia (or AMD) graphics card at the more traditional $500-700 price range. As Jeremy Laird wrote in his column in our December 2022 issue, the RTX 4090 is a beast that just about justifies its astronomical price, but the RTX 4080 is somehow worse value on a performance-to-dollar ratio.

As I write in my Lab Note this month (page 72), the RTX 4080 is clearly hobbled and priced in a way designed to push people (especially early adopters) to the more expensive RTX 4090—a classic marketing tactic most famously used by Apple and its "Pro" products. The fact that this strategy is being deployed during a cost of living crisis isn't a good look though, and I feel journalist sentiment towards Nvidia (us included) is starting to turn.

At the time of writing, I'm encouraged by the fact that it's relatively easy to get a hold of an RTX 4080 or 4090 partner board—the eBay scalping was fairly short-lived. Hopefully, that's a sign that maybe they're not quite as hot as Nvidia thinks they are and also that they're not worth the hassle of scalping. Who knows, maybe we'll even get a discount soon? Wishful thinking, I know!

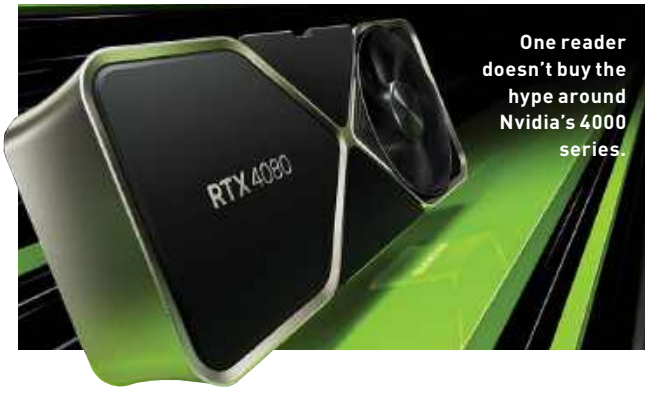

One reader doesn't buy the hype around Nvidia's 4000 series.

© NVIDIA

My hope is that the sub-$1,000 AMD RX 7900 XTX and RX 7900 XT will make the $1,200 RTX 4080 look even less enticing than it does now when it launches, and who knows, Intel might continue its hot streak with its Arc GPUs. For all our sakes, Nvidia needs some competition right now.

Multiple GPUs?

I'd like to build a PC dedicated to games, but keep my old monitor for office work, surfing the net, and anything else non-gaming. I want to put my new, powerful graphics card on the new screen and my old graphics card on my old screen. Is this possible? Will gaming performance be affected by using both at the same time? **–G. Bordier**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: You should be able to install two or more graphics cards in your PC without any issues, especially if they use the same graphics platform (e.g. Nvidia or AMD). However, I can't think of one good reason why you would want to do that unless you were using two of the same graphics cards in tandem, also known as Nvidia SLI or AMD Crossfire. Even so, this style of setup has become pretty rare over the last few years, to the point where GPU makers have phased it out, and the RTX 4000 series doesn't support it at all.

Using two graphics cards will use more power, take up more space in your case, and there could be driver conflicts or settings to change when playing games. Your new graphics card will have three or four video outputs, so you can run the work monitor and the gaming monitor simultaneously if you want to, and have options to add even more monitors in the future too.

THE BUILDS

THIS MONTH'S STREET PRICES...

BUDGET

IN OUR INTEL BUDGET MACHINE, we see a welcome price drop on both our SSD and GPU this month. This frees up some budget to play with, so we thought we'd switch things up and modernize this rig. The first thing we changed was the motherboard—we've stuck with ASRock but swapped from a B660 chipset to a more powerful Z690 Phantom Gaming 4/D5. Not only does this upgrade our chipset but we also now have DDR5 support. These DDR5 sticks have dipped in price since their release and are now more affordable, so it's time to introduce them to our budget machine.

We swapped out our Corsair DDR4 3200MHz RAM for 16GB (2x 8GB) of Kingston FURY Beast @ 4800MHz. We also changed our previous Intel Core i5-12600 for the K-variant, a faster and stronger chip than before, so updating from DDR4 to DDR5 means the memory speed can increase from 3,200MHz to 4,800MHz which the 12600K supports. In terms of changes, the rest of this machine remains the same and with the fluctuating prices from last month's budget machine, all these changes and updates ended up costing us $14 more than the previous build.

For the AMD machine, we've kept things pretty much as they were. At the time of writing, we're waiting for next-gen budget GPUs to be released—this may see a shift in prices that will allow us to upgrade this build. If we can free up some budget, we may look at bringing in the newest budget AMD chip, the Ryzen 5 7600X. As that doesn't come with a stock cooler, we would have to bring in a CPU cooler to accommodate it and also an AM5 motherboard. That said, we've upgraded our mobo to an X570C chipset with the ASRock Steel Legend—around $15 more than our previous B550M board. We also swapped the MSI Radeon RX 6650 XT MECH 2X OC for a Gigabyte Radeon RX 6650 XT EAGLE OC 8GB as the former was out of stock. Fortunately, there was no change in price as the Gigabyte is $260 on NewEgg.

AMD INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	550W Cooler Master MWE Gold 80+ Gold	$60
Mobo	ASRock X570 Steel Legend ATX AM4 Motherboard NEW	$134
CPU	AMD Ryzen 5 5600X	$160
GPU	Gigabyte Radeon RX 6650 XT EAGLE OC 8GB NEW	$260
RAM	16GB (2x 8GB) PNY XLR8 Gaming EPIC-X RGB @ 3200MHz	$60
SSD	500GB Corsair MP600 PCIe 4.0 M.2 SSD	$65
HDD	4TB Seagate SkyHawk Surveillance 7200	$88
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$954**

INTEL INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	550W Asus ROG Strix 80+ Gold	$70
Mobo	ASRock Z690 Phantom Gaming 4/D5 ATX LGA1700 NEW	$156
CPU	Intel Core i5-12600K NEW	$237
GPU	Zotac GeForce RTX 3060 GAMING TWIN EDGE OC 12GB	$350
RAM	16GB (2 x 8GB) Kingston FURY Beast DDR5 @ 4800MHz NEW	$89
SSD	500GB MSI Spatium M450 PCIe 4.0 M.2 SSD	$40
HDD	4TB Seagate SkyHawk Surveillance 7200	$88
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$1,157**

AS WE IDEALLY WANTED to shift the new AMD Ryzen 5 7600X from this build into our budget AMD machine, we've taken it out of our mid-range machine and introduced the AMD Ryzen 7 7700X instead. This chip costs $30 more than the Ryzen 5 7600X last issue. Thankfully, the price of the Ryzen 5 7600X chip is now around the $230 mark, suiting the budget build more. This 7700X is a faster processor and contains two more cores and four more threads than the budget-oriented machines. We also changed the motherboard for our mid-range machine, bringing in the ASRock B650E PG RIPTIDE WIFI ATX mobo. Previously, we had the Asus PRIME B650M-A AX Micro ATX board here, so we decided to try and kill two birds with one stone by upgrading our chipset and also improving compatibility and future-proofing by going from Micro-ATX to ATX.

The change of motherboard and processor makes no difference to compatibility with the rest of the components. Our DDR5 Corsair RAM still works just as well with the new chip and board too. At the time of writing, we're still holding on for the latest AMD GPUs to hit the shelves so our Radeon RX 6800 XT won't be leaving this machine. This OC variant from ASRock is still a powerful card capable of pushing high FPS at high settings too. Like the GPU, we keep the rest of the build the same.

As for the Intel mid-range machine, we featured the Intel i5-13600K in our cover build for this issue, where it proved capable of being able to sit alongside an RTX 4090 and produce some great benchmarking figures. That said, we can't put an RTX 4090 in the mid-range build now, can we? We'll have to pair it alongside Gigabyte's Gaming OC RTX 3070 Ti for now, as RTX 3080s are still too expensive to include in this section. The rest of our build contains a well-balanced bunch of components that don't need much further attention this time around.

AMD INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow	$130
PSU	750W Corsair RM750 80+ Gold	$80
Mobo	ASRock B650E PG RIPTIDE WIFI ATX AM5 NEW	$230
CPU	AMD Ryzen 7 7700X NEW	$330
Cooler	Cooler Master MASTERLIQUID ML240L RGB V2	$98
GPU	ASRock Radeon RX 6800 XT Phantom Gaming D OC 16GB	$515
RAM	32GB (2x 16GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	$135
SSD	1TB Crucial P5 Plus NVME M.2 PCIe 4.0	$94
HDD	4TB Seagate SkyHawk Surveillance 7200	$88
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $1,732

INTEL INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow	$130
PSU	750W Gigabyte P750GM 80+ Gold	$110
Mobo	MSI PRO Z690-A ATX LGA 1700	$190
CPU	Intel Core i5-13600K	$300
Cooler	Cooler Master MASTERLIQUID ML240L RGB V2	$98
GPU	Gigabyte GeForce RTX 3070 Ti GAMING OC 8GB	$610
RAM	32GB (2x 16GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	$135
SSD	1TB ADATA XPG GAMMIX S70 M.2 PCIe 4.0	$99
HDD	4TB Seagate SkyHawk Surveillance 7200	$88
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $1,792

WE ARE STILL waiting for the latest AMD GPUs to arrive, so will start with the Intel build. With an RTX 4090, it is hard not to get excited about this turbo build. We've been lucky enough to test two different RTX 4090 cards in separate builds and, boy, can this card perform. It isn't without its building challenges though.

Our PNY XLR8 GeForce RTX 4090 Gaming VERTO EPIC-X in this month's cover build needed four individual PCIe ports on the back of the PSU, while our CPU also required two. To pre-empt this issue in other RTX 4090s, we have selected a PSU with six available PCIe/CPU ports on the back and with at least 1,000W. Specifically, Corsair's RM1000 80+Gold PSU.

We also swapped out the MSI RTX 4090 for a Zotac Gaming Trinity OC RTX 4090 due to limited stock of the former. We will be swapping out a Z690 with a Gigabyte Z790 AORUS ELITE AX ATX LGA1700 mobo to keep that as up-to-date as possible.

On the AMD build, we spruced up the turbo machine last month, adding the AMD Ryzen 9 7900X. We also gave it the best AMD chipset with the ASRock X670E PG Lightning ATX motherboard—and both of these aren't going anywhere just yet. The Radeon RX 6950 XT GPU, the best we can include in this rig, is far cheaper than the RTX 4090 in the opposing camp. The ASRock OC Formula 6950 XT was also out of stock so we swapped that out for MSI's GAMING X TRIO Radeon RX 6950 XT.

We stuck with our trusty 6TB WD Blue 5400 HDD for the turbo and mid-range machines and are also using 1TB of SN850X on both turbo PCs—as luck would have it, a sale brought the price down to $100 from $135. Another component shared across these two builds is the RAM, and we've stuck with 64GB of Corsair's Vengeance DDR5 @ 5200MHz. With the various price drops and sales, the overall price of the AMD machine fell to $2,445 from $2,710, though it may rise when we add the new AMD GPUs to the mix. There's a huge difference in price between the two turbo machines, but that's down to the crazy RTX 4090 prices.

AMD INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$190
PSU	1000W Corsair RM1000 80+ Gold NEW	$168
Mobo	ASRock X670E PG Lightning ATX AM5	$260
CPU	AMD Ryzen 9 7900X	$440
Cooler	Cooler Master MasterLiquid ML360R RGB 66.7 CFM 360mm	$145
GPU	ASRock OC Formula Radeon RX 6950 XT 16GB	$780
RAM	64GB (2x 32GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	$250
SSD	1TB WD Black SN850X M.2 PCIe 4.0	$100
HDD	6TB WD Blue 5400 HDD	$80
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$2,445**

INTEL INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$190
PSU	1000W Corsair RM1000 80+ Gold NEW	$168
Mobo	Gigabyte Z790 AORUS ELITE AX DDR5 LGA 1700 NEW	$250
CPU	Intel Core i9-13900K	$660
Cooler	NZXT Kraken X73 73.11 CFM 360mm	$185
GPU	Zotac GAMING Trinity OC GeForce RTX 4090 24GB NEW	$2,160
RAM	64GB (2x 32GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	$250
SSD	1TB WD Black SN850X M.2 PCIe 4.0	$100
HDD	6TB WD Blue 5400 HDD	$80
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$4,075**

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US, Inc., 10th floor, 1100 13th Street, N.W Washington DC 20005, USA. Website: www.futureus.com. Periodicals postage paid at New York, NY, and at additional mailing offices. Newsstand distribution is handled by Curtis Circulation Company. Basic subscription rates: one year (13 issues) US: $30; Canada: US$30; Foreign: US$50. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R128220688). PMA #40612608. Subscriptions do not include newsstand specials. POSTMASTER: Send changes of address to *Maximum PC*, PO Box 2024, Langhorne, PA 19047, USA. Future US, Inc. also publishes *MacLife*, and *PC Gamer*.

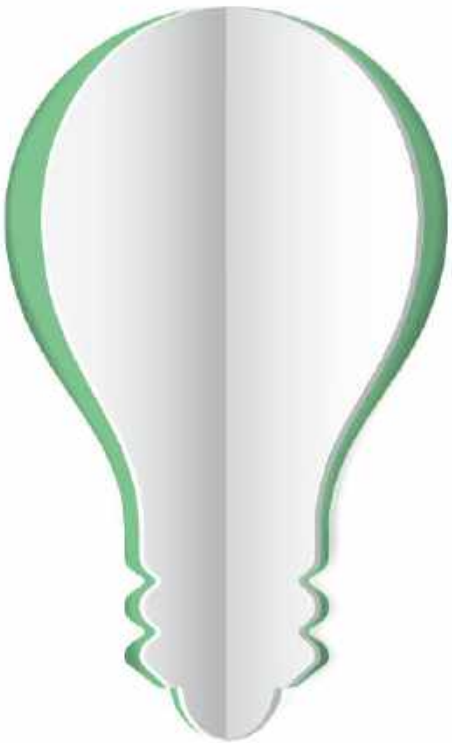

PAPER POWER

More than half of the energy demand at U.S. pulp, paper and paper-based packaging mills is met using renewable, carbon-neutral biomass energy.

Source: American Forest & Paper Association, 2022

Discover the story of paper
www.lovepaperna.org
Scan for paper facts, activities, blogs and much more!

LOVE PAPER®

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside