LIQUID-COOL YOUR PC PG. 32

MASTER WINDOWS **COPILOT AI** PG. 42

RTX 4080 SUPER REVIEWED PG. 74

MAXIMUM PC

MINIMUM BS • MAY 2024 • www.maximumpc.com

BUILT & TESTED STEP-BY-STEP GUIDE PG. 16

BUILD A NEXT-GEN PC

No Holds Barred 4K Gaming

- RTX 4080 Super + i9-14900K
- PCIe 5.0 SSD + Overclocked RAM
- Dynamic chassis cooling

GAME FASTER IN WINDOWS PG.60

PHOTOSHOP Vs GIMP PG.92

FUTURE
VOL 29, NO 5

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

The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

EFF.ORG

ELECTRONIC FRONTIER FOUNDATION

Protecting Rights and Promoting Freedom on the Electronic Frontier

inside

MAY 2024

SCAN TO GET THE TOM'S HARDWARE WEEKLY NEWSLETTER

QUICKSTART

Wireless gaming keyboards offer a degree of flexibility and freedom that you just don't get with a wired keyboard.

R&D

LETTERS

© CAPCOM, ASUS

IN THE LAB

AMAZING
SUBSCRIPTION OFFERS

SUBSCRIBE TO ANY OF OUR **BEST-SELLING** MAGAZINES

SHOP THE FULL RANGE

WWW.MAGAZINESDIRECT.COM/**SUBSCRIBE**

Meet the technology experts

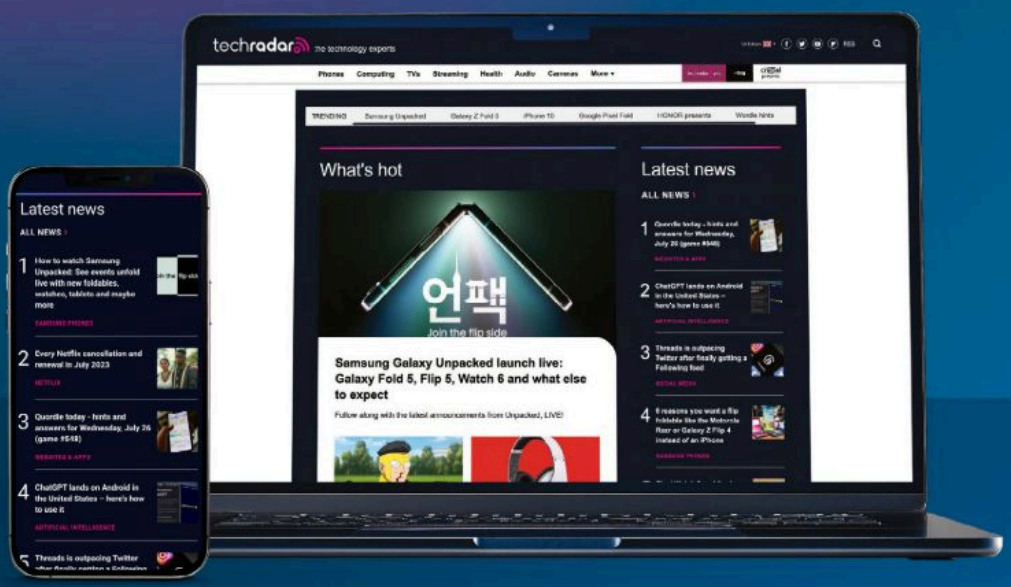

- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

techradar
the technology experts

www.techradar.com

MAXIMUMPC

EDITORIAL
Editor-in-Chief: Guy Cocker
Contributing Writers: Ian Betteridge, Fraser Brown, Tim Danton, Nate Drake, Ian Evenden, Dave James, Jeremy Laird, Chris Lloyd, Aleksha McLoughlin, Nick Peers, Nik Rawlinson, Zak Storey, Adam Timberley
Production Editor: Steve Wright
Editor Emeritus: Andrew Sanchez

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

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
CEO: Jon Steinberg
MD Tech: Paul Newman
Group Editor-in-Chief: Graham Barlow
Group Art Director: Warren Brown

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

Future US LLC, 130 West 42nd Street, 7th Floor, New York. NY 10036. 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 May 21, 2024

EXTRA DIGITAL FEATURES

AUDIO FILE | PHOTO GALLERY | VIDEO FILE

Guy Cocker

SUPER HERO FATIGUE

THIS ISSUE, we wrap up the RTX 40-series Super voyage we've been on for the last few months, as we've now built machines featuring the enhanced versions of the RTX 4070, 4070 Ti, and now 4080 cards. Each card has something interesting, whether it's a price cut or a performance boost, but the 4080 Super is the least easy to pigeonhole in terms of appeal. On paper, it should be an easy recommendation—the Super has all the power of the standard 4080, but with a $200 price cut, making it the $999 card it should have been from the start. In reality, it's difficult to find it in stock, and those who have it are charging $1,300. Meanwhile, the original 4080 is equally hard to find.

It's a shame, because as you can read in Zak's review on page 74, and his build on page 16, a $999 4080 Super would make sense for the dedicated gamer who isn't quite ready to spend $1,599 on an RTX 4090. Given the GPU price situation, Zak has decided to go all out on his other components in this issue's build—or "balls to the wall", as he calls it in his feature. We're talking dual PCIe 5.0 SSDs, overclockable DDR5 RAM and premium RGB fans, as well as the powerful Intel Core i9-14900K. It's a system that both looks and acts the part, and pushes that 4080 Super as far as it can go. You can find Zak's verdict in our cover build feature, and what's next from Nvidia now they've unveiled their new Blackwell architecture. Given the success the company is having in the AI market, I just hope they continue to make consumer graphics cards, to be quite honest.

Speaking of AI, we have a great feature on page 42 looking at how best to take advantage of Microsoft's Copilot features across Windows and Office, whether you're a free, 365, or professional user. Personally, I'd been putting off going near the Copilot button, which has been making itself more prominent in Windows, just because it felt like a 'beta' product at first. As we find out in the feature, it still has its quirks, but there are some seriously impressive things that you can do with it, particularly in Word and Teams.

We have so much other great stuff this issue, including how to play retro games on your PC, and whether you should still be liquid-cooling your computer in 2024. There are a ton of great tutorials as well, including how to remove the bloatware preinstalled on your PC, how to optimize Windows for gaming, and even how to get into game-making using Nvidia RTX Remix.

I don't usually highlight our Blueprints section, but it definitely deserves looking over for anyone in the market for a new PC. That's because Zak has updated the builds with some pretty interesting changes, including two Intel Arc-based graphics cards, plus Intel 14th gen and AMD Ryzen 7000 series CPUs across the board.

Next month, we aim to bring things down considerably on our build budget to answer the question: do you still need a dedicated GPU in 2024? With AMD's 8700G promising 1080p gaming in *Valorant* and *Rocket League*, should you ditch the discrete graphics card? Or are you better off going with a cheaper CPU, and pairing it with a low-end GPU? I can't wait to see this answered by Zak, so look out for that next month.

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

Nvidia's AI Superchip

The company is no longer about gaming

ON THE FIRST DAY of its annual GPU technology conference, Nvidia CEO Jensen Huang showed off the company's next big thing: Blackwell. Nvidia is keen to point out that Blackwell is a platform, not a GPU, and will power the next generation of AI accelerators and 5000-series graphics cards. Basically, it is the basis for everything from modest graphics cards to monster data-center kit, designed to run the world-changing AI that seems destined to dominate things. It is named in honor of David Harold Blackwell, the mathematician who specialized in game theory, probability, and statistics.

Nvidia claims that Blackwell has "six revolutionary technologies" designed for AI training and real-time LLMs. At its heart, the Blackwell silicon is two of the largest possible physical chips unified into one 4nm monster GPU with 208 billion transistors. It has support for a second-generation transformer engine and a fifth-generation NVLink, enabling up to 1.8TB/s throughput per GPU. There's a dedicated RAS engine, which stands for reliability, availability, and serviceability. A secure AI system is designed to protect models and data. Finally, there's a dedicated decompression engine to accelerate database queries and data analytics. Nvidia threw around some numbers for how much improvement Blackwell will offer, starting at 2.5 times faster than Hopper, going up to six or seven. When the metrics switch to LLM AI, the numbers go through the roof to 25 times faster. A lot of effort has been put into shifting data about, too. The platform includes a new NVLink chip, which has 50 billion transistors. The plan was to get every GPU in a Blackwell system to 'talk' to every other GPU at full speed, making one giant GPU.

The initial Blackwell chips have arrived in the form of the GB200 Grace Blackwell Superchip. This is designed for multi-node, liquid-cooled rack mounting in systems destined to run the most intensive workloads, meaning a lot of AI. It has two Blackwell B200 Tensor cores and a Grace CPU, all connected by a 900GB/s interconnect. This 40 petaFLOP monster is billed by Nvidia as the world's most powerful chip. Put 72 of these together, and you have the first exaflop supercomputer that fits on a single rack. In context, the first machine that could manage an exaflop was switched on in May 2022, and has 74 racks. Admittedly, Blackwell can only manage the feat by running inference FP4 instructions rather than full FP64, but it is still a staggering show of power. Less frightening is the B100, a replacement for the H100, A100, and B200. These consist of a single Blackwell GPU.

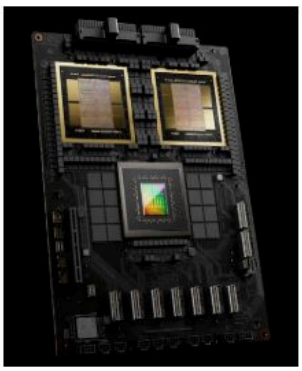

The GB200 Superchip, destined to power the world's most powerful data centers.

Interesting early projects for Blackwell include Gr00t, which stands for Generalist Robot 00 Technology (that's the story they are going with). This humanoid robot project has already produced some startling results. Each robot is powered by a single Blackwell GPU running multi-modal AI models (ones that combine inputs such as video, pressure sensors, audio, and so forth).

This is impressive stuff, especially if you're an Nvidia shareholder, but what about gaming? Well, we'll have to wait for Blackwell to filter down. Will it keep the two chips fused into one design? That might prove tricky to get to work in a gaming GPU, as well as being expensive, and possibly overkill. One half of a Blackwell GPU is 104 billion transistors—that's 25 percent more than the current RTX 4090, without factoring in efficiency gains.

The rumors have already started about Blackwell-powered GeForce cards. It's claimed that the RTX 5080 will be faster than the current RTX 4090. The RTX 5090 will be faster than anything else by some margin; it will be the biggest generational jump in performance in Nvidia's history. We can also expect a switch to GDDR7 memory, and that it will be designed to compete not just with AMD's RDNA 4, but RDNA 5, too.

Nvidia is not a gaming company anymore; it supplies the chips that will fuel the AI revolution. What gaming cards we get are increasingly less important to Nvidia. Blackwell cost billions to develop, but will make Nvidia many more over the next few years. In the process, the company is set to become one of the most important on the planet. **–CL**

> "This 40 petaFLOP monster is billed by Nvidia as the world's most powerful chip

WHAT IS AN AI PC?

WE'VE HEARD THE BUZZWORD, so what does it mean? Intel has defined the 'AI PC' for us, via the company's Robert Hallock. He claims that it consists of four things: a neural processing unit (NPU), a GPU, the capability to run Vector Network Instructions, and DP4a instructions. All are found in Intel's newest chips, naturally. Thus, it's the ability to run AI functions with some degree of aplomb. Intel went one step further at a recent AI conference, claiming that it also means a NPU with a processing power of 40 TOPS—trillions of instructions per second. This level of power is a little beyond what's available now, but many of the next-gen CPUs will count. Microsoft says it's also a system running Copilot with a dedicated Copilot key, plus 16GB of memory. All this means that just about every decent next-gen PC will have the hardware, which would make them all AI PCs. This is why Intel doesn't plan to brand anything specially. In a few years, everything will count, including your phone, rendering it a redundant epithet altogether. **–CL**

AMD UPGRADES FSR

AMD'S UPSCALING TECH, FidelityFX Super Resolution, has reached version 3.1, and is now in the hands of the developers. First to use it is *Ratchet and Clank: Rift Apart*, with an update in July. The big change is the decoupling of FSR from FMF (Fluid Motion Frames). You can play without upscaling and keep frame generation, or use Nvidia or Intel's upscaling tech with FMF, even on non-AMD cards. We're also promised less shimmer around objects, reduced ghosting, and better detail preservation. The next step for FSR is rumored to be switching to an AI-powered system. AMD is alone in not using AI hardware here. This means that FSR works on older cards, but leaves a lot of power on the table with newer silicon. AMD's chief technical officer said AI was AMD's main focus for 2024. Some AI upscaling is expected with RDNA 4, also due this year. Microsoft is working on new upscaling tech called DirectX Super Resolution, which won't care what hardware you're packing. Intel, Nvidia, and AMD may go all-in on hardware AI systems, leaving software upscaling to DirectX. **–CL**

Tech Triumphs and Tragedies

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

TRIUMPHS

QUANTUM APPLICATIONS
Google has launched a competition with $5 million in prizes to find uses for quantum computers.

10 PETABYTE DRIVES
Huawei claims to have invented a magneto-electric disk 0PB drive, drawing 2KW.

STABBING A VIRUS
Researchers in Australia have made an experimental surface covered in silicon nano spikes that physically disable viruses.

TRAGEDIES

BLAME THE USER
Microsoft claims Copilot criticism is because people don't know how to write prompts to full effect.

RIP WINDOWS ANDROID
Support for installed apps has a year to live—the rest of it is dead.

SPY IN YOUR CAR
Got an app for your vehicle? The data it collects is for sale, with insurances firms among the interested parties.

BIG TECH IN TROUBLE IN EUROPE

The EU isn't happy with the big players

THE EUROPEAN UNION has a piece of legislation called the Digital Markets Act, aiming to address the practices of large online companies, the so-called 'gatekeepers'. If you qualify as one, you must comply with a set of rules and obligations. You can guess who qualifies: Alphabet, Apple, Google, Amazon, Meta, and ByteDance, the Chinese firm behind TikTok. Five instances of non-compliance form the opening salvo. Apple and Alphabet are claimed to not allow apps to freely communicate, steering people away from alternatives. Apple isn't giving enough choice of apps, specifically a browser choice. Meta is making people pay to avoid ads, and Google is pushing its own goods and services over rivals. The punishment for non-compliance is a fine of up to ten per cent of the company's turnover—this can double for repeat offenders.

Previous concessions to EU rulings haven't always been generous. Apple drew fire for the way it allowed alternative app stores, but only with a pricing structure that discourages any development. Meta's European subscription model gave a binary choice: pay a fee, or we collect your data. At least it offered to reduce the fee from about $10.80 to $6 a month. Holding big tech to account is daunting, but the EU looks to be serious about this, which has implications elsewhere. The subjects of this scrutiny have started to protest in a series of public statements. Apple says it is "confident our plan complies with the DMA". Meta says it uses a "well-established business model". Google claims to have already made "significant changes to the way our services operate". However, there is a growing distrust with the gatekeepers, and the EU seems more determined than ever to bring them into line with the law. **–CL**

DOJ GOES FOR APPLE

This time it means business

APPLE IS WORTH $2.5 trillion, and has weathered many challenges, but the latest legal assault looks to have weight behind it. The Department of Justice, working with 16 state and district attorney generals, has filed an anti-trust case that claims it uses its dominance of the smartphone market to control it. The case cites section two of the Sherman anti-trust act, the federal law that limits the power of monopolies.

Among the areas of (alleged) abuse are using the app review process to stifle potential rivals, making it difficult to connect non-Apple smart watches to an iPhone, blocking access to Apple Pay, blocking 'super apps', and using green bubbles to stigmatize messages that aren't from another iPhone. Essentially, it calls out the ways Apple keeps others out of the lucrative ecosystem of the iPhone, or locks them into it. In the first three months of this year, 72 percent of new smartphones sold in the US were iPhones. Apple will defend itself vigorously, of course. It has enough spare cash to pay for a huge legal team and keep it employed for years, stretching out every aspect of the case. Behind that warm, friendly Apple store is a ferocious defender of territory.

The DOJ isn't the only legal authority taking a close look at Apple (the European Union isn't happy, either—see page 9), but it might have the clout to get it to actually change. Fines haven't worked in the past; Apple took one for $2 billion from the EU over music apps without blinking. It'll take actual legally enforceable changes in policy and practice to open out the world of Apple to competitors. If you have a long memory, you may recall the DOJ taking on Microsoft in the '90s, claiming it used the dominance of Windows to stifle competitors' web browsers. It took a long time to change things, but it did eventually. If it's not a warning from history for Apple, it should at least be a concern. **–CL**

Steam Updates Sharing

VALVE HAS COMBINED its Steam Family View and Family Sharing features into Steam Families. It lets you add up to five others to your Steam Family who have access to 'shareable' titles in your game library, and you to theirs. Each member gets their own saved games and set of Steam achievements, though not all games can be shared due to 'technical or other reasons'. The biggest issue with the previous system has been fixed—if somebody was playing a game from your library, nothing else could be played from it. Now, all members can play games from any member's library at once. There is also a new set of parental controls, including limiting playtime. To prevent hopping between multiple families, there is a one-year wait before you can join another family, or your place be taken by another. There is an exception if you leave and are invited back into the same family—then there's no wait.

However, the owner of a game is still held responsible for the conduct of the other members. If one gets banned from a game for whatever reason, you get banned, too. This could be harsh, particularly when you get thrown out of your expensive new purchase because somebody else says something thoughtless in chat. Other members can still play the game, which seems odd, but the culprit and owner are out of luck. To be fair, the previous system had a similar problem. Valve says Steam Families is designed for close family members, and will monitor and adjust requirements to keep usage in line with intent. It is currently in beta, so you have to opt in for now. Just be careful who you let on board, even if they're siblings. **–CL**

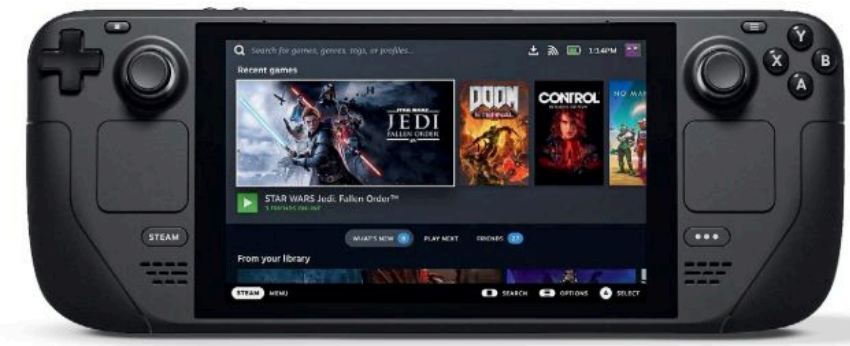

Happy 35th Web Anniversary

Tim Berners-Lee had some thoughts on the internet at this year's 35th anniversary, and it wasn't all complementary. He bemoans that it has become the playground of big business using "exploitative business models" based around data collection. It is bedeviled by poor leadership, anemic regulation, and profit-driven agendas. It also seems poised to create political turmoil. It's not all bad, though, and he points out some highlights, but it could be so much better, starting with taking back control of your data: something his Solid (Social Linked Data) project aims to do. His conclusion is that the web is too important to be left to the whims of big corporations. **–CL**

Windows Updates and BSOD

No Windows update goes unpunished. The latest of these is KB5035853, which has been the root of a number of snafus. There have also been a raft of complaints citing sluggish performance, stuttering audio, boot times, Explorer being unresponsive during copying, and a transparent Taskbar. The initial fix is the usual one: remove it. Microsoft released an optional patch on March 29 that is said to fix most performance issues, particularly with AMD systems (KB5035942), but other issues are yet to be sorted. Testing for all configurations is a daunting task, but Microsoft isn't winning any friends when an automatic update can render your whole rig inoperable. **–CL**

Jarred Walton

TECH TALK

Nvidia drops an AI bombshell with Blackwell

IT'S BEEN TWO YEARS since Nvidia revealed its Hopper H100 GPU architecture, currently one of the most sought-after processors for AI workloads. In fact, it's in such high demand that individual H100 accelerators can cost $30,000–$40,000. It has also been banned from export to China.

The mid-cycle refresh H200 has just started shipping. This means it's time for Nvidia to reveal its post-Hopper data center GPU plans.

Meet Blackwell. As we mentioned on page 8, the work of its namesake, David Blackwell, has had an impact on the research and development of artificial intelligence. That's fitting, considering the new GB200 GPU is set to power the next generation of massive AI supercomputers. Nvidia hasn't spilled all the beans, so we don't know the die size or number of processing units. However, we know that Blackwell has 208 billion transistors, and will be built on TSMC's N4P 4nm node.

We say 'combined', because the Maxwell GPU is composed of two die, linked together via a new Nvidia High Bandwidth Interface (NV-HBI). The maximum die size of a chip is around 858 mm2, but anything above 800 mm2 is effectively at the reticle size limit. Nvidia's Ampere GA100 chip was 826 mm2, made on TSMC's N7 node. The Hopper H100 is an 814 mm2 chip fabricated on TSMC's N4 node. TSMC N4P won't allow for substantially more transistors in a given area, so Nvidia's solution is to bind two chips together. The cost per Blackwell GB200 GPU is more than twice that of Hopper H100.

Each Blackwell die has four HBM3e 12-Hi stacks of memory—12GB each with 1 TB/s of bandwidth. That's two fewer HBM stacks per die than Hopper, which allows for more die area to focus on improving the compute. That's still 192GB of total memory and 8TB/s of bandwidth—over double the memory capacity and bandwidth of the highest-performance H100 solution.

> Blackwell has 208 billion transistors, and will be built on TSMC's N4P 4nm node

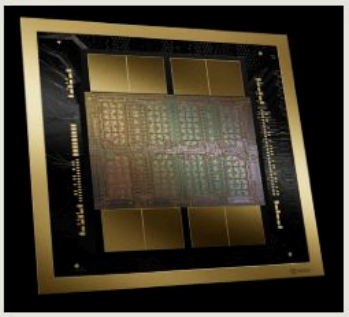

This comprises two full-reticle sized dies linked together via a 10 TB/s NV-HBI interface.

Nvidia also adds support for new FP4 and FP6 number formats, with its upgrade Transformer Engine helping developers leverage new formats. These will be mainly for inference workloads, and each GB200 GPU can provide up to 20 petaflops of FP4 compute. If you want to compare formats, Blackwell will provide up to 5 petaflops of FP16/BF16—2.5 times faster than Hopper.

Blackwell has also overhauled GPU-to-GPU connectivity and NVLink. Each Blackwell GPU has 18 NVLink connections, providing 100 GB/s of bandwidth. In aggregate, that's 1.8 TB/s of bandwidth for GPU communications—18 times more than on Hopper.

Nvidia's new Grace Blackwell 'superchip' will house two B200 GPUs with a Grace CPU. The main building block for Blackwell installations will be GB200 NVL72, which has 18 GB200 1U liquid-cooled trays, each with two GB200 Superchips and two Grace CPUs. Each tray can be configured for up to 2700W of power draw.

GB200 NVL72 will also have nine liquid-cooled NVLink 1U trays to link the GPUs together, with 14.4 TB/s of NVLink bandwidth per tray. That's 130 TB/s of aggregate NVLink bandwidth. Each NVLink v5 chip is 50 billion transistors on its own—nearly as many transistors as the old GA100 architecture.

Nvidia claims Blackwell GB200 will be up to four times faster than H100 on AI training workloads, and up to 30 times faster on inference workloads. Pricing and availability have not been announced, but we suspect a GB200 NVL72 rack could cost $5–$10 million. With eight of those, you get the Blackwell GB200 SuperPOD, with 576 GPUs, 288 CPUs, 240TB of HBM3e memory, and 11.5 exaflops of FP4 compute.

It's a tour de force in the AI space, set to enable the training and use of AI data sets with trillions of parameters.

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

THE LIST

THE BEST WIRELESS KEYBOARDS

WIRELESS GAMING KEYBOARDS offer a degree of flexibility and freedom that you just don't get with a regular wired keyboard. Freeing your keyboard from the chains of its cable means that you'll have the option to use it wherever you like, from your lap to the sofa and beyond—within range, at least. There's plenty to consider, from that very range and battery life to weight, build quality, programmability, and keystroke feel. Here are our five favorites.

5 STEELSERIES APEX PRO TKL WIRELESS

We recommend this for its compact frame and competitive gaming edge. That's because this SteelSeries board comes with magnetic OmniPoint 2.0 switches, which allow the user to set the actuation point of each switch anywhere between 0.2–3.8mm via the SteelSeries software. If you have your keys set to 0.2mm actuation, the Apex Pro TKL will take some beating for speed, that's for sure. **$249, www.steelseries.com**

4 ASUS ROG AZOTH

As enthusiast keyboards go, the great typing experience and fantastic build quality means the Asus ROG Azoth is a tough act to beat. From the sleek lubed switches under the surface to the neat OLED screen with equalizer and multifunction tactile, everything on the Azoth makes you want it more. What Asus gets extremely right is the feeling of each key press, which comes down with a soft and satisfying clack. It's gorgeous to type, and more than responsive enough for gaming. **$209, www.asus.com**

3 LOGITECH G915

This wireless keyboard has been around the Sun a few times, but it doesn't cease to stand out. It remains our pick for the best low-profile wireless gaming keyboard. A choice of switches, great battery life, and lag-free wireless connection make it a strong performer. The fact that it has extras like a volume wheel and dedicated media keys, plus per-key RGB lighting, sweeten the deal. **$186, www.logitech.com**

2 KEYCHRON K2 (VERSION 2)

Proving that great wireless keyboards don't need to cost the earth, the Keychron K2 is an entry-level mechanical keyboard that hits all the highlights. It redefines affordability in the wireless keyboard market, partly courtesy of 240 hours of battery life (provided you don't use the backlight). You don't even have to ditch mechanical switches. It offers both Wi-Fi and Bluetooth. **$69, www.keychron.com**

1 ASUS ROG STRIX SCOPE II 96 WIRELESS

This takes top spot thanks to its fully lubed switch feel, excellent build quality, and solid sound dampening. It's also got PBT keycaps, an adjustable media wheel, and hot-swappable switches, plus a numpad and programmable multimedia wheel. It uses Asus's Armoury Crate, supporting macros and RGB lighting customizations. What's not to love? **$159, www.asus.com**

Jeremy Laird

TRADE CHAT

Intel's crazy plan for 1nm silicon

FIVE NEW CHIP PRODUCTION NODES in four years. That was Intel's plan to get back to technology leadership. It has now added a new 1nm lithography node to its roadmap, due in 2027. But here's the thing: in terms of products to buy, Intel has achieved little of its original plan. So, what's going on?

First, we need to understand the true implications of Intel's existing plan for new chip nodes. In reality, it added up to far fewer than five nodes. Officially, the new nodes in question are Intel 7, Intel 4, Intel 3, Intel 20A, and Intel 18A.

Intel 7 is really just a rebrand and tweak of its existing and troubled 10nm node, Intel 3 is derived from Intel 4, and the same applies to Intel 18A in relation to Intel 20A. In which case, at the time of announcing that seemingly ambitious five-nodes-in-four-years roadmap, Intel was only committing to two fully new nodes: Intel 4 and Intel 20A.

But it was still pretty bold, given Intel's terrible recent track record with new silicon. Back in 2012, Intel was planning on unleashing 10nm CPUs as soon as 2015. In reality, it didn't launch a truly commercial 10nm product until September 2019 with Ice Lake. It was then over four years before it released a CPU on the next truly new node, known as Intel 4, when its latest Meteor Lake mobile chips stumbled onto the market at the end of 2023.

But Meteor Lake only contains a slither of Intel 4 silicon. Most of the chiplets that make up a Meteor Lake CPU are produced by TSMC, not Intel. Only the compute tile is an Intel 4 chip, so you could say that all Intel has achieved since CEO Pat Gelsinger took over the company and rolled out the new plan, at least in terms of chips to buy, is a tiny volume of Meteor Lake CPU dies on that Intel 4 node.

Sure, there have been other CPU launches, including Alder Lake and Raptor Lake, but those

> Here we are in 2024, and Intel has until the end of next year to deliver

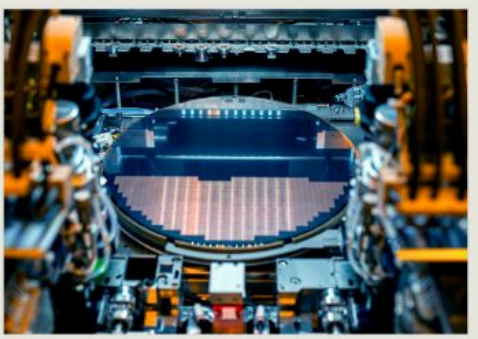

The Intel 7 node is really just a rebrand of its existing 10nm

didn't get Intel any closer to delivering on that roadmap, based as they are on older 10nm tech. Here we are in 2024, and Intel has until the end of next year to deliver.

The picture gets even weirder when you consider Intel's most recently revealed plans for its chip fabs. At the Intel Foundry Direct Connect last month, Intel showed a graph mapping out its planned chip production capacity up to 2029, and it was a little bit shocking.

Intel's capacity to produce chips is shrinking over 2023 and 2024, and won't again exceed 2023 levels until 2027. The same graph shows that in 2025, chips produced on the latest Intel 4 and 3 nodes (which it brackets together for this data) will be a minority of overall output and, oddly, will have been slightly overtaken by the then-brand new 20A/18A node (again, bracketed together because, really, they are two versions of the same node).

Capacity for both the Intel 4/3 and 20A/18A nodes will grow slightly in 2026, but it's not until 2027 that Intel expects those new nodes to replace its current mainstay of 10nm technology as the majority of its output.

Put another way, that bold plan back in 2021 now adds up to two nodes in four years, neither of them in large volumes for six years. In that context, Intel's announcement at the same event that the company is plotting a new '10A' node, which is equivalent to 1nm, for 2027 takes rather different implications.

Intel might be trickling out a few 1nm chips in 2027, but if its planned roll-out of the Intel 4 node is anything to go by, it won't be making 1nm at scale until 2031—and that's a best-case scenario. The bottom line is that Intel is at one and the same time nearing the end of its original plan, and yet still has almost everything to prove.

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

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> StopCrypt fears
> Tiny server build
> Backup bandwidth

Keep ransomware at bay

I've just read up on a new piece of ransomware that can evade detection and is aimed at consumer users. How can we protect ourselves from it?

—Samuel B Marshall

THE DOCTOR RESPONDS:

The ransomware you talk of is actually a variant of an existing form: StopCrypt. It has cleverly evolved to take a softly, softly approach to bypassing security measures, employing multiple delayed stages of execution to sneak on to a system before infecting it. StopCrypt doesn't generate the same level of media attention (surprise, surprise) because it doesn't involve huge sums of money being extorted from businesses. Instead, it focuses on us mere mortals, looking to extort hundreds of dollars in return for providing the key to decrypt your files and return them to a usable state.

While it's worrying how ransomware has evolved, the chances of you catching it are slim unless you break all the rules of good security. For example, StopCrypt is commonly found in malvertising and dodgy downloads advertised as free and cracked software. If you're partial to such activities, then you'll already be dicing with infections from all kinds of malware, not just StopCrypt. Even if your security software has kept you safe so far, it's not a smart way to operate, particularly if you do so on your main PC.

The obvious thing to do is steer clear of dodgy websites (or access them on a machine physically isolated from the your network). Second, make sure you have strong third-party anti-malware software, and run regular scans with another security package—the Premium version of Malwarebytes Anti-Malware is hot on blocking dangerous websites, and can be installed alongside existing security software. Failing that, run weekly scans using Malwarebytes Free, and consider adding the Emsisoft Emergency Kit (www.emsisoft.com/en/home/emergency-kit), a free portable anti-malware scanner that combines two antivirus engines—its own and Bitdefender's—for a second (and third) opinion.

We also recommend finding ways to protect yourself against malvertising—look for ad blockers like Disconnect (https://disconnect.me) and AdBlock Plus (https://adblockplus.org), which can strip out a lot of unwanted advertising, including ads containing hidden malware. If you want a universal solution to protect your network, read the January 2024 issue for our tutorial on setting up Pi-hole.

Ultimately, you also need to protect the target of any ransomware: your personal data. One obvious remedy is to make sure it's regularly backed up—and in a location that ransomware can't easily reach, such as offsite using a cloud backup tool like OneDrive. If that location is a network share, make sure it's password protected, and that you've not checked 'Remember my credentials' when logging on in the past. If you have, open Credential Manager and select Windows Credentials to view and remove any stored passwords. This prevents ransomware from being able to use these credentials to connect to remote shares and encrypt them, too.

Having to input your username and password each time is fiddly (although some backup software, including Macrium Reflect, can store your credentials within the program itself), but it might one day save you hundreds of dollars in extortion fees.

Micro server build

I enjoyed your latest NAS server build, but realize that I have no need for all the extra storage you've packed into the case. What I'd love to build is something akin to a tiny NUC PC, but using the components you did. I'm hoping I can do this with the AS Rock N100DC-ITX. Do you concur? **—Daniel J Powell**

THE DOCTOR RESPONDS:

You've got it in one, Daniel. Because the N100DC-ITX comes with a built-in PSU that's powered from a standard laptop connector, you can house it in a similar style of case to those used by NUC PCs. The trick is not to look for a 'mini ITX case', because that'll send you down the path of towers and NAS-style cases, like the Jonsbo N2 we featured in last month's build. Instead, you need to think along the lines of HTPC.

↘ submit your questions to: doctor@maximumpc.com

One case worth checking out is MITXPC's MX500-USB3 (https://mitxpc.com/products/mx500-usb3). It's got a tiny footprint (228.4 x 194.4 x 61.8 mm) and is designed for mobos with laptop power connectors, like the N100DC-ITX. Despite its tiny footprint, there's space for two 2.5-inch drives (plus the board has a single NVMe slot), so you can still max out its limited storage potential at some point if your needs change. There's also room for up to four 40mm case fans, although if the ventilation is as good as the Jonsbo N2, the CPU won't max out above 50C/120F, even when taxed. MITXPC sells the MX500-USB3 direct for $49.95 plus shipping, or you can buy it from Amazon (www.amazon.com/dp/B0728DX73X) for $59.95.

Although it doesn't offer any space for a PSU, MITXPC offers a choice of optional PICOPSU power supply boards (80W-120W) for pairing with an ATX-powered mobo like the ASUS Prime N100I-D4. Prices start from around $29.95 for the PICOPSU board (see https://mitxpc.com/collections/power-supplies)—you can either supply your own laptop power block, or purchase through MITXPC, which charges $94.95 for the PICOPSU board plus 120W laptop power block.

Throttle Duplicati backups

I've installed Duplicati (May 2023 issue), but have run into a problem with backing up to OneDrive. As soon as the backup starts, my internet connection grinds to a halt. I've tried setting a universal limit using the throttle option button, but it has no effect (ironically, it slows down all my local backups). A quick search online mentions some advanced settings, but I don't know which ones to try—and even if they work or not. Can you advise?

—Justin Cone

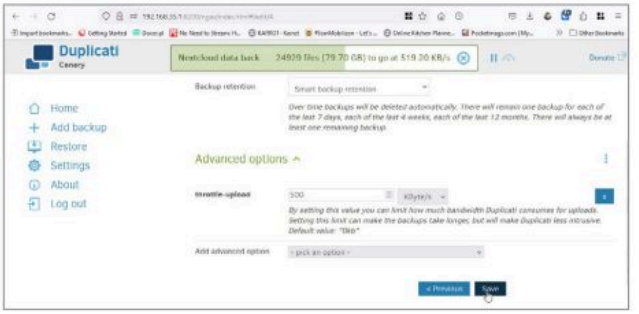

Throttle individual backup jobs to protect your bandwidth.

THE DOCTOR RESPONDS: It can be a pain configuring Duplicati to work with cloud backups without it saturating your bandwidth, particularly on slower connections. Things are doubly complicated when there are multiple places where settings can be configured. We recommend ignoring Duplicati's main Settings screen to focus on applying limits to individual backup jobs where possible.

To do this, first stop any cloud backups currently running, then expand the backup job on Duplicati's home screen, and click Edit under Configuration. Click Next to jump to the Destination screen, where your online storage should already be configured. Expand the 'Advanced options' and click the '- pick an option -' dropdown menu. Select 'fragment-size: Fragment size for large uploads'. You'll see that it's set to 10MB by default—try a lower setting (say 5MB) and click Next, skipping through the next few screens to land on the final Options page.

Here you'll find yet another expandable Advanced options section. Clicking '- pick an option -' here reveals an even bigger list of configurable options to choose from. Look for the 'throttle-upload' option and set it to a suitable level—say, 400 KByte/s—and click Save.

Once done, click 'Run now' next to the backup job, and monitor its upload speeds through the Duplicati interface. It should now respect your choice. Next, run a bandwidth test using www.speedtest.net—if there's still a noticeable lag, then edit the settings again, choosing a lower figure. On the other hand, if performance is fine, try higher figures until you strike the balance between getting your backups done in good time while maintaining a usable internet connection. Don't forget to stop and restart the backup each time to apply any changes.

Upgrade key problem

I've purchased a Windows 11 Pro key to upgrade my Windows 11 Home machine. I've followed the instructions, but when I enter the key, I get a 'Windows upgrade failed' message, with an error code of 0x80070490. Can you assist? **—William Headrick**

THE DOCTOR RESPONDS: This is a common issue, particularly if you're looking to upgrade without nuking your current installation and starting from scratch. The problem has been linked to an old version of the .NET Framework lurking on PCs—usually installed by a program that required it. It's version 3.5 or lower, and you can resolve the issue and upgrade by opening the 'Turn Windows features on or off' Control Panel (search for 'Windows Features'). You should see a '.NET Framework 3.5 (includes .NET 2.0 and 3.0)' entry—clear the checkbox and click OK.

This should clear the error and allow you to upgrade to Windows 11 Pro; if it still doesn't work for any reason, download the Windows 11 Media Creation tool (https://go.microsoft.com/fwlink/?linkid=2156295) and use that to perform a repair install. It should clear any incompatible software, and once done, you'll be able to upgrade successfully via the System > Activation section of Settings.

Can't remove USB

Whenever I attempt to eject my USB thumb drive, Windows refuses to let me safely do so, claiming something is using it. What can I do to resolve this without having to either take a risk each time or shut down my PC to unplug the drive?

—Karen R Gillespie

THE DOCTOR RESPONDS: If you're unable to easily identify the app or process responsible for locking the drive, you can uncover it using Event Viewer. Type 'event' into the Search box, and launch Event Viewer from the results shown. Expand the Windows Logs section on the left and select System from inside it. Next, click 'Filter Current Log...' in the Actions pane on the right, type 225 into the <All Event IDs> field, and click OK. All events with ID 225 refer to failed attempts to eject or remove a USB device. Simply examine each one, and you should be able to identify the offending app or process.

From here, you have a couple of choices—if there's a reference to a specific file on the drive that's been locked, try a free tool like Lockhunter (https://lockhunter.com), which should be able to free up the file. Failing that, if a specific app is referenced, simply close all instances of it (check Task Manager if necessary). Finally, if it's a process, it's probably best to log off or put your PC to sleep temporarily. You can then remove the drive safely without having to shut down and restart.

BALLS TO THE WALL

THE ULTIMATE GAMING PC

Strap in, as our **RTX 4080 Super** build has finally landed, says *Zak Storey*

It can be argued that Nvidia's 40 series Super cards have been incredibly disruptive since their launch. We've focused these last two issues on the 70 and 70 Ti Super, and the value that they represent, pairing them with suitable processors, and apt systems built around ensuring they provide the best possible value out of the lot.

Interestingly, the RTX 4080 Super, although initially launching as a big value option with a $200 RRP drop, has become expensive, and because of that, we decided to push the limit and see exactly what we can achieve if we threw everything into a build with it in. We picked up the Asus ROG Strix Gaming OC variant, with its massive triple-fan white cooler, and set to work building a spec list that would pair well with the GPU.

This build features one of the latest and greatest Core i9s from Intel, 4TB of PCIe 5.0 storage, a 1200W PSU, a stunning Z790 motherboard from NZXT, and perhaps more interestingly, a full-tower, 90-degree shifted chassis, allowing us to take advantage of thermal convection with a chimney style design. This build is incredibly exciting, and a bit of a first for this journalist in terms of working with a chassis like this. So with that out of the way, let's dive in and find out exactly what makes this build tick.

INGREDIENTS

PART		PRICE
CPU	Intel Core i9-14900K	$540
Motherboard	NZXT N7 Z790 ATX	$297
CPU Cooler	Corsair iCUE Link H150i LCD Black	$255
RAM	32GB (2x16GB) Crucial Pro Overclocking DDR5 @ 6000 C36	$105
SSD 1	2TB Crucial T705 M.2 PCIe 5.0 x4	$337
SSD 2	2TB Crucial T700 M.2 PCIe 5.0 x4	$295
GPU	Asus ROG Strix Gaming OC RTX 4080 Super	$1,300
Case	Geometric Future Model 8 ATX Mid-Tower	$150
PSU	1200W Corsair RM1200x Shift 80+ Gold	$190
120mm Fans	2x Corsair iCUE Link QX120 RGB Starter Kits	$280
TOTAL		**$3,749**

PRICES CORRECT AT THE TIME OF PUBLICATION

media://content.jwplatform.com/videos/qoG9J8fa-u2lN49He.mp4
Please type this URL into your browser if the link is broken

HARDWARE HAVOC

CPU

INTEL CORE i9-14900K

Intel's Core i9-14900K might be one of the hottest chips around (seriously, we've not actually seen this thing not run at 100C under full load), but it's one of the fastest. With eight performance cores, 16 threads, 16 efficient cores, and 36MB of smart-cache, this insane little number will quite happily crank itself up to 6 GHz, depending on workloads, and rip-roar through any task you throw at it.

It unequivocally dominates the gaming scene right now, and although it isn't exactly the best value option out there, if you're after the ultimate performance on a mainstream platform, then the 14900K is the chip to beat. ***$540, www.intel.com***

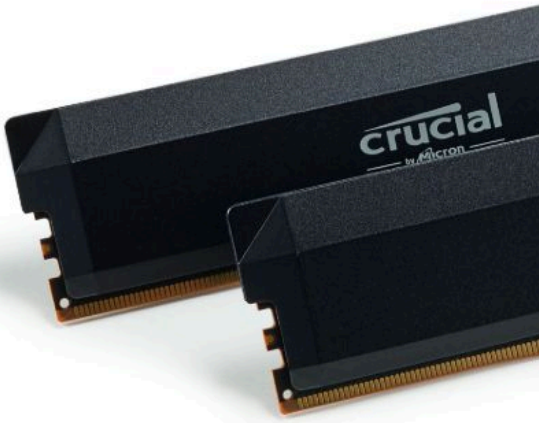

Motherboard

NZXT N7 Z790X

It might be on the 'cheaper' end of the Z790 spectrum, but this motherboard hits the right notes. With support for Intel 14th and 13th gen chips, it also packs some epic connectivity. You've got support for PCIe 5.0 graphics cards, along with two PCIe 4.0 M.2 SSDs, DDR5 support up to 7200 MHz, and some tasty rear I/O, including seven USB ports, a USB C, HDMI, 2.5G Ethernet, 5.1 Audio Out plus Optical audio, WiFi 6E, and a clear CMOS button. On top of that, it has a 16+1+2 power phase setup in a sleek, motherboard design, complete with full board brushed aluminum armor. You can also get the N7 in black or white to match your build. *$297, www.nzxt.com*

BUILT & TESTED
Step-By-Step Guide
PG. 22

RAM

32GB (2X16GB) CRUCIAL PRO OVERCLOCKING DDR5 @ 6000 C36

Our memory pick for this build actually came as an accidental shipment alongside the drives we're going to be using this time around, and that's Crucial's Pro Overclocking DDR5. This is an incredibly discreet kit of 32GB 6000 MHz memory, complete with C36 latency, giving it a 12 ns real-world latency. In other words, it's pretty slick, particularly at that price.

Similar to our last build, there are no frills—just pure, pristine, black heatsinks finished with a crisp logo. Intel chips don't massively benefit from faster memory speeds, but it's still a nice bonus to have, particularly for video rendering, where frequency and capacity work incredibly well together. *$109, www.crucial.com*

SSD 1

2TB CRUCIAL T705 M.2 PCIE 5.0 SSD

Our SSD of choice for our OS is Crucial's T705 M.2 PCIe 5.0 SSD. This blisteringly fast 5.0 drive clocks in with sequentials topping out the 5.0 bandwidth at a stunning 14.5 GB/s. It's not the cheapest drive out there, and at $337 it certainly does take its toll on the old bank account, but if massive file transfers are your thing, then the T705 is the perfect drive for you.
$337, www.crucial.com

SSD 2

2TB CRUCIAL T700 M.2 PCIE 5.0 SSD

Our second SSD is another PCIe 5.0 drive, but with a caveat. We initially picked this drive, knowing that our sequential performance would be capped at 7.5GB/s or so, as it would be sitting in a PCIe 4.0 M.2 slot, the major difference being that the random 4K performance (not hindered by bandwidth limitations) would be higher, and that is a major performance indicator for a game's load time. Both of these drives also future-proof us significantly once we jump up to Z890 and Intel's 15th gen in the future, or AMD's X770.
$295, www.crucial.com

ALWAYS CHECK THE FINE PRINT

We're never ashamed to admit when we're wrong, or when we make a mistake here at *Maximum PC*, and sadly, we were caught off guard in this build with a tiny detail that's somewhat hammered our SSD performance. By default, the NZXT N7 Z790 motherboard does 'technically' support PCIe 5.0 M.2 SSDs as a boot drive. It's listed in the advertising for it, and in the specs tables. However, when you read the finer print, it turns out that PCIe 5.0 drives are in fact only supported in the topmost PCIe slot, ie. not the M.2 slot, the x16 GPU slot... yeah. Less than helpful, that NZXT.

In our test results (and we'll go into a little more detail on this later), you'll see that maximum sequentials are around half what we'd expect, as unsurprisingly, they're running in a PCIe 4.0 board. We knew we'd be getting that with the T700, as it was chosen more for its Random 4K performance for gaming rather than sequential file transfers, but it's a well-earned reminder that you should always check the fine print.

Case

GEOMETRIC FUTURE MODEL 8 ATX MID-TOWER

A bit of a wild-card choice for the case this month. We've gone with the Geometric Future Dharma chassis. Launching in mid 2022, this is a highly intriguing case. The company, only founded in 2020, has come a long way in a very short amount of time.

What makes the Model 8 unique, however, is its internal layout. Featuring a very similar design to that of most modern towers, the internals have been rotated 90 degrees. What that means is that your motherboard I/O, GPU ports, and everything else are actually routed out of the top of the chassis itself. Admittedly, this is a bit of a hassle, but what it does do, with enough floor-mounted fans, is give us access to some incredible cooling potential.

In fact, we managed to cram a total of four 120mm fans in the floor, three in the 'rear' exhaust, and one in the roof. Couple that with some slick material usage and a generally impressive finish throughout, and it might be the most overlooked chassis of the last two years.

$150, www.geometricfuture.com

BUILT & TESTED Step-By-Step Guide PG. 22

120MM FANS

2X CORSAIR ICUE LINK QX120 RGB STARTER KITS

Last but by no means least are the fans for this build. For this one, it's probably fairly obvious that to match our iCUE Link cooler, we've paired it up with two sets of Corsair's QX120 RGB Starter kits. To be clear, you don't need to do this; it's just the easiest way for us to get the hardware in. We're utilizing only five of the extra fans total, plus a single iCUE Link hub. You could just use the single starter kit + two expansion kits, but this way gives you a ton of extra cables to help with the cable management situation, that's for sure.

$280, www.corsair.com

GPU

ASUS ROG STRIX GAMING OC RTX 4080 SUPER

Here is the showpiece of the build, the Asus ROG Strix Gaming OC, in all its magnificent glory. This is a chunky GPU, that's a fact—significantly larger than the standard Founders Edition, the Strix Gaming OC, aside from being white, fits a triple fan cooling solution, alongside some surprisingly tasty RGB as well.

As for the GPU itself? Well you can find our full review on page 74, and the benchmark results for this build on the same page, but what we will say is this thing absolutely demolishes frame rates at 4K. Whatever you think of the 4080 Super's current pricing scheme, there's no denying that outside of the ridiculousness of the 4090, this is the card to pick. *$1,300, www.asus.com*

PSU

1200W CORSAIR RM1200X SHIFT 80+ GOLD

We've gone with another tried and true PSU for this build, and that's Corsair's 1200W RMx Shift power supply. Delivering a huge amount of power on demand, and at a fairly decent efficiency, its modular cable ports have been rotated around to the side, giving us fantastic clearance, and making it a super-easy plug-and-play solution in the rear of our case. Corsair is utilizing a new design of cable to facilitate this: the Type 5, which effectively act as micro-connectors, too. It's fully ATX 3.0 compatible, and also comes with a direct PCIe to 12VHPWR600 cable, so there's no need for those pesky adapters/converters and extra points of failure. *$190, www.corsair.com*

CPU COOLER

CORSAIR ICUE LINK H150I LCD BLACK

We've got one heck of a chip to cool with that Intel Core i9-14900K, so if we can improve our odds and ramp up the cooling, we've got to commit to it. More cooling means higher clock speeds for longer and better performance overall.

To that end, we've gone with a classic cooler—our go-to for any and all test-bed scenarios currently, and that's Corsair's iCUE Link H150i LCD in black. Its 360mm bulk of a rad gives us a ton of surface area—combine that with the triple QX fans, and ease of setup, and it's easily the best pick for this build right now (also, after the 32-cable RGB setup, a sleeker daisy-chain setup definitely feels nicer to work with). *$255, www.corsair.com*

ROTATIONAL MADNESS

LENGTH OF TIME: 1-2 hours
DIFFICULTY: Easy

With the hardware selected (although mildly erroneous on the PCIe SSD side, but we'll leave that to the end), it's time to start building this beautiful thing, and trust us when we say that it's going to look like one heck of a machine when it's done.

The bulk of that is thanks to the incredible design of the Geometric Future Model 8 Dharma. Okay, it may have a funky name, but this thing really is phenomenal. From a purely design aesthetic, it's just good to see a manufacturer try to do something new and different. That goes for the majority of the company's cases, too. All of them are somewhat odd, weird, and outlandish, with warping designs and differing materials (one literally has a cloth front). They deserve credit for taking those risks.

The majority of the time, the vast swathe of cases out there from the big manufacturers are pretty much the same. You get the same internal layout, cable management, and power supply shroud—even the same colours and chassis designs. There's a reason the market is flooded with fishbowl cases and Lian Li PC O11 Dynamic imitators.

Capitalism drives competition and brings down prices—usually. It's the reason we won the Cold War, and a defining element of Western civilization; one that very much drives progress. But when it comes to larger companies with significant investment behind them, there typically comes a desire to have predictable returns and continual growth—say, 10 percent gains year on year, rather than erratic swings one way or another. That's where consistent and safe strategy is preferred over bold and dynamic plays and risk-taking.

What that tends to lead to is one smaller manufacturer taking a risk, trying to make a name for themselves, and designing something remarkable that usually challenges that status quo and becomes massively popular as a result, only for every other major manufacturer to then go out of their way to imitate that design, a year or so later, improving on it slightly, adding their own finer details, and pricing it accordingly.

That then ironically makes it the new status quo, once that initial risk has been taken by another manufacturer. We've seen it with PSU Shrouds, vertical GPU mounts, cable bar management, RGB fans, daisy-chain solutions—you name it. It's not exciting, but it is predictable. That's not necessarily a bad thing, either, as it adds refinement and develops the style and tech further. However, there's a risk that we're missing out on unique styles, designs, and features that might actually be better, simply because the big manufacturers that actually have the cash to invest in development simply don't want to take risks.

Seeing something like this, then, the Dharma, which incorporates modern design features, mixed with an internal layout that harkens back to a very rare style from the mid 2010s (which from a thermal design standpoint actually massively benefits modern-day GPUs), is refreshing. Is this going to be the new fishbowl chassis? Unlikely—Geometric doesn't quite have the marketing clout of those bigger brands, but nonetheless, it's

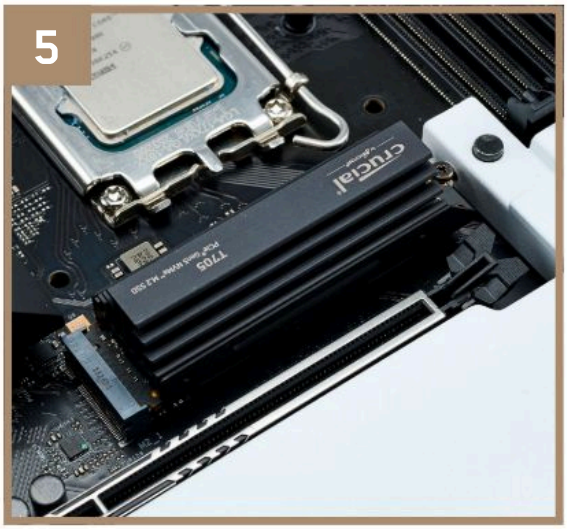

well worth considering. Anyway, that's enough of our philosophy behind good and evil on case design. Let's get down to the cold hard brass tacks of this build, shall we?

STRIPPED TO THE BONES

As with all good *Maximum PC* builds, our first and most important step up is the chassis strip-down [Image 1]. We do this with pretty much every build we undertake here at the mag, for a number of reasons. First and foremost, it helps protect the vast majority of those external panels from scratches and any potential damage. Also, by removing any and all componentry, panels, mesh filters, or things we're not going to use, it makes it far easier for us to work inside the case, and to finish up the cable management. It's particularly useful if you're looking at liquid-cooling PCs, too, as every inch of clearance counts in builds like that.

The Dharma is a fairly straightforward build—there's not a huge amount for us to remove. The topmost panel slides back and off, there's a dust filter underneath [Image 2] that you can remove just by sliding out, and the rear cable-management panel is held in place with thumbscrews—the same as the glass panel. You can remove the front panel, too. That's a solid sheet you can pop off once the two side panels are removed, and is held in place by pop rivets.

One side note here—the Model 8 has some serious cooling potential (there's space for eight 120mm fans, with a touch of modding), three or four in the floor, three in the side/rear, and one in the roof. It's worth noting, however, that only the floor has a dust filter in it. Clearly, Geometric expects you to use the floor as intake only, and the rear as an exhaust, otherwise you'll be pulling a not insignificant amount of unfiltered air straight through the rear of the chassis and into your graphics card. That's a bit of a shame, as it would be fantastic to really overload the positive pressure in this chassis and have seven 120mm fans drawing air in the bottom and rear of the case, with a single 120mm pushing it out of the roof.

MOTHERBOARD WOES

With the chassis stripped, we moved onto the motherboard prep [Image 3], this time remembering to nudge our photographer into ensuring the photo contained our anti-static workbench, ie. the NZXT's Z790 cardboard box.

We prepared the socket first, lifting up the retention arm on the LGA 1766 socket, and lifting the bracket up, in preparation for the CPU. After that, it's simply a case of carefully placing your CPU (in our build, the Intel Core i9-14900K) into position, ensuring you're aligning it correctly. There's a golden triangle on the corner of the chip, and on the motherboard socket to help with that, along with specific notches on the chip and the socket. Alternatively, you should always be able to read the branding on the CPU correctly if the top left of the motherboard (the very top left of the rear I/O, near the CPU power) is in a portrait orientation.

Place the CPU into the socket, give it a careful little wiggle with your finger to make sure it's snug and secure, and then lower the bracket back down, before re-securing the retention arm back into place. At this point, the plastic cover should pop off. If it doesn't, you can remove it by pulling it off the socket [Image 4] .

Going back to Image 3 quickly, you'll notice that we've also removed the topmost

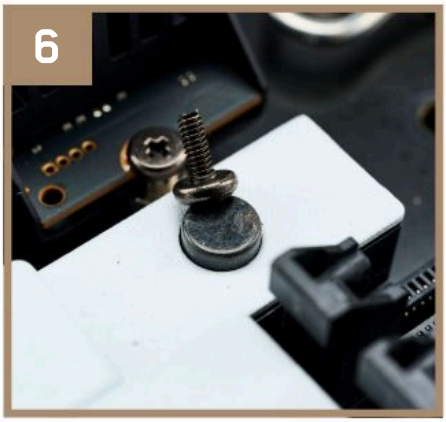

PCIe M.2 slot cover. These 'armor' plates are actually magnetized into position and come off fairly easily simply by pulling at them. That said, the magnets are pretty tough, so you shouldn't have trouble transporting the thing.

SSD 1 & 2

So, onto our contentious little drives. We've gone with the Crucial T705 as our primary OS drive, and given that it's the heatsink version, it presents us with a few problems. Unlike the T700, its heatsink isn't easily removable without prying it apart, and given that this is a $300 SSD at the height of its performance, that's not exactly something we're keen to do. That means we're going to have to remove the armor plate entirely to fit it [Image 5]. There are two small Philips screws holding it in place—simply remove those, and you can take out the bulk of the actual heatsink (this is a traditional M.2 heatsink, with a thermal pad underneath it).

With that done, it's simply a case of sliding your M.2 into position (taking care to pay attention to the notch orientation), and pushing it into place. Once in, you can secure it back down with one of the screws you used earlier. Little fun fact [Image 6]: those magnets are unfortunately right next to the M.2 slot, and since they're quite strong, as you're trying to push them into position to secure the M.2 drive into place, they have a bad habit of being pulled directly onto the magnet. It's a pain to deal with, but there's not a whole lot you can do about it other than persevere.

For our second PCIe 5.0 drive, the Crucial T700, we can actually disassemble this and hide it under one of the integrated heatsinks. It's still not something Crucial really wants you to do; you need a T5 torx screwdriver adapter, rather than a Philips head, to get the four screws holding it in place to come loose. We've used our Corsair kit for that, but any small precision toolkit will do the trick [Image 7], then you'll need to carefully pry the heatsink off the drive. It is stuck together with some rather curious-looking thermal pads that do cool some of the componentry underneath, so bear that in mind when you place it in your machine, but otherwise it's a simple disassembly.

With the drive now free, we move onto the motherboard, utilizing the M.2 slot below the GPU. There are no additional securing screws here—simply lift the magnetic armor plate up, place the drive into its M.2, reattach the armor plate back down, and you're done [Image 8].

DDR5 CONFUSION

The sheer dread we had when initially prepping this board was, well, impressive. At this point we had both our SSDs installed and had gone ahead and installed our CPU Cooler standoffs as well (four screws and the backplate does the trick nicely, which is easy enough with the iCUE Link kit). If you do need to swap to AMD, there are separate stand-offs that fit into the stock AMD motherboard backplate, and a replacement bracket that attaches to the CPU block itself, only to then move onto memory.

On every DDR5 motherboard we've had, the DDR5 slots have been in a particular orientation. That means memory manufacturers typically design their RAM so that the timings, voltages, and clock-speed labels face back towards the CPU (so the end-user can't see it), and the big brand logo faces forwards out of the case (so you can see that instead), which makes sense.

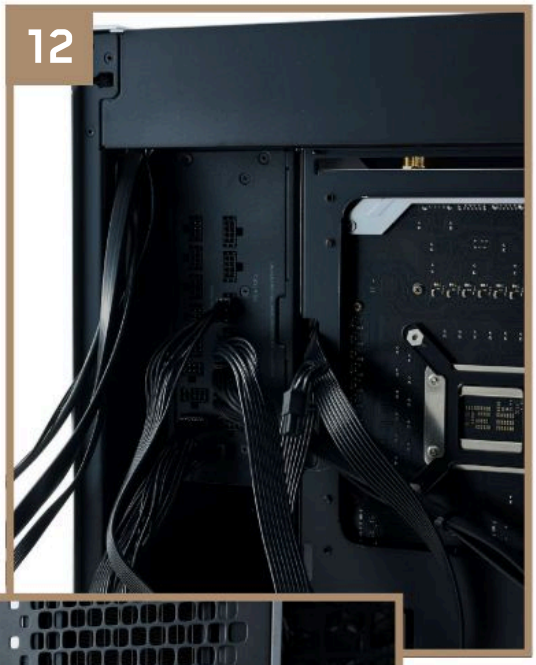

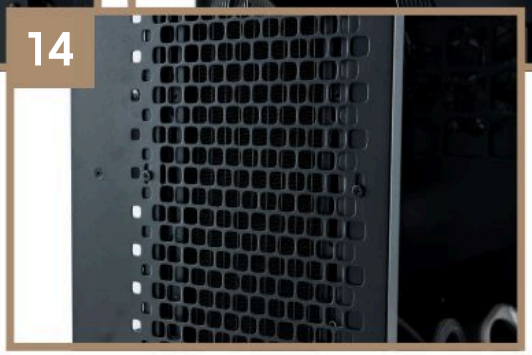

So we go to install the DDR5 as normal, orienting predominantly using that logic first, only to find that the notch doesn't fit. Panic sets in (we've checked the specs, right? This fits DDR5?) Yep, it does—what we didn't take into account was NZXT rotating the DDR slots the other way round. A quick memory shuffle later, and our DDR5 was in position, just going to show the value of those offset notches [Image 9]. Is it annoying that our labels are facing the wrong way? Well yes, but given the unique nature of our chassis, we've kind of lucked out here. Let us explain.

IN SHE GOES

With everything prepped, the motherboard goes into the chassis. We drop the case down carefully onto our work surface (definitely not the brand new clean perspex—our photographer has told us that we can't scratch it). Then, lower the motherboard into the chassis. The unique element of the Model 8 is that the motherboard itself is entirely rotated, so the rear I/O ports are facing upwards [Image 10], eliminating our memory label problem from earlier, and giving us an incredibly unique design.

There's a ton of room in here, too. The Model 8 supports E-ATX motherboards up to 12 x 13-inches, GPUs up to 400mm or 15.7 inches in length, and CPU towers up to 170mm or 6.69 inches tall. It's a beast.

NEXT STEP QUANDARIES

Here's the thing: this is the moment when we needed to take a step back and think about our next step. Getting the motherboard in is almost always our third step, but from here, it'll depend on clearances and access.

In this chassis, we decided that the smart move was to go with the power supply next. Interestingly, the PSU is held in the front of the chassis in that black metal PSU cover. Again, it's rotated, with the kettle-lead facing up and the fan facing inside of the case itself, acting as a secondary exhaust. That's a pretty genius design, as it effectively gives us access to an additional 140mm exhaust here.

Getting the PSU in, however, given the odd orientation when standing up, was a little finicky. We had to rotate the chassis so that its front panel was sat flush to the floor, then slide the power supply in from the side, and secure it in place with the included thumb-screws in a similar manner to how you'd do it in a traditional chassis. The advantage was that we had gravity to assist us in that endeavor, rather than trying to hold the PSU in place, and awkwardly secure those screws one-handed from the top [Image 11] .

The Corsair Shift works a treat. There's a decent amount of space in the rear of the case to allow for cable management, so those side-oriented ports make it easier to manage, particularly in the long term. Otherwise, all our ports would face down into the cable management space in the PSU shroud area.

After that, we installed the cables and began to pre-route them. The idea was that we could get all of this out of the way before moving on with the cooling solution (Corsair's iCUE Link setup always being a bit of a brain teaser). The CPU passthrough area is a bit of a squeeze too, and was a bit finicky to work around [Image 12]. As well as dealing with all of that, the front panel cables were hectic. You can spot them in the top left of that imagery, trailing off to the left (we've purposely moved them out of the way while we route everything into position).

FOR THE CHOP

At this point in time, and continually impressed with the flexibility of the chassis in question, this journalist knew it was going to be his main rig, and so as to avoid any hassle with certain connectors and cables, things got a little drastic in the *Maximum PC* lab. Here's the thing: the Model 8 has a ton of front I/O: one USB Type C 3.1, two USB Type A 3.0 ports, two USB 2.0 ports, HD audio passthrough, power, and reset. The thing is, none of these ports on our rigs ever get used. The USB C and 3.0 are handy, sure, and the power, but 2.0 ports? In 2024? Not for us. So out came the scissors, and off came the cables for both the HD Audio and USB 2.0. That's two less cables for us to worry about or manage [Image 13].

In the past, and in liquid-cooled builds where cable management space is at a premium, we've often removed pretty much every single cable possible, leaving only the power in place. The logic is that 99 percent of the time, the rear I/O is always decently accessible, faster, and drives better performance (particularly for audio compared to an HD passthrough), so why bother?

COOLING EXPERIMENTATION

One area that does somewhat let the Model 8 down is this rear exhaust fan mounting solution. It's a little odd, if we're honest. Geometric has added two 'rails' that you attach to your fans or radiator, which then attach to the case exterior using these Philips head screws [Image 14]. It works, but it's incredibly finicky to get right. It does give you some wiggle room when it comes to moving the brackets and the attached fans up and down, but it seems a little odd to add an extra step when you could just run fan mounting locations along the rear of the chassis in a more traditional manner.

The only logical reason we can think of is that this gives us more compatibility with 140mm fans, but given the solutions on other case designs that facilitate both fan-mounting sizes that simply fasten into the chassis directly, or with full-sized removable radiator brackets, it seems like a misstep—a minor one, but still.

Regardless, we got the iCUE Link cooler in and prepped, fixing the fans as exhausts drawing air from inside the chassis to the rear of the case. We slid the AIO unit all the way down the bracket [Image 15] in preparation for our bottom fans.

After that, it was onto the CPU bracket. Placing a small blob of thermal paste on the 14900K, and with the chassis on its back, we placed the Corsair CPU block on its mounting standoffs and secured the thumb screws in place. It's best to do this in a diagonal fashion, carefully tightening one side, then the other, a little bit at a time so as to not put too much pressure on the CPU itself during the procedure. You're not really likely to damage it, but it's better to be safe than sorry. Additionally, always do this manually with a screwdriver—not with any power tools or electric screwdrivers, as the torque can very easily overtighten it and potentially bend the PCB in the process [Image 16].

With the AIO in place, we moved on to the bottom fans, installing three 120mms in the floor, all nicely daisy-chained together. Judging by the product imagery, Geometric expects users to install the AIO in the floor first, as there's a ton of extra space for the radiator here [Image 17]. At this point, we still hadn't tinkered with the iCUE Link cabling either, but this was more just a preliminary fit.

Lastly, we added one extra 120mm into the roof of the chassis, which if we're honest, looks a little lonely. In hindsight, and given the huge amount of clearance between the motherboard I/O and the top of the chassis (a good 2.36 inches), it might have given us a cleaner look to install that last fan on the other side of those mounts instead. That way, it would still provide an exhaust and an additional light source, but with the added bonus of looking a little cleaner and less lonely at the top [Image 18].

CABLE HEADACHES

With our fans in position, we started on the cable management, and it was hectic. Sadly, there's very little support for this in the rear of the case. It could really use some cable tie-mounting locations—we jerry-rigged a few through certain gaps and mounting locations and made it work [Image 19], but it does feel a little messy compared to more modern designs. As a minor mod, you can also effectively get some 3m double-sided tape cable clips from Amazon to help, but again, this is not something you should be expected to do when purchasing a modern case.

As you can see, though, the RMx Shift is really pulling its weight here, and having all that space below it allowed us to hide not only all of our excess PSU cables, but also our iCUE Link hub. Speaking of iCUE Link, we've got it and the LCD screen connected via the two USB splitter to a single USB 2.0 port on the motherboard, then the PCIe power going to the hub. As each hub supports up to seven devices on each 'rail' or side, the AIO and top fan are daisy-chained together, then connect to the hub, and the three bottom fans then connect to the hub on the second rail.

MORE FANS?

That's when we realized, sat there, tinkering around with the cabling, that there's still a ton of space down in the floor, and Corsair's fans don't exactly have a ton of cables. Maybe we could fit a fourth 120mm fan in the bottom?

We grabbed one from our store cupboard, raised the rear exhaust AIO to accommodate it, and then added it to the daisy chain on the bottom block. Four 120mm fans locked secure in the base of the case and acted as intakes, providing a ton of light [Image 20]. One problem: those pop rivets in the front panel made contact with that last fan and pushed the front panel out. We could force the panel back in, but it would in turn push that last fan up at a slight 10-degree angle.

Taking the front panel off, we grabbed a hacksaw and cut off the bottommost pop rivets from the front panel. They're littered all over the front panel from top to bottom, and with those bottom ones cut out (again, the ones making contact with the fans). It sat perfectly flush again, and didn't bulge out. So yes, technically the Model 8 does only support three 120mm fans in the floor, unless you're a bit thuggish with a hacksaw and your cable management.

With that complete, we laid the chassis back down on its rear, removed the rear PCIe slot covers for the RTX 4080 Super, and slid it into place, re-securing it back down and installing the 12VHPWR600 cable into position. Given the size of the chassis, and where the PSU is situated, we did have to run the cable through one of the bottommost cable cutouts to make it reach, but otherwise, it's perfectly safe, and fits nicely. [Image 21]

SLEEK AND SUBTLE SOPHISTICATION

ALRIGHT, you can probably tell by this point that we're in love with this build. In fact, what you can achieve with this chassis is nothing short of epic. From a building standpoint, although it does lack the cable-management finesse we're perhaps more used to in the modern era, that rotated motherboard design and thermal channel, mixed with the absolutely insane amount of fans that we managed to install into it, made it a joy to work in. Combine that with the insane level of internal space to play with, and the surprisingly small footprint it takes up on a desk, and it's a remarkable piece of tech.

We did have issues, of course. That front panel and those pins were certainly among them, and the cable management quandaries with the lack of lockdown points, and the length of some of the items was mildly problematic (particularly for the 12VHPWR), but that's nothing some—you guessed it—custom cables couldn't fix.

Similarly, the topmost fan in a chassis like this looks very much out of place. In a more traditional orientation in any other case, it'd make a lot more sense in that position, but as a single exhaust in the roof, it just looks bizarre, despite being very useful for encouraging airflow. It performs fine, it just looks odd.

On initial boot, we did actually have a few issues with those rear exhaust fans (on the AIO) not powering on at all. Hilariously, this was due to them not being plugged in. For whatever reason, we forgot to actually plug the fans into the passthrough hub attached to the AIO, and just plugged the top fan into the hub, and the hub into the link. It's those little details that catch you off-guard using daisy-chaining solutions. Going from perpetually building PCs with masses of cables in a direct manner to a serial solution is just baffling at times. Still, a quick cable re-route, and one extra chucked in, and it was running like clockwork.

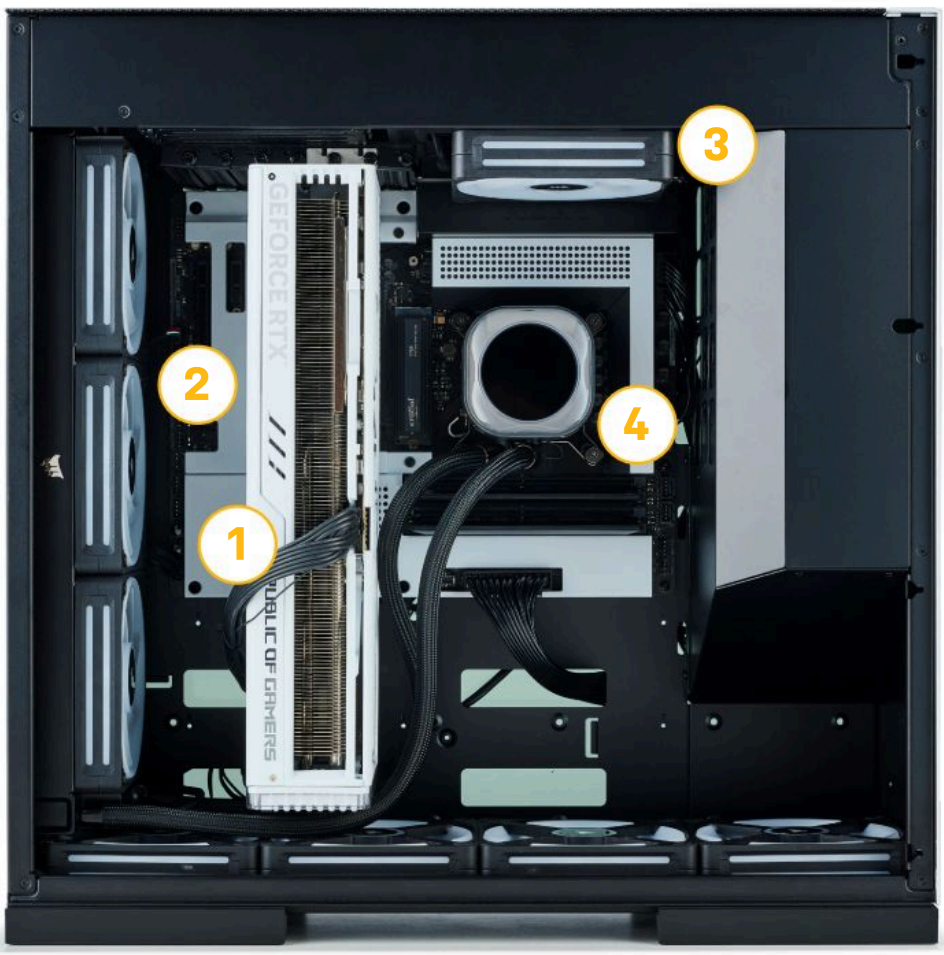

1 The power cable connected to the GPU is a little bit of a tight fit here. A custom cable would do it wonders, and help give us a little extra length to work with in the process.

2 Annoyingly, although the NZXT Z790 motherboard technically supports PCIe 5.0 SSDs, it doesn't on its M.2 slots, so we'll definitely be changing out the board as soon as possible.

3 This top-mounted fan looks a little lonely in comparison to the rest of the build. Mounting it externally will make a big difference to the overall look without compromising cooling.

4 Corsair's iCUE Link AIO CPU block has an option in iCUE to rotate the display. Some LCDs don't, so always double-check before you make your AIO purchase if you're using a rotated case like this.

BENCHMARKS

	ZERO-POINT	
Cinebench R23 Single-Core (Index)	1,909	2,083 (9%)
Cinebench R23 Multi-Core (Index)	22,221	35,600 (60%)
CrystalDisk QD32 Sequential Read (MB/s)	7,007	7,119 (2%)
CrystalDisk QD32 Sequential Write (MB/s)	6,026	6,819 (13%)
3DMark Fire Strike Ultra (Index)	17,670	17,612 (0%)
Cyberpunk 2077 (fps)	69	67 (-3%)
Cyberpunk 2077 RTX (fps)	54	51 (-6%)
Metro Exodus (fps)	87	79 (-9%)
Metro Exodus RTX (fps)	53	53 (0%)
Total War: Three Kingdoms (fps)	73	73 (0%)
Core Price	$2,620	$2,579 (-1.5%)

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

Our zero-point consists of the Hyte Y70 Touch build. Featuring an Intel Core i5-14600K, Gigabyte GeForce RTX 4080 Aero OC, MSI Meg Z790 Ace Max motherboard, 32GB of Corsair Dominator Titanium DDR5 @ 7200, and a 1TB Adata Legend 960 Max PCIe 4.0 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. "Core Price" refers to the key components generating performance (CPU, GPU, Mobo, OS SSD, RAM), not accessories.

THANK YOU, BUT PCIE 5.0 IS IN ANOTHER CASTLE

NOW WE GET to the pain. The frustration of finding out only after the fact that your motherboard doesn't actually support M.2 PCIe 5.0 is frustrating. Always read the fine print. It makes it doubly worse, as we're using two PCIe 5.0 drives, not just a singular drive in this build.

That said, it's not all doom and gloom, as those sequential numbers are still remarkably high. Additionally, the big advantage we have is that this build is predominantly designed with gaming in mind. The reason we picked those drives was because the Random 4K performance, particularly on single-thread, is considerably higher than that of the competition available for PCIe 4.0, and lo and behold, under scrutiny these things absolutely rip through our random 4K results. Although we can't include it in the table above, the T705 scored 89 MB/s on random 4K read and 300 MB/s on write, and the T700 netted 73 MB/s on read and 242 MB/s on write. Not by any means slow figures, although certainly not quite as quick as on a true PCIe 5.0 slot.

To that end, then, the only major change we'd make to this build is swapping the motherboard itself out for something of a slightly higher, more modern caliber, that has true M.2 PCIe 5.0 support. Still, we can't entirely blame NZXT on this one—this journalist should have read the fine print, after all.

Otherwise, overall performance is where we'd expect, the 14900K monsters through Cinebench clocking in a phenomenal 35,600 points in the multi-threaded task, and the single core slips in at just over 2,083.

In-game performance is also about where we'd expect to see it, with the RTX 4080 Super falling short of our overclocked RTX 4080 by a few percent here and there. The only major difference between the two again is that our Aero OC is overclocked by 6-7 percent in contrast.

Ultimately, our Geometric build is a build of two halves. On the one hand, it's been a fantastic all-round system-building experience, and we've been left with a phenomenal-looking, super-clean gaming PC. On the other, some hardware slip-ups and missteps on components have bitten us in the butt ever so slightly; something we'll definitely bear in mind moving forward.

THE 40 SUPER SERIES: SUCCESS OR FAILURE?

NVIDIA'S SUPER CARDS HAVE MADE WAVES, BUT IS IT ENOUGH TO GET YOU TO UPGRADE?

WELL, that wraps up our in-depth look at Nvidia's Super series cards: three builds, each with a unique GPU at their heart, and each GPU in turn unique in its strategic approach.

Without a doubt, these three cards are exactly that—unique in how they've been applied. Given its market-share dominance currently (Nvidia GPUs account for 76.92 percent of all GPUs, according to Steam), Team Green has had the opportunity to try out new strategies and techniques like this to see what boosts sales, and what doesn't. Each card went about its debut in a very different way, despite launching at similar times.

The RTX 4070 Super, on the one hand, had a huge internal hardware bump, bringing it far closer in performance to the older 4070 Ti, but at the same price as its predecessor. A 25 percent CUDA core increase is nothing to sniff at, and it really did drive up those frame rates. The RTX 4070 Ti Super was the middling child, with only a 10 percent hardware increase, but with an added extra 4GB of VRAM to better facilitate 4K and longer-term gaming. This is dependent, of course, on if texture size continues to dramatically increase in that favorite hobby of ours, but again, all at the same price as before. Lastly, the more awkward GPU of the three, the RTX 4080 Super—a meager five percent hardware increase, but with a supposed $200 price drop. It's arguably the most contentious of the three, and easily could be seen as a cash-grab/relaunch and marketing strat, rather than any serious new player in GPU development.

It's the last one that's particularly turbulent, as its success is far less dependent on its overall design, but rather on market conditions at the time. Right now, available stock of the $999 RTX 4080 Supers is pretty much non-existent. Cards that were retailing around that mark or just above have shot up by $200-300 in most cases, and the remaining RTX 4080s have disappeared as well. Whether that's due to a lack of stock, or malicious scalpers shorting the market to resell cards at a far higher value is up to you to decide, but it doesn't change the fact that these cards, and the significance of the RTX 4080 Super launch, is somewhat dulled because of it, which is a shame.

PURE PERFORMANCE

Market conditions and availability aside, the performance bumps that the two more junior cards have received is nothing short of impressive. The fact that we've finally got a $600 graphics card that's

RTX 40 SUPER SERIES SPEC LIST

	Nvidia GeForce RTX 4090	Nvidia GeForce RTX 4080 Super	Nvidia GeForce RTX 4080	Nvidia GeForce RTX 4070 Ti Super	Nvidia GeForce RTX 4070 Ti	Nvidia GeForce RTX 4070 Super	Nvidia GeForce RTX 4070
GPU	AD102-300	AD103-400	AD103-300	AD103-275	AD104-400	AD104-350	AD104-250
Transistors (Billions)	76.3	45.9	45.9	45.9	35.8	35.8	35.8
CUDA Cores	16,384	10,240	9,728	8,448	7,680	7,168	5,888
Tensor Cores	512	320	304	264	240	224	184
Ray Tracing Cores	128	80	76	66	60	56	46
Base Clock / Boost Clock	2,230 MHz / 2,520 MHz	2,205 MHz / 2,550 MHz	2,210 MHz / 2,505 MHz	2,340 MHz / 2,610 MHz	2,310 MHz / 2,610 MHz	1,980 MHz / 2,475 MHz	1,920 MHz / 2,475 MHz
Memory Size	24GB	16GB	16GB	16GB	12GB	12GB	12GB
Memory Type	GDDR6X	GDDR6X	GDDR6X	GDDR6X	GDDR6X	GDDR6X	GDDR6X
TDP	450W	320W	320W	285W	285W	200W	220W
RRP (USD)	$1,599	$999	$1,199	$799	$799	$599	$599

capable of smashing 4K gaming at 60 fps is pretty epic to see. Yes, it's not quite flawless in that regard. Arguably that's still a huge investment up front for a GPU, particularly when something like a PS5 will set you back far less for similar levels of performance (if not graphical fidelity), but nonetheless it does show that we're seeing significant progress still, if only to mildly keep AMD at bay with its RX 7000 series units.

The Ti Super was a little more lackluster than the stock RTX 4070 Super, that extra 4GB of VRAM didn't exactly translate into an absolute barrage of extra frames, but it's more about future-proofing than diving into the heady realms of performance. Whether that's worth the extra $200, however, is difficult to justify, particularly if you're on a budget.

WHAT THE FUTURE HOLDS

The real trick will be to bring that 4K 60 fps mark down in the next generation. The GTX 1060 when it launched revolutionized 1080p gaming by consistently smashing that 60 fps barrier in a number of different configurations for less than $200, and that's where the RTX 5060 needs to really land if we want to see 1080p relegated to the pages of history for PC gamers.

Right now, there's not a whole lot to go with regards to the 50 series. We expect Nvidia to launch the cards likely some time at the end of 2024, possibly early 2025, likely utilizing the Blackwell architecture on TSMC's 3nm process. How Nvidia and TSMC are circumventing the quantum tunneling effects that happen at such a small transistor size will be anyone's guess, but there are rumors that the architecture could bring a doubling of overall performance compared to its current generation.

It also looks like Nvidia could be moving away from its monolithic architecture design, and instead opt for a chiplet-style design, similar to what AMD is using both for Ryzen and its RDNA GPUs. This does give the company a lot more flexibility in its GPU design, as you can keep certain elements of the GPU itself (for instance, I/O and other, less performance-heavy operations) on easier and cheaper to manufacture larger transistor sizes, and then focus your more complex and performance-oriented hardware (CUDA, Tensor, and RTX cores) on the 3nm process instead. This could, and should, bring down the costs for the end user while boosting performance. Given that Nvidia is the last manufacturer to still use monolithic designs, and that it's also using multi-chiplet designs in its enterprise-grade AI solutions, it's almost guaranteed that this shift is going to happen.

The Blackwell architecture is currently only found in Nvidia's high-end AI model enterprise units.

WHAT IS QUANTUM TUNNELING?

Quantum tunneling has been a headache for Moore's Law enthusiasts for some time now, particularly in regard to processors and transistor size. In a very simplified explanation, how a transistor works (effectively flipping from 0 to 1) is that a current passes through a gate (an insulating material, or semi-conducting material, in our case, silicon) that is heated up via a secondary current. As it heats up, the resistance decreases and the primary current can pass through switching that transistor from 0 (no current passing through) to 1 (current passing through).

The problem, however, is that once you get to such small transistor sizes, electrons can effectively jump through the semi-conducting gate, without that gate being opened by any secondary current. That's called quantum tunneling, and as you can imagine, creates all sorts of problems for programs that are dependent on absolute 1s and 0s.

At the moment, this is commonly seen at transistor sizes smaller than 1nm. We have seen some significant developments in the field to attempt to get around this right now. However, nothing concrete has arrived just yet that's economically viable.

Instead, scalability and multi-chipset designs have been proposed to counteract this limitation by simply increasing the physical size of processors, but with a caveat of increased heat and power-draw as a result.

At this point, until quantum tunneling has been sidestepped with traditional methods, brute-forcing and double transistor size is effectively out of the question. Manufacturers will need to work smarter, not harder in how they approach their CPU and GPU designs.

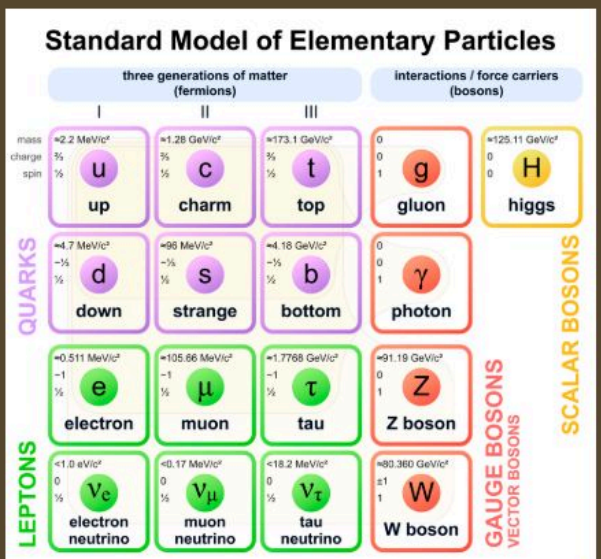

Those pesky electrons are causing all manner of problems.

LIQUID LUNACY: ADVANCED PC COOLING IN 2024

Is it worth liquid-cooling your PC in this age of non-overclockable chips?

WE OFTEN DESCRIBE PC building as Lego for adults. For the most part, modern systems are incredibly easy to construct. There are, of course, certain pitfalls that you can encounter, but most of it comes down to cable management, clearances, and hardware compatibility. Graphics cards, CPUs, DDR memory, M.2 SSDs, and even the very connectors that power them, are designed to make the build process as simple as possible. Even modern cases encourage users to build their rigs in a certain way to avoid the pitfalls and errors that used to be commonplace.

On the one hand, this makes our hobby incredibly accessible. There's room for refinement, of course—cable-tidying tricks, design choices, hardware combos—but at its core, it's an enjoyable recreational pastime that with enough knowledge is fairly easy to accomplish. On the other, for some, it lacks the thrill, panache, and knowledge of implementing a skill set you might find in other hobbies.

That's where custom liquid cooling comes in. It is the absolute pinnacle of PC building, giving you an unprecedented amount of freedom in creating something that's unique to you. In an age of auto-overclocking and temperature-dependent CPUs and GPUs, pulling as much heat away from these as fast as possible leads to far better performance. Because of that, liquid cooling has never been as valuable as it is today.

But where do you start? How do you make that step from off-the-shelf products to custom bespoke designs? Let's slide into the pool and take a look, shall we? You won't regret it. –ZAK STOREY


START
RESET
HYPERX
HYPERX
HYPERX
HYPERX
GEFORCE RTX
GEFORCE RTX

A HISTORY OF LIQUID COOLING

Where it all began

FOR AS LONG AS there have been computers, cooling has been a problem, particularly as CPUs have continually increased their transistor density, clock speed, and power draw. In the early 1950s, machines like the UNIVAC1 and IBM's System 360 famously utilized liquid cooling to alleviate the excess heat build-up in their machines. In fact, IBM saw that this problem was only going to get worse all the way back in 1965, and invested heavily in researching solutions for the issue, developing its own 'Thermal Conduction Module', which acted as a rudimentary water block, built with thermal pins surrounded by helium gas to accelerate heat transference into the block and away from their machines.

Most famously, the Cray-1 was also one of the first to use a liquid-cooling-esque style solution, utilizing a refrigerant coolant design to cool the machine itself. This very much acted like a large-scale traditional air-cooled heatsink, however, with the refrigerant being channeled through pipes connected to cooling bars, which directly made contact with copper cores attached to the modules within the machine. That coolant was then circulated back to a condenser, where the heat was extracted and dumped into the surrounding air.

The Cray-2 supercomputer was far more challenging to cool than its predecessor, and the company opted to switch to full immersion cooling instead, moving away from the humble waterblock-style designs, and submerging the entire supercomputer into an inert coolant. That still requires circulation, and a pump is necessary to move it to radiators, where heat is pulled and then pumped into the atmosphere. However, the thermal capacity of that fluid was far greater than what could be achieved with prior techniques.

MODERN COMPUTING?

With the rise of the PC in the early '90s, and humans being the ever-adventurous species that we are, many home enthusiasts, eager to amplify the performance of their machines and frustrated with air-cooled solutions, began to manufacture their own liquid-cooling components and loops. These were typically made from repurposed plumbing supplies, car parts, and hand-made blocks designed to allow the CPUs of the day to perform at higher frequencies when overclocking while still remaining stable.

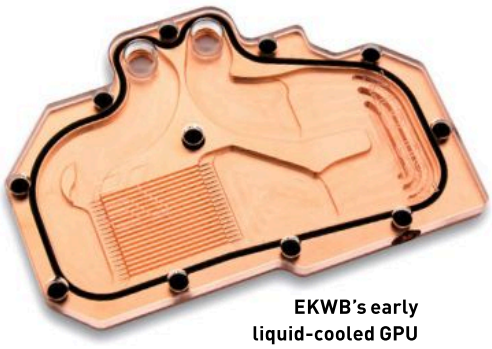

EKWB's early liquid-cooled GPU waterblocks were incredibly simple yet efficient by today's standards, this one's for the GTX 580.

Copper was used, and the G1/4 fittings in particular were commonplace and easy to come by. Some enthusiasts even devised a few pure silver water blocks instead of copper, as the thermal conductivity of it is far greater than that of the copper blocks we know today. However, that came with a significant increase in cost.

As CNC machines and software became more developed and easier to use, it led to a rise in far more sophisticated blocks, with smarter designs developed to minimize whirlpooling effects within the loop and other such problems that typically lead to a decrease in thermal capacity throughout the loop.

THE BIG MANUFACTURERS

It wasn't until the early 2000s that mainstream companies, such as EKWB and Alphacool, would enter the consumer liquid-cooling world. EKWB, first founded in 2003 by Edvard König in Slovenia, was one of the first to enter the consumer market, specifically with the aim of manufacturing what was considered at the time to be 'extreme-cooling' solutions, and improved waterblocks to minimize noise and improve thermal performance. It launched its first full-cover water block for GPUs in 2004 with the Wave Water Accelerator, along with a series of XTX radiators and soft-tubing solutions.

Alphacool similarly started operations in 2003 as well, developing and releasing a line of direct-on-chip CPU and GPU coolers (specifically cooling only the chips themselves, rather than any supplemental componentry, such as MOSFETs VRAM or any other capacitors). The company also acquired Laing and Xylem in 2006, who

The Cray-1 Supercomputer was the first to really take advantage of the liquid cooling that we think of today.

© CLEMENS PFEIFFER/ WIKIPEDIA/DEUTSCHES MUSEUM, EKWB

are most notable for the creation of the Laing D5 pump, still used in liquid cooling to this day. It wouldn't be until 2007 that Alphacool introduced its own full-sized GPU waterblock to the market, and 2008 with its first full radiators.

LATE 2010s

In the modern era, a whole swathe of manufacturers have arrived to the custom liquid-cooling game, from the likes of Corsair, to Thermaltake, Phanteks, Barrowch, and more. The field has blossomed massively, particularly with the help of YouTubers and modders focusing on it over the past decade.

A number of complete DIY kits are now available to make liquid cooling easier, along with a plethora of PC build configurators online.

ASETEK AND THE AIO PUMP DESIGN

We can't talk about liquid cooling without mentioning the Danish company Asetek, either. Founded in 1997 by André Eriksen, the company specializes in a mix of data-center and desktop-cooling solutions.

It famously patented the first OEM AIO pump design back in 2003, effectively meaning any manufacturer that attached a pump or impeller to a CPU block would have to pay royalties or find a different solution under EP 771 (European) or US patent 517924 P (with several others). Asetek has sued multiple manufacturers over the years for infringement of that patent, most famously losing to Cooler Master in 2017 in the Netherlands, and CoolIT in 2022.

Interestingly, the patent looks to be expiring in 2025, meaning we may very well see a spate of new pump designs in the years to follow once it finally lapses.

The Asetek pump design was famously used in almost all AIOs

BUT WHAT ABOUT THE COOLANT?

As for coolant, the development in this field has been a mix of two different angles. It's well known that one of the most efficient coolants out there, particularly for DIY customers, has been de-ionized or distilled water. Certain refrigerants are available, but these are typically reserved for top-tier enterprise-grade machines that require significantly greater thermal capacity than with the likes of any consumer PC.

Because de-ionized water is so readily available and affordable with such high thermal conductivity, the need for any advanced coolant, at least in the DIY PC space, has been fairly minimal.

For the longest time, Mayhems, a British company, set up by a slightly eccentric ex-army officer called Michael Wood, was the go-to brand for any would-be liquid cooler. The company famously developed its Aurora line of particle coolants that held suspended metallic flakes within a glycol-based coolant to not only sufficiently cool the components inside, but also produce a mesmerizing visual effect. It took a number of iterations to develop a composition that didn't build up internal gunk within the blocks themselves, though.

There were two variants available on first launch: Aurora 1 and Aurora 2. Aurora 1 was specifically designed for show-builds with larger particles giving a far greater dynamic effect in the loops in question, while Aurora 2 was a longer-life variant with smaller particles that lasted for a few months before the reflective flakes fell out of the fluid. Mayhems also produced a line of opaque coolants, known as the Pastel line, which provided a completely opaque solution for your coolant, instead of the transparent mixtures available at the time.

Mayhems also worked closely with the University of Cambridge to develop a 6nm 2d particle coolant (which you may have seen in our November 2020 issue) that had significantly better thermal conductivity than water, and actually improved the performance of some water blocks over time by filling in the imperfections in the CNC'd metal internally.

We tested Mayhem's XTR 6nm coolant back in 2020.

Aurora 1 was fantastic for show builds, but less so long term.

WHY SHOULD YOU LIQUID COOL?

What's the point of all this in 2024?

GOOD QUESTION. As you probably spotted, liquid cooling was first and foremost developed as an antidote to the excessive temperatures and noise generated by overclocking CPUs and GPUs. The more you increase the voltages and clock speeds going into those parts, the greater the excess heat produced. To go beyond the limits of our air coolers of the day, the logic was simple: greater cooling to facilitate higher clock speeds, at least until the chip itself became unstable.

In our own experience over the years, liquid cooling has allowed us to do that and more. We predominantly used it to expand our overclocking capabilities, particularly when paired with delidding CPUs (the act of removing the integrated heat spreader and replacing the standard thermal paste with liquid metal or something similar). Alternatively, if you were more focused on the audible sound of your PC at idle or in-game, you could undervolt your hardware instead, dropping temperatures further, and by increasing the surface area available through those radiators, significantly shift temperatures lower and have your fan-speeds reduced, too.

The biggest advantage, however, particularly in 2024, is just how far we've come in terms of auto-overclocking solutions. Almost every single Intel 14th gen chip can now quite comfortably sit at 100 Celsius for any length of time during heavier workloads. In fact, their ability to fluctuate clock speeds like that is now so dynamic that you'll see performance significantly change, depending on what kind of cooling the chip has to hand. Our 14700K build performed 10 percent slower under its air cooler in the ITX Hydra build than it did as standard under a 360mm AIO. The same goes for GPUs and AMD's Ryzen lines, too. The firmware, software, and BIOS are now so advanced, and the competition so fierce that edging out any temperature you can allows your CPU to really let rip at higher speeds for longer. Expanding that capacity with the help of more surface area and higher thermal capacity in your fluid should lead to significantly greater performance over longer periods of time.

No matter how you cool Intel 14th gen chips, they're going to hit 100 C.

HOW DOES LIQUID COOLING WORK?

Bizarrely, liquid cooling operates in a very similar manner to what you'd find in a combustion engine in a car. But before our motorized analogies, let's start with the basics. All cooling is based on the same principles—our aim is to transfer the heat generated from the component in question, and move it away and out of the chassis.

Regardless of whether you use an air tower, all-in-one liquid cooler, or full custom loop, the basic process is always going to be the same. You have a coolant or a vapor that's cool, it makes contact with the heat spreader or contact on the CPU or GPU, absorbs the heat from that component, and then, either via the power of convection or a pump, moves towards a radiator or fin stack of some description, where cool air (usually drawn over the fins by a fan) then pulls the heat out of the radiator and into the surrounding environment (like your office). In fact, this is why most data centers and server rooms have incredibly aggressive air conditioning.

In the case of an air tower (or any cooler or heatsink that utilizes heat pipes), there's a special liquid inside each pipe that has a significantly lower boiling point than water. Once it makes contact with the heat at the bottom of the pipe closest to the component, it heats up, evaporates into a gas, rises to the top of the heat pipe, where that cool air is passing over those fins, cools down, condenses back into a liquid and then falls back down to the CPU or GPU to start the procedure again. This is an incredibly rapid procedure that takes very little time to occur and uses convection as its basis, making it exceptionally reliable. The downside, however, is that heat pipes like this can only be built so large, and the maximum thermal capacity of the liquid inside is relatively minor in comparison to that of any solution found in a custom loop or AIO.

Speaking of, custom loops and AIOs operate in a similar manner to that of your car. The heat source (in this case, an engine or CPU) generates that heat, and then a series of pumps move a coolant

EKWB now sell Direct Die AIOs for Intel's 14th gen processors with specially designed CPU blocks for them.

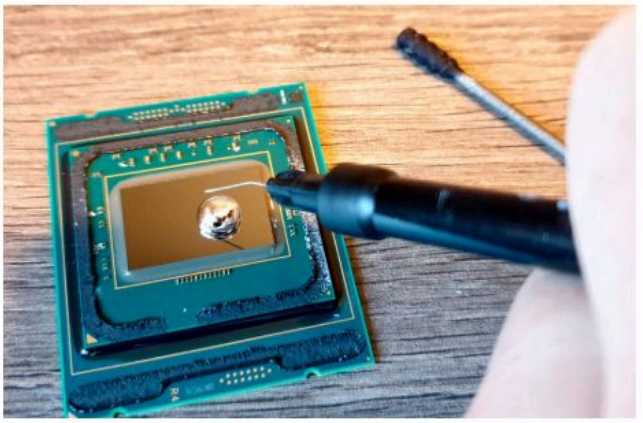

Delidding a CPU can be incredibly scary, but well worth the effort if you're replacing the thermal paste with liquid metal.

(typically glycol based) through the system to a radiator with fans attached that expel that heat into the air. The major difference between air towers, AIOs, and custom loops is the volume and capacity available to them. Regardless of what style you use, it's all about capacity in the fluid, and surface area on your fin stack or radiator. No matter how dense an air tower is, it cannot compete with a significantly larger AIO, particularly when we have radiators that are 360mm, 480mm, and even 520mm in diameter, with four fans of varying sizes. Couple that with the increased thermal capacity and volume of fluid available (disregarding any specialist coolants), and what you're left with is a far more efficient system.

IT'S ALL ABOUT THE BOTTLENECKS

There are limits to that formula, however, and if you're not careful you can get roadblocked pretty hard and fast by one bottleneck or another. In the case of a custom loop (and even an AIO), this bottleneck typically occurs between the CPU, its IHS (integrated heat spreader, the metal plate situated on top of the silicon itself), the TIM (thermal interface material sat between the CPU and its IHS), the water block, and the thermal paste connecting the two.

Transfering the heat between these elements has an effective maximum bandwidth, and regardless of radiator size or fluid in your loop, there comes a point where the heat can't be transferred fast enough. It's a case of diminishing returns, where the more radiators you add won't necessarily reduce temperatures or improve clock speeds further.

There are ways around this; delidding has become increasingly popular over the years. Depending on the processor, the TIM connecting the silicon to the IHS may be a basic thermal paste, rather than soldered directly on. Although cheaper to manufacture and produce, thermal paste chips have a lower thermal conductivity than direct metal solder. Delidding is the act of removing the IHS from the CPU, usually with a specialized tool, separating the glue from the IHS, removing the thermal paste, and then replacing that paste with a liquid-metal substitute.

In our own testing, we've seen temperatures on older processors drop by upwards of 10 to 18 C under load by doing this. There are a number of manufacturers out there, Der8auer being the most famous, that have the tools necessary to delid your CPU. Typically, liquid metal is applied, and the IHS then re-glued back into position during the process. However, EKWB and a number of others have created direct-die waterblocks and AIOs, specifically built to be used with a delidded CPU without an IHS at all, removing that bottleneck entirely.

Again, this is entirely dependent on the CPU in question. Many of AMD's and Intel's processors have been soldered in the past, and trying to remove the IHS in these conditions can significantly damage the chip if you try to do so. It's always worth double-checking online first.

THE FUNDAMENTAL BASICS OF ALL PC COOLING TODAY

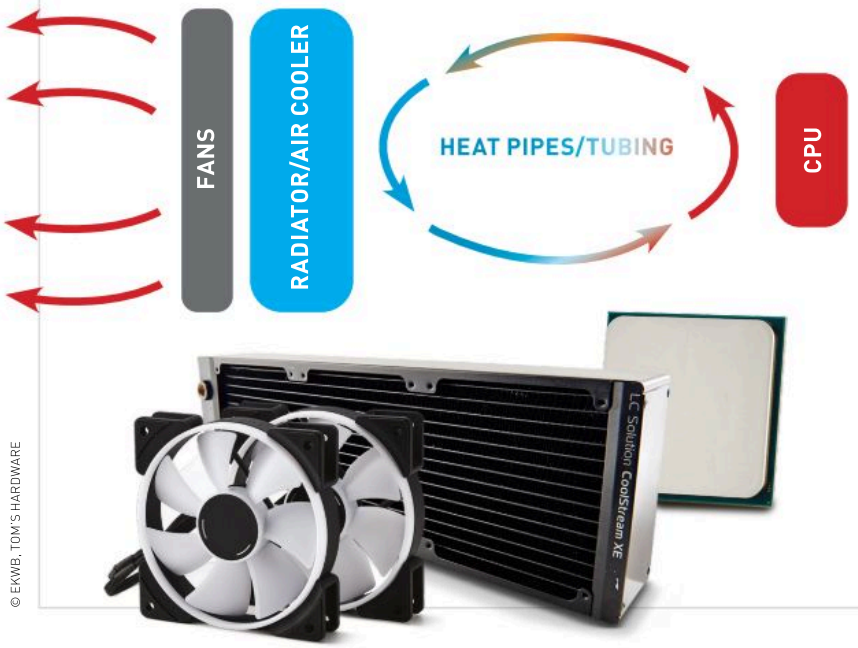

© EKWB, TOM'S HARDWARE

LET'S WATERCOOL OUR MACHINE

We've covered the why, but what about the how?

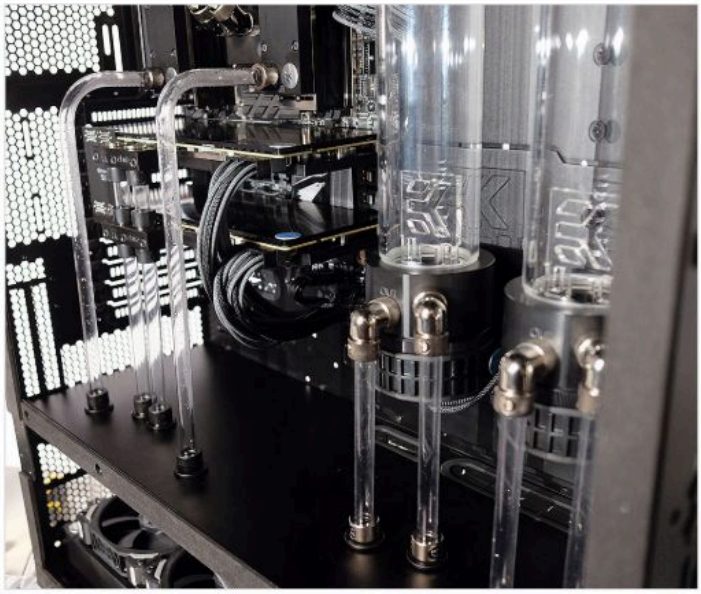

Dual loops—one for each component—can be done, but it does require a lot of extra hardware, and a ton of work.

WE COULD spend another six pages explaining the finer details on how to build the perfect liquid-cooled PC from start to finish, but as we're saving that for a future Dream Machine in an upcoming issue, we're going to give you the best crash-course advice we can on how to get yourself going on your own liquid-cooled adventure.

Let's face it: there are a fair few more considerations that need to be made before committing to the fully custom liquid-cooled build. These include what you're cooling, how you're cooling it, what components and case you're using, and even styling choices need to be planned well in advance—not to mention what your budget is.

At first glance, It can be quite daunting, so let's break this down into steps to make it a little more digestible.

1 THE PLANNING STAGE

First, you need to take a step back and really build out a plan. Decide exactly what your chassis will be. Ideally, for an easy build, you're looking for a large case with plenty of space for radiators, good airflow, and somewhere to put a pump and reservoir of some description, or a combo unit (alternatively, you can attach these to a radiator with a specialized bracket). It also helps to have sufficient space for cable management, as that extra componentry is going to have its own SATA power and fan headers.

You can build a liquid-cooled PC in an ITX chassis. However, space constraints and thermals are more challenging under these conditions, so we recommend an ATX mid-tower or larger.

Then you need to decide what you want to cool and how. Do you want to do just your CPU, or the GPU as well? Two separate loops, with one for each component? Or combine them into one continuous loop? All are feasible options, and shouldn't alter the overall thermal dynamics much, but a second loop will add additional parts and cost, which is always worth bearing in mind. We'd recommend withholding those fancier builds for a larger chassis.

Then, you need to decide what tubing you'd like to run. For a first build, we'd always recommend going with soft-tubing to start. It's easy to use (only requiring a good pair of scissors to cut), and can look great, particularly if you go with a premium solution, or prefer a hard industrial style with black rubber tubing (EKWB has a fantastic ZMT variant).

Hard tubing solutions are available both in the form of PETG and Acrylic (always go for the latter—it's tough enough, and holds up better in sunlight), but these can require tube bending, using a heatgun and a silicone insert, to create the angles you'll need to connect your components. Alternatively, manufacturers provide a wide variety of angled female-to-female connectors, allowing you to bypass the tube bending, but this will increase the cost of your build, as each hard 90-degree or 45-degree bend will require an additional two compression fittings, and one angled female to female connector.

Rubber tubing can look incredibly classy, even compared to hard tubing.

Beyond that, there's copper and glass pipe, both of which require specialist materials to shape, cut, and bend, but they do look incredible once finished.

2 DRAWING STAGE

Once those overarching design decisions have been made, your next action is to draw out your build. Take a pen and paper, use Photoshop, or spend far too long in CAD software, and draw the outline of your chassis, where your components will go, and how you want your tubing runs to connect. You can be quite flexible in how you plan your loop, but here are a few things to consider:

1. **Your pump should always be gravity-fed by a reservoir.**
2. **If using hard-tubing, try to stick to a single 90-degree bend per length.**

2.1 **If not using hard-tubing, bear in mind the rigidity of your tubing and where it might kink if the bend radius is too tight.**

2.2 **Ensure you have enough 90-degree and 45-degree fittings to accommodate the above.**

3. **Ensure you have a fill port and drain port for easy maintenance, preferably with a ball valve attached to the drain.**
4. **Make sure your case has enough clearance for your radiators and fans together.**
5. **The radiator's position in your loop doesn't matter.**

With your plan drawn up, you can count how many fittings you'll need. We recommend going with a good-quality compression fitting, making sure to match the size and type of the fitting to the tubing you're using. Hard tubing requires different fittings to soft tubing, despite both acting as compression fittings. The soft tubing will use a barb, plus the tube itself to act as the o-ring, followed by a compression cap that tightens over the top, while the hard-tube fitting will likely have multiple o-rings integrated into the fitting, along with a compression cap.

In a bare-minimum custom loop with both the CPU and GPU cooled, you should have one 360mm radiator, one pump and res combo (or separate products), one CPU block, and one GPU block. This will require a minimum of 10 compression fittings, not including any rotary-angled male-to-female connectors.

3 PRE-BUILD STAGE

You've planned your build and got all the parts. The next step is to take apart the graphics card, right? Wrong. Before going anywhere near liquid-cooled componentry, you need to get these components working in a normal system. That can be on air, with a spare AIO—anything. The last thing you want is to put your system together, build it perfectly, and fill the entire loop, only for the motherboard to fail, the memory to be incompatible, or the power supply to be a dud. So get those components tested, Windows-installed, and prep your workbench, because it might take a while.

4 BUILD ORDER

Build order is something that we stress flexibility on. In the case of liquid cooling, it's a little more rigid. Aside from pre-build testing, you want to disassemble your GPU and get your block attached, then prep your motherboard, M.2, CPU, and RAM in the same manner, get that all installed in the chassis, including the graphics card, then move on to your cooling solutions—particularly the radiator (with fans), reservoir, and pump.

Dry-runs are useful here. You may need to move critical components, and reinstall them after, depending on clearances and how things fit. Once that's done, install your fittings and angled connectors, and size up your tubing runs.

These can then go into the build after that. Then, and only then, do you start cabling, as it's typically far easier to install than your tubing runs will be. Once this is complete, double-check that every G1/4 port either has a nice tight tubing run and fitting attached, or that it's plugged. You can also use a loop leak tester (a pump that pushes air into your loop, with a dial on it indicating pressure). If the loop has a leak, the pressure will drop over time.

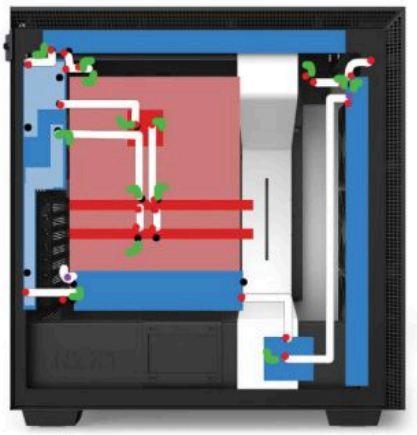

A good plan goes a long way.

Build order seriously matters when you're working on a liquid-cooled build.

After that, you can fill your loop. You'll need to either use a 24-pin ATX bridge on your build's power supply (remembering to unplug the CPU power), or do the same, but with a secondary spare PSU, and attach the pump power to that unit.

5 PUMP CYCLING

Inside your pump lies an impeller. The lubricant for this is typically the coolant itself. When you first fill your loop, you'll need to prime the pump. To do this, fill the reservoir until it can't hold any more coolant, then seal it. If gravity is feeding your pump, some of this coolant will fall into the impeller and protect it from damage. With your 24-pin bridged and pump connected to the power supply in question, you then need to flick the power switch on the back of the PSU.

Watch the coolant drain out of the reservoir and make its way around the tubing and blocks, then open the reservoir and repeat the process until you have a continuously flowing loop.

In an ideal world, you should let the loop run like this for 24 hours before configuring the system (placing paper towels under fittings and on top of components can protect them, and to also indicate any leaks). This way, you'll be able to spot problems, and allow the system to move air bubbles out of the blocks, as these will affect thermal performance.

Once this is done, you can then reconnect everything and get your system up and running.

PARALLEL OR SERIAL COMPONENT CONNECTIONS

Depending on your loop configuration, you can configure it so two tubes will run coolant to and from your components. For instance, you could run a pump reservoir to your GPU, up into your CPU, back down into your GPU, and into the radiator. This is what's known as a parallel setup. It's less efficient at cooling your CPU, but typically only if your coolant isn't moving fast enough around your loop as that heat builds up around your CPU.

Due to the fact that the loop itself is pressurized, and as long as the speeds the coolant is flowing at are fairly high, temperature in the coolant should remain fairly consistent. Still, if you don't want to run your pump at high speeds, we highly recommend running a serial configuration instead.

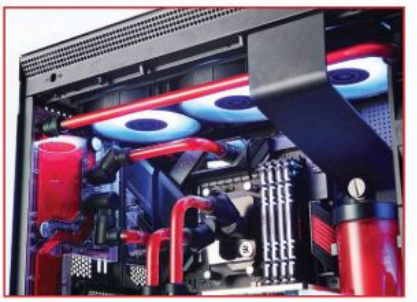

In our hypothetical scenario, that would involve running the pump reservoir to the GPU, the GPU to the CPU, and then the CPU to the radiator instead. Just make sure that you follow the directions on the CPU block regarding pump flow direction.

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

MAXIMUMPC

PREMIUM DIGITAL SUBSCRIPTION

As a premium subscriber you instantly get access to 100+ back issues

www.magazinesdirect.com/MPP

*Price in US $. Offer is valid on US orders only, visit us at www.magazinesdirect.com or call 1-844-779-2822 for other subscription options.
Offer valid until May 30 2024.

COPILOT PRO & COPILOT FOR MICROSOFT 365

Why Copilot is great, when it's awful and which version you should buy

CONTRIBUTORS: Ian Betteridge, Tim Danton, Adam Timberley

AH, MICROSOFT. How we love your ability to create world-dominating software with one hand and sow world-beating confusion with the other. In January, finally, it announced that Copilot for Office was available for all. Except, this being Microsoft, calling it Copilot for Office would be too simple. Instead, we now have the free Copilot, Copilot Pro, and Copilot for Microsoft 365 (and Windows Copilot, GitHub Copilot, and a trio of Sales, Service, and Security Copilots, but let's ignore these distractions).

Now, vanilla Copilot has nothing to do with Office. It's the successor to Bing Chat and, unless you choose to download the Copilot apps for iOS and Android, you'll access it via the web. As it's based on GPT-4 and GPT-4 Turbo, though, it remains a powerful tool, especially considering that Microsoft doesn't charge a penny for it.

Copilot Pro, which costs $30 per month, is a different beast. If you already have a Microsoft 365 Personal or Family subscription then a Copilot button inveigles its way into the main apps: Word, Excel, PowerPoint, OneNote, and Office. However, let's lay a big fat caveat here: it's only for the desktop apps on Windows for now. Apple fans are restricted to web apps and iPadOS. As you'll discover, this 'it's coming soon' mantra is a recurring feature for Copilot Pro.

There are two other good reasons to consider paying the $30 per month. One is that you get 'priority access' to GPT-4 during peak times, so you both get to feel smug and wait shorter times for results. The other is that it gives you 100 daily 'boosts' in Microsoft Designer, which is powered by OpenAI's DALL-E 3, compared to 15 from plain Copilot. A 'boost' translates into fast GPU time, so again you'll rarely find yourself waiting for images to be created.

Finally, we come to Copilot for Microsoft 365. This is Microsoft's AI heavyweight, adding full integration with SharePoint and adding Teams to the mix. For now, Teams is the best Copilot integration by a distance. As we'll discuss, it has the potential to save employees a lot of time, but also requires your business to invest heavily in IT time—particularly during the setup stage—and money.

Copilot sits at the heart of Microsoft's strategy and its Microsoft 365 offering.

	FOR INDIVIDUALS		FOR ORGANIZATIONS	
	COPILOT	COPILOT PRO	COPILOT	COPILOT FOR MICROSOFT 365
	Free	$20	Free	$30
Foundational capabilities	●	●	●	●
Web grounding	●	●	●	●
Commercial data protection	●	●	●	●
Priority GPT model access		●		●
Copilot in Outlook, Word, Excel, PowerPoint, and OneNote		●		●
Copilot in Teams				●
Microsoft Graph grounding				●
Enterprise-grade data protection				●
Customization (GPTs)		Copilot GPT Builder		Copilot Studio

Prices are monthly and direct from Microsoft

Until recently, Copilot for Microsoft 365 was limited to businesses with over 300 seats. Now, any size of business can buy it—even a one-seat business, as it works with any Microsoft 365 Business Standard or Business Premium subscription. Larger businesses will need a minimum of a Microsoft 365 E3 subscription.

So, what is Copilot? In short, it's complicated. Here to hopefully answer all your questions is our guide to what Copilot can do for individuals and for businesses.

COPILOT FOR MICROSOFT 365 OR COPILOT PRO?

Copilot Pro is designed for individuals. It gathers its data from the document you're working on and from the web; it won't look any further. Disappointingly, it doesn't yet search your personal OneDrive documents, although we expect this feature to land at some point this year.

Copilot for Microsoft 365 is the much bigger brother, aimed solely at businesses. With no set minimum of employees, however, any size of business can benefit from Copilot for Microsoft 365's key feature—that it can gather all your relevant data stored on SharePoint, whether it's in Word docs, Excel sheets, PowerPoint presentations, or Teams calls. The Microsoft 365 roadmap says that OneDrive will also be searchable in May.

Adam Timberley is an IT consultant, and one of his clients asked him to be part of an early team using Copilot for Microsoft 365. "I found Copilot was really good [for Teams]," he said, and its skills went well beyond transcriptions and action points. "You could go back and ask it questions. What did this person say? Why did they say it? You could even ask it for its opinion. It was sometimes useful, but it did go down a creative route sometimes and you ended up getting things wrong." So, is it as reliable as talking to a work colleague? Adam laughs. "Depends on the work colleague."

Copilot for Microsoft 365 is also powerful due to its access to your data. Microsoft calls this process 'grounding', so where Copilot Pro uses "web grounding"—that is, it uses the internet (and sometimes the active document) as sources to answer to your queries—Copilot for Microsoft 365 has access to your company data, as long as the user has the correct privileges. This can not only save you time, but may also draw in information you would otherwise not have noticed; perhaps drawn from a meeting on the topic that you didn't attend.

However, you can't 100 percent rely on Copilot, as Adam discovered when he tried to use AI to create a PowerPoint presentation for a military client. At first, it looked great: some carefully selected prompts, using DALL-E to create the images, produced a compelling set of slides. "Then I decided against it because it had come out with so many incorrect statements and so many pictures that were obviously AI generated. I thought, 'If I take that into a high-level meeting with a general and they reveal me to be using AI, I could be in serious trouble here.'"

This flags up two things. One is that Copilot for Microsoft 365 is much more powerful than Copilot Pro—yes, this can also create presentations automatically, but only from data drawn in from the web or the document you're working on—but the second is that it brings danger. In particular, if you fail to label a file as confidential then it could be surfaced by a Copilot query by anyone in your company. It doesn't take much imagination to see how badly that could go wrong. If your business has fewer than 150 employees, then you'll need to rely on your partners' expertise to leap that hurdle.

CREATE YOUR OWN COPILOT GPTS

Businesses using Copilot for Microsoft 365 can create their own AI assistants already, although 'AI assistant' pushes the definition: chatbot is a better word.

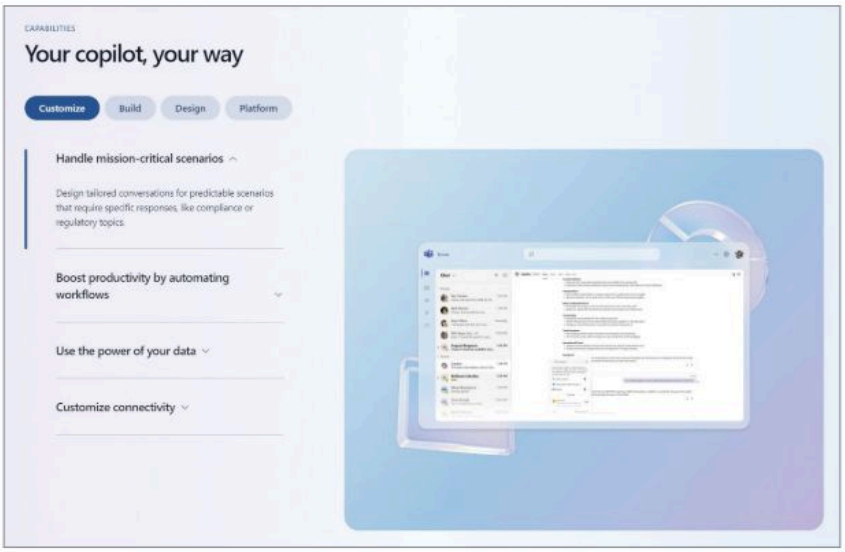

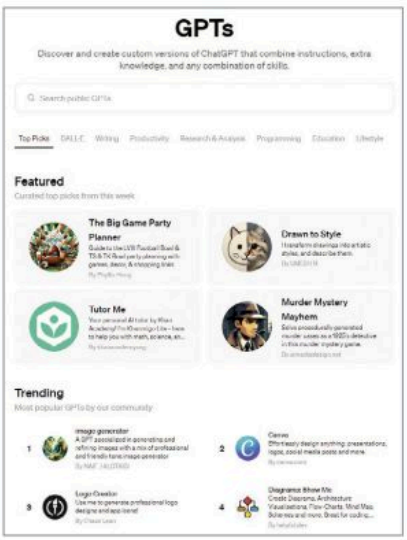

Copilot for Microsoft 365 users can create their own GPTs, but OpenAI has a headstart.

Essentially, the tool directs Microsoft's large language model at your choice of data, so if you sell bathroom fittings then you could point it to a brochure or your website, and the LLM will devour the data and answer questions on it. In this case, it's like a glorified user manual that you can chat to.

In relation to OpenAI, the two tools are essentially the same, but in different clothing. We're not criticizing it, though, as GPTs have the potential to be superb specialized personal helpers. You can also fine-tune them during and after their creation, as you spot things you don't like or if you want to add extra skills or sources.

In businesses, there are some immediate uses—employees could ask GPTs about HR information, background on clients, or a who's who of your organization. If you head to OpenAI's GPT Store (chat.openai.com/gpts) then you'll see the kind of things that have already been created by communities and companies. For example, there's a Wolfram GPT, a PowerPoint slide-making tool, and tutors in numerous topics.

For now, only Copilot for Office 365 subscribers can create their apps in Microsoft Copilot Studio (tinyurl.com/355studio). Copilot Pro subscribers must wait for Copilot GPT Builder. Just like the lack of OneDrive integration, it's frustrating that Microsoft has started to sell Copilot Pro without this facility. All we know is that it's "coming soon".

INTEGRATION WITH OFFICE APPS

Subscribers to both Copilot Pro and Copilot for Microsoft 365 can take advantage of the AI's integration with Office apps. As already mentioned, that means Word, Excel, PowerPoint, OneNote, and Outlook for Copilot Pro users, while business users get the huge benefit of Teams integration.

The strength of AI in those apps often comes in its ability to simplify information. For example, in Outlook, it can make sense of a multi-person conversation that's taken place over days; superb if you've just come back from holiday. In Word, it can condense a 30-page report into a one-page summary, complete with references. In OneNote, it can turn handwritten notes into action points.

Microsoft has done a good job of integrating Copilot into the Office apps. It sits in the Home toolbar, and the sole time it's pushy is when you open a new document in Word and a Copilot prompt appears. The only aggravation is that it's

FREE AND PAID-FOR COPILOT ALTERNATIVES

Ignore the seemingly endless number of companies creating AI models based on GPT—this handful of genuine Copilot alternatives are actually worth considering

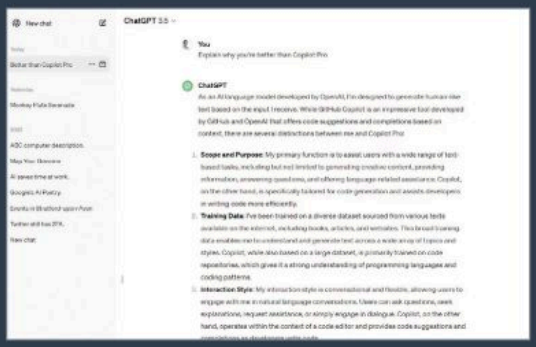

CHATGPT, CHATGPT PLUS, CHATGPT TEAM

PRICE ChatGPT, free; ChatGPT Plus, $20 per month; ChatGPT Team, $30 per month from chat.openai.com

Plain old ChatGPT is starting to look both old and plain, relying as it does on GPT-3.5 and—unless you download the iOS and Android apps—the stark web interface. Nor does it offer any way to create images. Where it wins compared to the free version of Copilot is that you get to keep your history, while there's no limit on queries or interactions.

Things get more interesting with ChatGPT Plus. The interface is the same, but you can create GPTs and connect with third-party services. That means it can do things such as plan holidays and make restaurant bookings. It's also brilliant at analyzing data, making sense of complicated spreadsheets, for example. Like Copilot Pro, it gives you access to DALL-E for image creation—but we still prefer Midjourney.

There's one final point in ChatGPT Plus' favor. At $20 per month, it's cheaper than Copilot Pro. OpenAI is hoping to lure more businesses to its service, too, with a new service called ChatGPT Team. This provides an admin console, the ability to share GPTs within your workspace, and higher caps on GPT-4 and DALL-E.

MICROSOFT COPILOT

PRICE Free from copilot.microsoft.com

There is a free version of Microsoft Copilot, and it's available to everyone via apps on phones and its website, copilot.microsoft.com. There are several key differences to Copilot Pro. First, it doesn't integrate with Office apps, which is, after all, arguably the biggest draw of the service. Second, you can't create Copilot GPTs. Third, during peak times, you'll likely be stuck with GPT 3.5 rather than the more powerful and up-to-date GPT 4. Fourth, image generation is restricted to 15 'boosts' per day, and you may need to wait longer (depending on demand).

GOOGLE GEMINI, GEMINI ADVANCED

PRICE Gemini, free; Gemini Advanced via Google One, $20 per month from gemini.google.com

The basic version of Gemini is free, as was its predecessor, Bard, until Google replaced it in February. This free version of Google's LLM isn't as powerful as GPT-4, but its not-so secret weapon is integration with Google services such as Calendar, Gmail, and Drive, So you can tell it to hunt through all your documents looking for mention of a certain keyword, and then produce a summary.

However, Google hopes that you'll upgrade to Gemini Advanced, which is now part of its One AI Premium offering. This includes what Google describes as its "most capable model, Ultra 1.0", and it certainly offers more nuanced, creative answers than plain Gemini. It's early days, though, with no capability to create images and no integration with Docs or Gmail. Fortunately, Google sweetens the deal with 2TB of storage and a VPN.

WRITESONIC

PRICE Free to $20 per month from writesonic.com

Occasional writers might find that this handy tool does what they need: kick things off. The Article Writer is a great way to turn an idea into—as the name suggests—an article, even if it's likely to be a tad bland. Or you can use Writesonic to generate a structure, to rewrite your copy or simply use it to generate ideas. The free version is restricted to GPT-3.5 and includes 10,000 words, while the Small Team version ($19 per month) gives you a choice of GPT-3.5 or GPT-4 (you get 200K words per month with the former, 33.3K with the latter) and extra tools such as 'brand voices'. If you need an unlimited number of words, but don't mind GPT-3.5, you have the option of paying $20 per month for the Freelancer option.

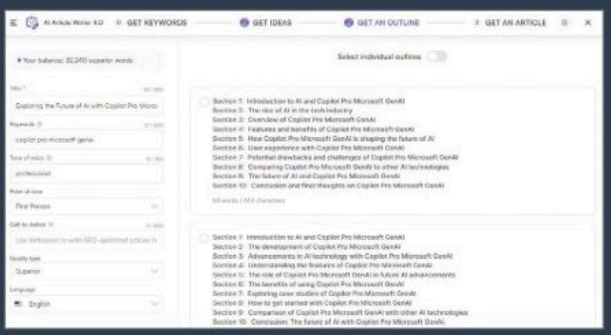

a cloud-based service, so you need an internet connection and will have to wait a few seconds for the results to appear.

Although it's subtle in appearance, Copilot (both Pro and for Microsoft 365) is the strongest reason to keep using Office rather than switching to one of the free alternatives. When OneDrive integration appears, it will be even more powerful.

CREATE IMAGES WITH DALL-E

Although Midjourney remains our favorite image-generation tool, if you don't fancy paying a minimum of $8 per month then Microsoft's Designer (designer.microsoft.com) is well worth visiting.

The nominal difference between Copilot and Copilot Pro is that while non-subscribers get 15 'boosts' per day for free, Copilot Pro users enjoy 100. In our tests, however, the benefits are that paid-for users get three or four images with a 16:9 aspect ratio (1,792 x 1,024), while free users are only offered two images, with a square aspect ratio (1,024 x 1,024).

The results are more than good enough to include in PowerPoint presentations, and Designer's big advantage over Midjourney is ease of use. It takes seconds to create your first image. As with all generative AI, though, the better you get at prompting, the better the results will be.

Where Midjourney excels is refinement. If you like an image, but it's not quite right, then you can iterate by adding an extra parameter. With Designer, you have to rerun the command with slightly altered phrasing. This generates a new set of images, unrelated to what you had before.

It's also harder to create photorealistic images with DALL-E than it is with Midjourney—somehow there's always an AI veneer—and that means your results can look samey, as if they've been drawn by one artist.

While we like that DALL-E comes as part of the paid-for Copilot package, we can't see Midjourney fans switching over.

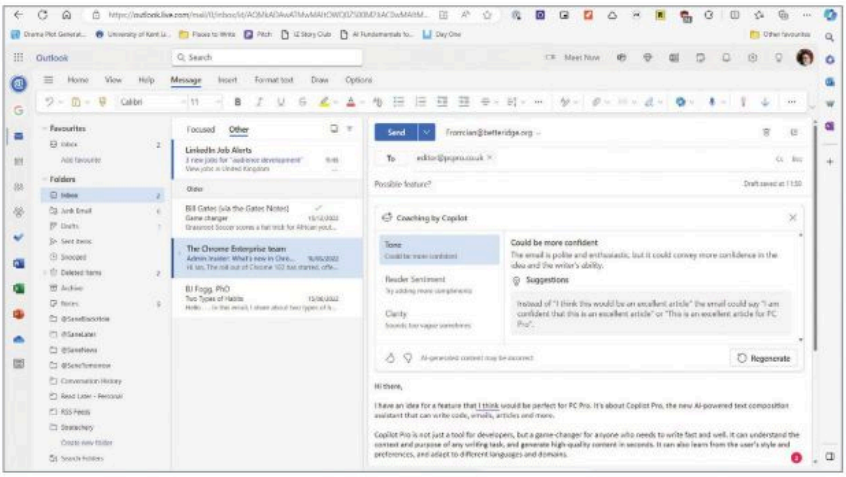

Choose the tone of your AI-generated emails carefully.

MASTERING PROMPTS

We keep talking about mastering prompts, and fortunately Microsoft is there to offer you help. One excellent place to start your Copilot journey is at support.microsoft.com/copilot. This includes several useful resources, including updates on what's new (Microsoft is always adding features), eight short videos showing Copilot in action—although these can be a little folksy—plus extra information that Microsoft thinks could be useful.

One example is a 'toolkit' about prompting, although we'd call it a four-page PDF. This attempts to explain the 'art and science of prompting', including a section that talks about the right prompt 'ingredients'. It suggests that there are four such morsels.

First, the goal. This is the response you want, such as 'generate five bullet points' or 'write a summary'. Then there's context, so why you need it and who is involved. Microsoft's chosen example here is 'prepare me for a meeting with Client X to discuss their 'Phase 3+' brand campaign'. Peak Microsoft.

Third, comes source—that is, what information, sources, or samples Copilot should use. In the example, it's 'focus on email and Teams chats since June', but that only refers to Copilot for 365, not Copilot Pro. For Pro, your source is your document or the web.

Expectations comes last. This is how you want Copilot to respond to best

AN AI PRESENTATION

'Create a six-slide presentation about processors' resulted in this effort, even if Copilot's choice of images is a little random (not unusually)

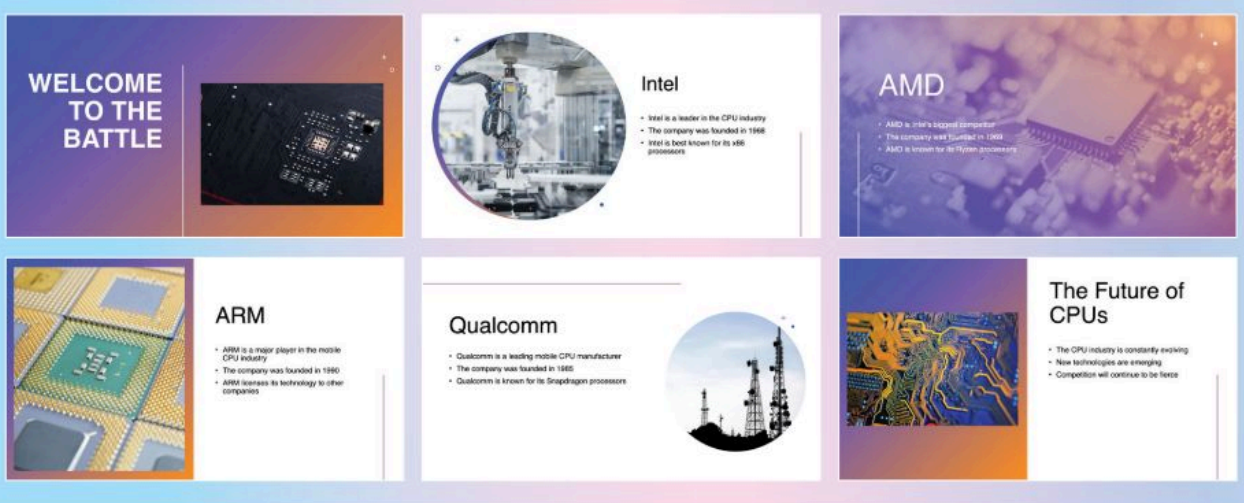

COPILOT PRO FOR WORD

A handy companion that can do the basics, but it won't turn you into a wordsmith

SCORE ★★★★☆

Copilot has three main uses in Word: creating a draft for either an entire document or for sections of it, based on a prompt of up to 2,000 characters; rewriting selected text according to a prompt; and answering questions about a document's content, including summarizing it.

When you create a new document in Word, you'll see how much Microsoft wants you to draft using Copilot—it's the first thing you see. Any time you make a new paragraph, the Copilot icon shows up in the left margin, letting you input a prompt to write a new section. If you select text, the icon shows an option to rewrite the selection. There's also a Copilot icon in the Home toolbar, which opens a sidebar so you can ask questions about the open document, summarize it, or write additional parts.

The output quality of any big language model depends on the prompt you provide. If you aren't specific or clear enough, Copilot Pro will produce dull texts that don't match your voice. You must give sufficient details and context for it to understand your purpose and style, as well as information about the preferences of the audience, all of which can be tricky in a prompt limited to 2,000 characters.

Copilot for Word works best when creating an outline rather than a document.

Occasionally, it's vague, drifts off-topic, or entirely ignores explicit instructions—for example, about word counts or facts you have told it to include. You must check and edit the texts that Copilot Pro writes to make sure they're right for your goal and audience.

Copilot has its own little linguistic ticks, which you will probably need to be explicit about in prompts. It adores bullet-pointed lists, and will include them every time unless you tell it not to. Similarly, and likely reflecting the web content on which it's trained, it loves to include 'hints and tips' sections.

Rewriting is also erratic. Although its grammar is impeccable (with a penchant for Oxford commas), Copilot won't catch and fix all the problems in your paragraphs; think coherence, structure, and flow. It's no replacement for a dedicated tool such as Grammarly or LanguageTool. I even found that Copilot added in things that Microsoft Editor thought were errors.

Copilot can still be handy in Word if you use it for what it's good at. For example, instead of making it write a whole document, ask it to create an outline for you to work. This delivers solid results, and if nothing else, avoids the blank sheet of paper issue.

I also found it handy to keep the sidebar open when working on a long piece; in one case, I needed to make sure that every chapter included a call to action at the end, and that's something Copilot is great at. It's also handy for those moments when you want to get suggestions for something, whether that's as trivial as looking for an antonym or as complex as 'give me five typical Russian male first names, popular in the 1950s'.

Considering this is Microsoft's first stab at Copilot for Word, it's an impressive effort. –IAN BETTERIDGE

meet your expectations, such as 'please use simple language so I can get up to speed quickly'. The full example is: 'Generate 3-5 bullet points to prepare me for a meeting with Client X to discuss their 'Phase 3+' brand campaign. Focus on Email and Teams chats since June. Please use simple language so I can get up to speed quickly.' Note the use of full stops to keep things simple.

There are more tips in the document. One is to start off with broader requests, and then give more details. This makes the most of Copilot's iterative nature. Once you know what you want, the more details you can give. Other tips? Using quotation marks if there's something specific you want Copilot to deal with, and simply being as clear as possible.

It's also worth noting what you shouldn't do. Being vague is likely to be your biggest problem, but you should also avoid jargon or slang words. Finally, remember that Copilot is expecting your follow-up question to relate to what came before. For a new topic, write 'New task'.

But we know the real question you want to ask: how good are the Copilot implementations in the Office apps? That's what the rest of this article will answer.

COPILOT PRO FOR POWERPOINT

There's little sign of actual intelligence here, but it's useful for improving existing documents

SCORE ★★★☆☆

If you want an example of what not to use Copilot for in PowerPoint, I have one here.

You might think that asking Copilot Pro to 'create a presentation about Copilot Pro in PowerPoint' would be an easy win for the AI. You would think wrong. Instead, it came up with a presentation about a fictitious piece of software that appears to combine advanced aircraft flight planning with a collision-avoidance system, which it illustrated with a picture of a car.

Of all the Office applications, PowerPoint probably benefits least from Copilot Pro when it comes to creating documents from scratch. This is down to it not having one of the key features that business users of Copilot get: the ability to take a Word document and turn it into a PowerPoint deck. PowerPoint is, fundamentally, an application for taking information and presenting it in a visual format, and a 2,000-character prompt just isn't enough for anything except the shortest presentations.

What Copilot is good for, though, is taking an existing deck and improving it. If you don't like a visual being on the left-hand side of a slide, you can just ask Copilot to move it to the right, although you also need to tell it to move whatever is already on the right to the left, too, if you don't want a visual pile-up. Similarly, asking it to change the headline font on all slides to, say, Constantia—and it makes the changes in seconds—feels magical.

Microsoft includes a set of pre-made prompts to get you started. For example, it can scan a deck for deadlines and list them, helping you avoid that moment

when you realize that you've combined two people's work and got different deadlines for a project. Stick to the preset prompts at first, and Copilot Pro in PowerPoint is a useful tool.–IAN BETTERIDGE

COPILOT PRO FOR EXCEL

The surprise hit of the package, Copilot Pro for Excel is a great tool for less savvy users

SCORE ★★★★☆

Unlike the other Copilot Pro tools, Copilot for Excel is labelled prominently as 'beta'. But even in this state, it has the promise of being a game-changer for anyone who needs to work with data, but doesn't want to become an expert in writing formulas, working out the best way to pivot data or spotting trends in large data sets.

Copilot for Excel exists in the toolbar, but sometimes it's grayed out. That's because it only works on .xlsx or .xlsm files saved in OneDrive or SharePoint. When the button is green, hitting it allows you to write natural language instructions to create formulas, analyze data, or highlight cells according to whatever criteria you want. If you've ever struggled with creating a complex formula (or even a simple one), you're going to love it.

The first thing to note is that Copilot only works within tables; if your data is unstructured, it won't let you do anything with it. This probably isn't a big deal for most people (every serious Excel user I have known has lived or died by tables), but it's a limitation that may affect some users of large and disparate data sets.

If your data is in a table, Copilot is miraculous. You can ask simple queries, such as how much you spent between two dates, or what categories you spend the most on. It's also great at creating graphs: just type, 'make a pie chart showing expenditure based on month'. You can manipulate data by asking it to use conditional formatting to highlight ranges, something that's easy to get wrong for people who aren't Excel experts.

You can also use Copilot to look for outliers in the data and highlight them, which helps if you're trying to clean up a dataset or are finding that the results aren't what you think they should be. If you're using a big dataset, that's a time saver. It's nothing that you can't do with Excel's existing tools, but for users who don't live in Excel, it would involve quite a bit of searching to know where to start.

The real power kicks in when you ask it to make formulas for you. I created a table of expenditure for a small business, tracking spending across categories, but I also wanted to have a column indicating the running total of my expenditure. This isn't easy to make if you don't know much about Excel, because it involves a SUMIF function that's based on the row's date, comparing it to others in the table.

I asked Copilot to 'add formula columns to summarize total expenditure so far in this date sequence', and it created a formula and added it to the table.

This kind of formula creation is going to save users of Excel hours. If you regularly work with data, but wouldn't call yourself an expert, it's probably worth the money on its own. Even on occasions where I wanted to do something I knew how to do, I found myself using Copilot instead, because it created better results than the various hacks, shortcuts, and cheats I'd learned. –IAN BETTERIDGE

DOES YOUR BUSINESS QUALIFY FOR FREE FASTTRACK SUPPORT?

When Microsoft announced the general availability of Copilot for Microsoft 365 in January, it also removed all limitations on rollout sizes. Up until that point, you needed to commit to 300 seats for a minimum of a year. At $30 exc VAT per user per month—on top of the cost of Office 365 itself, which is a minimum of $23 exc VAT per user per month for the E3 version—that added up to a $108,000 investment over the course of a year. Minimum.

The sweetener to that deal is that companies received direct support from Microsoft via its FastTrack experts. Now, that offer extends to companies with 150 paid-for Office 365 (or Microsoft 365) licences—if, that is, you fulfil four other criteria.

The first is easy: you must already use the Office apps. The second is that you are 'Teams ready', which means that you share content in meetings (such as Word docs) and hold meetings with at least three people in attendance.

The third is trickier: you must commit to monthly channel refreshes of Office software. Many businesses prefer the six-monthly refresh, as this gives them time to test line-of-business software. Switching to a monthly refresh comes with extra hassle and potentially higher costs.

Microsoft's final criterion revolves around labeling. You must either have an existing MIP labeling policy (MIP stands for Microsoft Information Protection and relates to labelling around Azure data) or have your own labeling policies in place. If you fulfil those criteria, get in touch with Microsoft directly.

'Teams-ready' businesses may be able to access free support.

COPILOT PRO FOR OUTLOOK

Of greatest use to people who are sent long, rambling emails or struggle to compose quick replies

SCORE ★★★☆☆

There are three key features in Copilot for Outlook: summarization, drafting, and coaching. Summarization is probably the feature you'll encounter first, as every email you receive has a prominent 'Summary by Copilot' bar at the top.

Click on this, and it creates a summary of the key points in the email. How useful this will be depends on the kind of mail you get, but if you spend a long time reading complex emails and trying to work out what the point is, you'll love it.

The second main feature, drafting, is like the drafting feature in Word—you give it a prompt, and it writes the email for you. You can vary the tone using pop-up options—direct, neutral, casual, formal or, erm, 'make it a poem'—and set the length as short, medium, or long. Outlook uses the last tone you selected, so if you decide to write a poem, change it before drafting an email to your accountant.

As with Word, I'd categorize the results as 'something to start with and

COPILOT FOR TEAMS

A brilliant personal assistant for meetings, but don't let it replace your critical thinking skills

SCORE ★★★★★

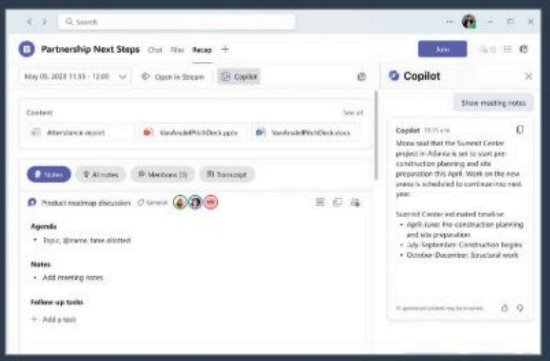

Copilot for Teams is a genuinely beneficial use of AI in the workplace.

Long before the arrival of Copilot, Microsoft Teams became my go-to app. It's the tool I use to stay updated about my organization's activities, and where I plan my daily tasks, share content, have meetings, and plan goals. Now Microsoft has added Copilot, it has become even more indispensable.

This is most obvious in meetings. The responsibility of capturing important meeting details typically rests on a single person, a task that's rarely enviable and often challenging. Even with transcription features enabled in Teams, sifting through conversations to find pertinent information was daunting. This is an area where Copilot excels, making it easy to summarize the discussion and action points.

I also find it useful within meetings, especially long ones that drag on, when your attention can wander. If I briefly zone out, or I need to head off for any reason, Copilot can offer summaries at any juncture. If you arrive late, or simply don't get to the meeting at all, being able to ask Copilot for a summary of the meeting with the key action points and follow-up tasks is a huge boon, as is its ability to revisit discussions and extract insights.

This isn't just theory. Over the past six months, I've asked Copilot to pinpoint staff disagreements, outline risks discussed, verify project timelines, and summarize crucial decisions. Using Copilot for these kinds of enquiries saves time by reducing the need to consult busy colleagues, delivering on two fundamental promises of AI: it streamlines communication and enhances my efficiency.

Copilot packs in several other useful features. Managers may love (a little too much, perhaps) its ability to analyze participation in meetings by generating metrics showing the percentage of words spoken by each colleague. One standout feature in Teams Premium, not available in the standard version, is the intelligent meeting recap. It parallels Copilot's meeting notes and holds the information for 30 days. The 'follow meeting' function in Teams Premium allows users to capture meetings, pose questions to Copilot post-event, and even analyze the sentiment, although this yields mixed results and is a little creepy.

Copilot in Teams can also save you time. It adeptly transforms meeting notes into Office 365 documents and drafts emails from discussions. It also uses Microsoft Graph to link and track communications across documents, meetings and chats on specific subjects. This is great if your memory isn't too good, so this is my favorite feature.

As Microsoft states, however, "Copilot is an assistant, not a replacement for human oversight". It's crucial that you review Copilot's output, particularly when it pertains to significant decisions, to ensure accuracy and relevance. So while Copilot boosts efficiency, there is a risk that users may become too dependent on the tool, which could lead to a decline in critical thinking and problem-solving abilities.

Finally, in organizations where security is paramount, outputs generated by Copilot don't automatically carry over the security classifications of the original files. This poses a risk when dealing with sensitive information. Consequently, it becomes the responsibility of the employee to review the AI-generated content to confirm the data is correctly classified and evaluated for potential risks.

To misquote a certain masked superhero, with great AI power comes great responsibility to check its results. –ADAM TIMBERLEY

personalize' rather than the finished article. Tonally, it veers towards the extreme: formal is very formal, and casual is fine for sending to family and friends. In my tests, direct produced the best results, although some of its phrases required toning down to stop them sounding like the kind of email you get from the bailiffs chasing you for a late payment.

My favorite feature is coaching. This checks the content of an email and gives you tips on how to improve it, with clear advice that's actually useful. It advised me to make my tone more confident, which professional writing coaches have told me in the past. –IAN BETTERIDGE

COPILOT PRO FOR ONENOTE

Don't expect anything radical, but this is a handy enhancement for inveterate note-takers

SCORE ★★★☆☆

It could be because I tried Copilot in OneNote after all the other apps. It could be a because I'm not a regular OneNote user. Or it could be that this is a, 'Oh, I suppose we should add it to OneNote as well' kind of effort.

Copilot appears in OneNote the same way as most other apps: as a button in the Home ribbon. You can do all the things you'd expect, such as summarizing meeting notes and creating action points. But where in Teams it feels powerful and integrated, here it feels like an add-on.

For example, I hoped that it would tie in with the transcription option, perhaps creating a rival to Otter.ai that not only delivered the notes but gave me something extra. But no, it couldn't even remove the timestamps from the text when I accidentally kept them in.

You'll also need to be careful about where you file notes, as Copilot has some difficulty discerning between the active note and all the ones filed within a project. Mind you, the fact that it can work across a bunch of collated notes has its advantages, too.

Copilot in OneNote performed best when working with my scrawled notes from a meeting, deducing that when I wrote some names next to "Who?" that these were indeed the people involved.

If you're a OneNote fan then you'll probably like what Copilot Pro can do. But I can't see this making any new converts. –TIM DANTON

CENTERFOLD

PERFORMANCE GEAR LAID BARE

1 WIND UP THE WATTAGE

One of the more striking features is a high-performance 54W mode for its AMD silicon. Most equivalent handhelds are limited to 30W TDP. Arguably, the 54W mode isn't a realistic goer on battery power, but it should squeeze out a few more frames when plugged into the wall or a dock.

2 TRACKPAD TIMES TWO

Ayaneo has decided that the Steam Deck's dual trackpads are a good idea. Certainly, they work similarly to conventional trackpads on laptops and help with navigating not only Windows, but game UIs. The latter all seem to be designed with the assumption that absolutely everyone is going to be using a mouse.

3 WHAT A VIEW

The monster 8.4-inch 2,560 by 1,600 screen is the Ayaneo Kun's centerpiece. It makes for a beast of a handheld PC, but also increases gaming immersion. The only problem is that the overall scale has you wondering if a gaming laptop isn't a better bet.

Ayaneo Kun

KING of the handheld gaming PCs? That seems like the plan for the Ayaneo Kun. Sure, it has the same AMD Ryzen 7 7840U silicon as a million other gaming handhelds. Ayaneo itself has a host of quirky handheld devices, from the tiny Air 1S, to the Flip in two designs, and the Slide.

The Kun is the most traditional of the lot, but has its own claims to innovation, the unique 54W mode on its 30W APU being one of them, plus a honking great 8.4-inch screen. That makes for a big beast, but it feels good in the hand in a Steam Deck-y kind of way.

It's also the only other handheld we've seen so far to pick up Valve's use of twin trackpads and run with it. The size of the Kun means it can get away with sticking two pads, one below the left D-pad and the other under the right-hand thumbstick. When you're using a Windows-based handheld, they come in handy for ease of navigation in those awkward launchers that resolutely adhere to the idea that if you're playing a PC game, you're going to be using a mouse as your primary input.

The core spec itself is impressive, too. Alongside the ubiquitous Ryzen 7 7840U and its 780M iGPU, there's 32GB of LPDDR5 and a 2TB Lexar SSD. Ayaneo has jammed a hefty battery into the Kun, too, which bodes well for longevity on the move. –MAXIMUM PC

RETRO GAMING ON WINDOWS 11

Want to enjoy classic games on your modern PC? *Nate Drake* can help get you playing

EVERY MODERN GAMER knows that it's hard to replicate the excitement we all felt as kids when we started playing videogames. True, many of these had rudimentary graphics and may have even had no network play, but who can forget the peculiar feeling of satisfaction that came with facing down the alien menace in *Half-Life*, armed only with a crowbar? What modern game can even come close to the joy of 'insult sword fighting' in *The Secret of Monkey Island*?

As modern hardware evolves and 64-bit systems become more commonplace, it's actually getting harder to run titles like these out of the box. Modern hardware often lacks optical disk drives, and original game cartridges for older consoles like the NES and SNES are selling for ridiculous amounts.

In this guide, we'll explore some of the ways in which having a modern PC is no bar to playing your favorite retro games. We'll explore simple ways to feel nostalgic by playing titles online via the Internet Archive or dedicated platforms like GOG and Steam. We'll also explore emulators like DOSBox and RetroArch to allow direct play on your device.


SIGN UP | LOG IN
UPLOAD
Search
VOLUNTEER
PEOPLE
CONTACT
JOBS
DONATE
HELP
Favorite
Share
Flag
acer
NZXT

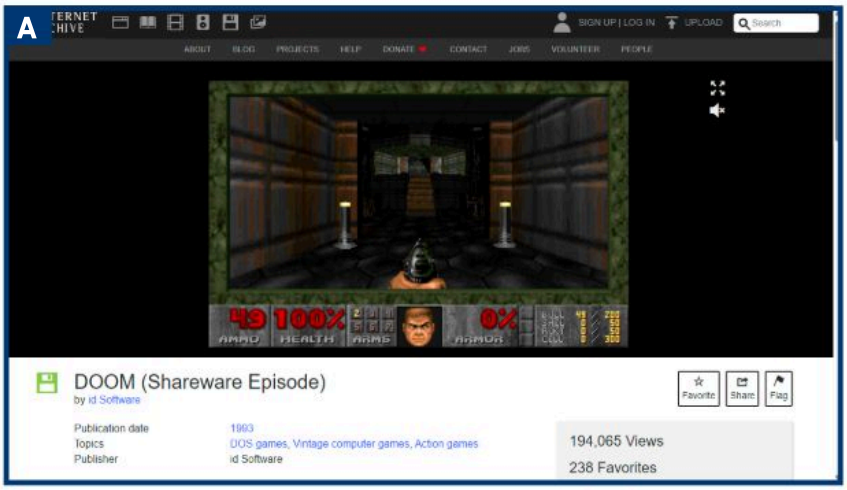

ONE FROM THE VAULTS

One of the simplest ways to relive the glory days of retro gaming is to open your browser and head over to the Internet Archive (https://archive.org). Simply clicking on the 'Software' category will automatically show links to the Archive's own 'Internet Arcade', as well as 'Classic PC Games' and MS-DOS games.

If you do select any of these, make sure you mark the checkbox 'Software' in the 'Media Type' section to display only videogames and not related media like screenshots.

Once you've selected your chosen title [**Image A**], click anywhere on the screen area to fire up the Internet Archive's browser implementation of DOSBox (see below), and play the game via the website.

From our research, the games seem to play faithfully, though it's always advisable to scroll down to the comments section to see if someone has listed which particular keys are used for gameplay. If this isn't present, then run a search in the field at the top right for the game manual, eg. 'Wolfenstein Manual', to see if the documentation has been uploaded.

Using the Internet Archive to play retro games is generally the quickest and safest way, particularly given that it was officially registered as a library in the state of California in 2007. This has led some retro gaming redditors to conclude that playing copyrighted games on the main website is legal, as it amounts to the same thing legally as borrowing a book. Take your own legal advice before deciding to play.

The Archive also plays host to a 'Vintage Software' collection (https://archive.org/details/vintagesoftware), which is chiefly 'Abandonware'—software that is no longer current, and whose owner either no longer supports it or no longer exists. The legality of playing such games exists in a gray area, so make sure once again to check copyright law where you are.

As the Internet Archive blog notes, if you choose a title that was made available on CD-ROM, the 700+ MB file can take some time to load into the online emulator.

As you're playing the game in-browser, certain features, like saving or multiplayer, also won't function. If you need these extras, scroll down the Internet Archive page to 'Download Options' to access the original game files. Be warned that they may require an emulator to run correctly in modern versions of Windows.

If you don't want to grapple with the command line and emulator settings just yet, Good Old Games (www.gog.com) has an excellent selection of DRM-free classic games, including abandonware [**Image B**].

The team at GOG go to great lengths to track down the original rights holders to titles, so it's very unlikely that you'll have any copyright issues for games you purchase. GOG developers also 'pre-patch' and restore older titles, such as *System Shock 2*, using the game's original source code, making it possible to run on Windows without any compatibility issues. If customizing or reverse engineering a game's source code isn't possible, it's usually sold prepackaged with an emulator, allowing it to run automatically upon install.

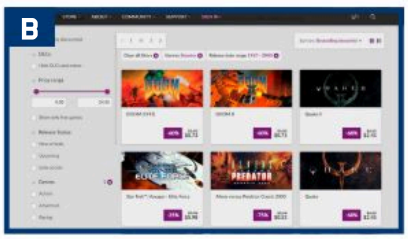

Though they don't quite have the same reputation for lovingly restoring vintage games, there are also a number of retro titles available in Steam, such as *Deus Ex*. To play these, you'll need a Steam account and also to have downloaded the Steam Client. As of January 2024, this means your PC will need to be running at least Windows 10.

Whether you use the Archive to download games or purchase via GOG/

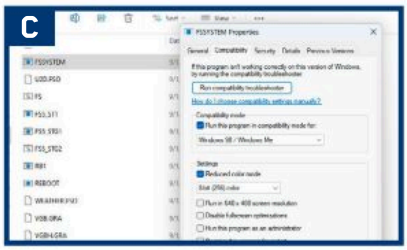

TAKING CONTROL

For a truly authentic gaming experience, nothing beats the tactile sensation of a controller.

If you're a purist, websites like Retro USB (www.retrousb.com) have USB adapters for old-style NES, SNES, N64, and Genesis controllers, meaning you can add old-style gaming ports to your computer.

This can be useful if you have a particularly exotic type of controller you'd like to use with a game, such as the NES Powerglove. Be warned, though, that not all adapters work with all accessories. In the case of RetroUSB, for instance, the 'Super Retro Port' works perfectly with SNES controllers, but will not function with peripherals like the Super Scope.

A less expensive and more convenient option is to scour eBay and Amazon for USB-equivalent controllers. For instance, during our research, we discovered Saffun sellling a two-pack of SNES-style controllers on Amazon that work wirelessly for under $29.

If you enjoy playing titles across multiple consoles, it can get expensive to buy equivalent controllers for each one. Fortunately, some modern retro controllers, such as the 8BitDo Pro 2 (www.8bitdo.com/pro2), are designed to allow play on multiple systems, as they support multiple button configuration profiles.

Steam, you may still find compatibility issues when attempting to launch on Windows 11. You can sometimes bypass this by running programs in compatibility mode [**Image C**].

To get started, right-click your chosen program and choose 'Properties'. Next, click on the 'Compatibility' tab. You'll notice a button marked 'Run compatibility troubleshooter'. In theory, clicking this will launch a wizard that automatically detects and configures compatibility settings for you.

In our tests, however, we found this to be useless. From here, though, you can check the box marked 'Run this program in compatibility mode for', then select your chosen OS from the drop-down menu, eg. 'Windows 98/Windows ME'.

You'll also notice that there are other options, such as 'Reduced Color Mode' (you can choose between 8-Bit and 16-Bit), as well as the option to run the game in 640 x 480 resolution. Use trial and error with different settings to determine what works best for your chosen title.

UNPACK DOSBOX

If you want to run classic PC games on your own device, DOSBox takes the gold medal. The emulator name is something of a misnomer, as while it can emulate DOS systems easily enough, it can also run almost all Windows 3.0 and Windows 9x games as well.

Before getting started, make sure that your PC can access the actual game files. As DOSBox involves working with the command line, we advise keeping things simple by creating a 'GAMES' folder in your root system drive, then placing the game folder inside it, eg. 'C:\ GAMES\QUAKE'.

If your chosen game files are bundled in an ISO image, Windows 11 has a built-in 'mount' option, which you can access simply by right-clicking. Make a note of the virtual drive letter.

Next, go to www.dosbox.com, and click on 'Downloads' to access the installer via SourceForge. Launch and choose 'Next', then 'Install' to continue. You can now run DOSBox via the desktop shortcut.

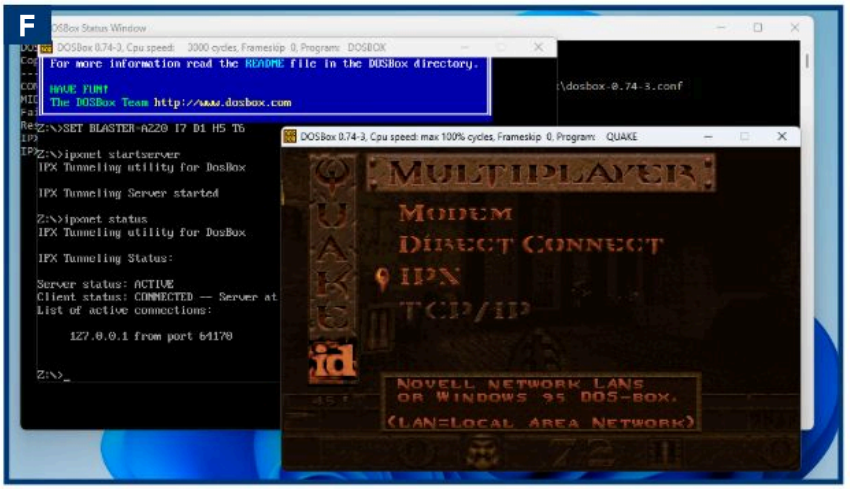

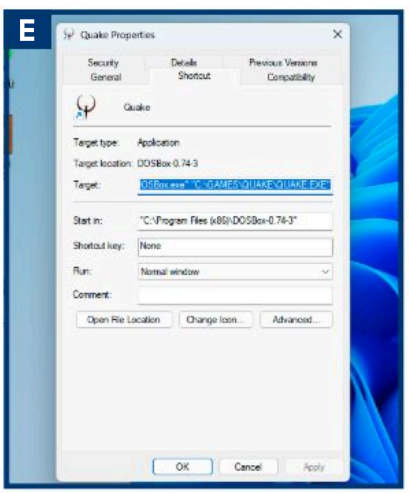

To get started, you first need to set your 'GAMES' directory to be mounted as the 'C' drive in DOSBox via the 'mount' command, eg: [**Image D**]

```
MOUNT C C:\ GAMES\
```

Next, enter

```
C:
```

to move to your new virtual drive. From here, you can use 'CD' to switch to your specific game directory, eg:

```
CD QUAKE
```

Use the 'DIR' command to remind yourself what files are installed here. You can then run the relevant executable simply by entering its name, eg:

```
QUAKE
```

Mounting virtual drives and searching folders manually each time you want to play can get very tiresome. Fortunately, once you understand how DOSBox works, there's a much easier way [**Image E**].

First, exit DOSBox, then open File Explorer. Navigate to 'Program Files' and find the exact name and location of the program, eg:

```
C:\Program Files (x86)\DOSBox-0.74-3\
DOSBox.exe
```

Next, navigate to your chosen game folder, and do the same for the launcher, eg:

```
C:\GAMES\QUAKE\QUAKE.EXE
```

Armed with this information, right-click anywhere on your desktop (or folder of choice), and select 'New' > 'Shortcut'.

In the field marked 'Type the location of the item', enter the path DOSBox in quote marks. Leave a space, then do the same for the game executable, eg:

```
"C:\Program Files (x86)\DOSBox-0.74-3\
DOSBox.exe" "C:\GAMES\QUAKE\
QUAKE.EXE"
```

Click 'Next' to continue. Type a suitable name for your shortcut, eg. 'Quake', then choose 'Finish'.

You can now use your shortcut to automatically launch your chosen game

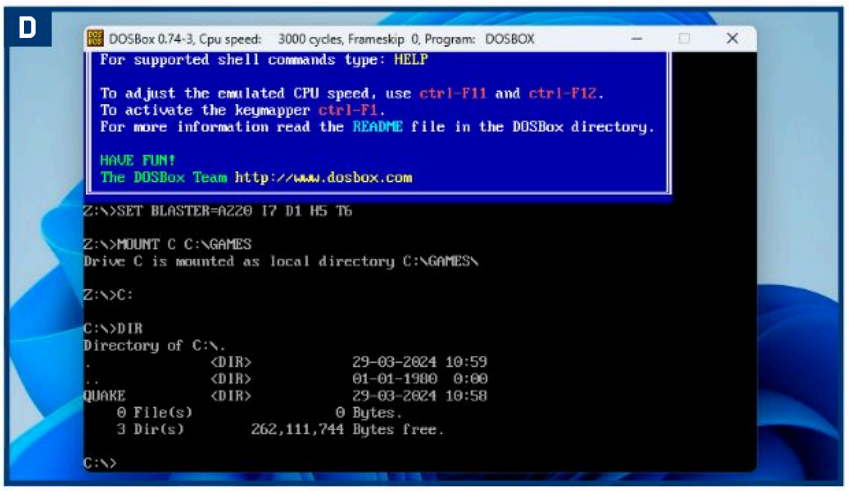

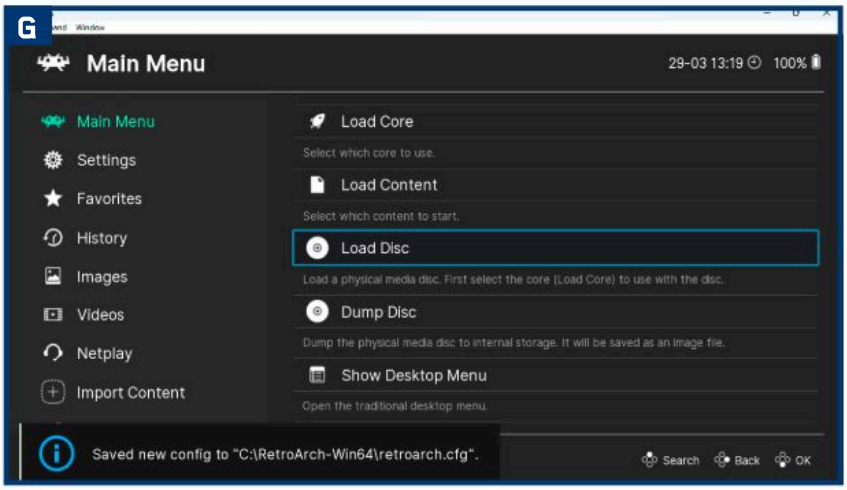

using DOSBox without going via the command line. By default, the shortcut will use the DOSBox icon, but you can change this to one in the game folder via right-click > 'Properties' > 'Change Icon'.

When browsing the DOSBox folder in 'Program Files', you'll likely have seen the batch file named 'DOSBox Options'. Launch this to open 'dosbox.conf', which allows you to alter game settings. You can find a full rundown of these in the DOSBox Wiki (www.dosbox.com/wiki/Dosbox.conf), but some options, such as fullscreen=false, are self-explanatory.

If you do open the .conf file, make sure to change 'ipx=false' to 'ipx=true'. This will allow DOSBox to emulate this now-defunct protocol to allow network/internet play with other users.

If you do want to go head to head with fellow retro gamers, after enabling IPX, simply launch DOSBox.

Start the IPX Server with:

```
Ipxnet startserver
```

By default, the server starts on UDP port 213, but you can manually specify another if you prefer, eg:

```
ipxnet startserver 19900
```

When Windows' firewall prompts you, click 'Allow' to authorize connections.

If you're working behind a router and want to play over the internet, make sure your chosen port is also open and correctly forwarded.

You can double check [**Image F**] that the IPX server is running correctly via:

```
IPXNET STATUS
```

Next, have the other players start their versions of DOSBox. They then need to run the following command to connect to your server:

```
ipxnet connect <IP> [UDP port]
```

eg:

```
ipxnet connect 192.168.2.103 19900
```

You can now launch a multiplayer match via the game's own settings.

CONSOLE CAPERS

While the DOSBox emulator is excellent for reliving old DOS, Win 3.x, and Win 9x titles, it can't natively run videogames designed for home consoles.

If this is your preference, then there are thousands of various emulators available, though our recommended choice is RetroArch [**Image G**]. Besides being free and open-source, RetroArch uses the Libretro API to run various 'cores' in a very efficient way, meaning that it's a capable frontend for a huge number of console emulators.

The software itself doesn't come with any games. If you have access to the original game CD or DVD then this won't pose a problem, as RetroArch can perform a 'dump' of all data to a local hard disk. If you want to play cartridge games, you may need to use specialist hardware to do this, like the Kazzo 'INL Retro' dumper programmer (www.infiniteneslives.com/hardware.php).

Once you have your game files, head to https://retroarch.com and choose 'Download' to access the installer. This is a straightforward process, but if you see the alert saying that DirectX 9.0c runtime isn't installed, select it from the optional components before clicking 'Next'.

On first launch, if you've already mounted your game CD, scroll down to 'Dump Disc' to make a copy of its contents.

Next, scroll down to 'Online Updater'. Select this, then 'Core Downloader' to view a list of systems that RetroArch can emulate. Scroll down to your chosen system, such as Nintendo - NES / Famicom (Nestopia UE). In the case of consoles like the NES, you will see the same system listed multiple times, as there's more than one popular emulator for it.

Click the RetroArch icon at the top right to return to the previous menu, then

TO CRT OR NOT CRT?

Die-hard retro gaming fans sometimes complain that one downside to playing retro games on modern hardware is that the look and feel is slightly off due to the rise of LCD and HD monitors over bulky CRT models.

In our opinion, the difference is negligible and barely alters gameplay. Some frontends, like RetroArch, offer 'scalers', which can add CRT-like effects to games like scanlines.

It's also true that coders in the '80s and '90s sometimes made use of blurry CRT effects for in-game graphics. Lightguns also don't work with modern monitors, so you'll need a CRT if you want to play classics like *Duck Hunt* or *Time Crisis*.

If you're determined to go down this route, we strongly recommend placing an ad in a newspaper or on Facebook Marketplace asking for a CRT to save on the huge postage costs of shipping it via online marketplaces like eBay.

Depending on the age of your PC, you may need to invest in an HDMI to composite adapter (VGA or DVI) to make sure your graphics card can output to the CRT monitor.

Older monitors have a 4:3 aspect ratio, so you most likely will need to make changes in your chosen emulator to have games display correctly on the screen.

RETROARCH TEARDOWN

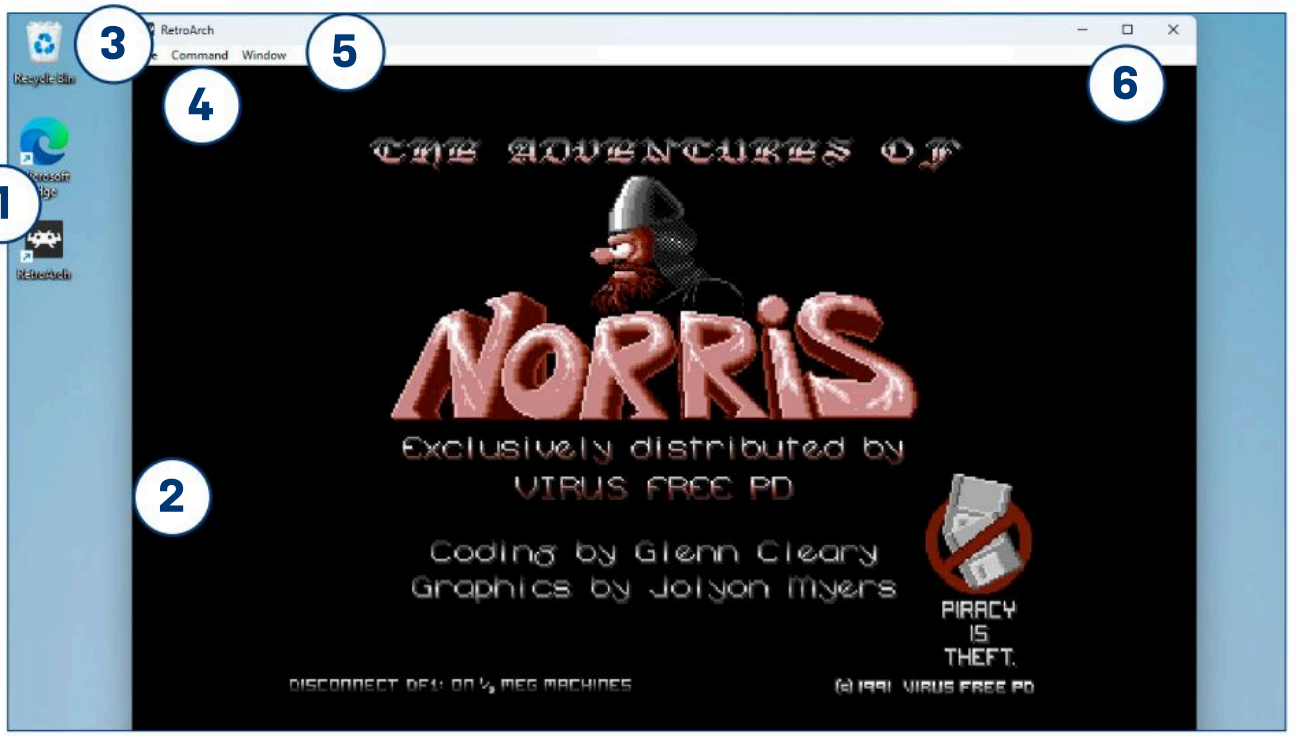

1. RETROARCH LAUNCHER
RetroArch serves as a frontend to multiple emulators. While running, you can download multiple 'cores' to run titles from various systems. These will then be accessible via the main menu.

2. GAME WINDOW
When 'content', i.e. a game, is loaded, it will play automatically in the main window. There are also various hotkeys, such as F2 and F4, to save and load game states.

3. FILE MENU
RetroArch users can access new cores and content from the main menu. Once these are installed, however, the 'File' menu can be used to load a new 'core' (system) or 'content' (game).

4. COMMAND MENU
This menu contains basic options such as pausing and restarting the game. From here, users can also configure audio and disk options, as well as manage save states.

5. WINDOW MENU
The Window menu can be used to adjust the window scale, as well as launch the 'Desktop Menu'. This offers an excellent overview of game playlists, history, and any saved images and videos.

6. WINDOW MANAGEMENT
Retroarch windows can be maximized, but won't automatically scale unless users have configured it so in the Windows menu. You can also toggle fullscreen using Alt + Enter.

choose 'Update installed cores' to make sure that you have the very latest version of the emulator.

Click the icon again to return to the main menu. From here, you can choose 'Import Content' from the left-hand pane. The easiest way to do this is to have your game files in one location, eg. 'C:\GAMES\NES'. If so, select 'Scan Directory to navigate to your chosen folder, then click 'Scan this Directory'.

You can now return to the main menu to 'Load Content' [Image H]. During gameplay, click into the 'File' menu at the top left to load another 'core' (console) or content 'game'.

Take some time to explore the 'Command' menu, particularly 'Save State Options'. During gameplay, you can use F2 and F4 to save and load your progress

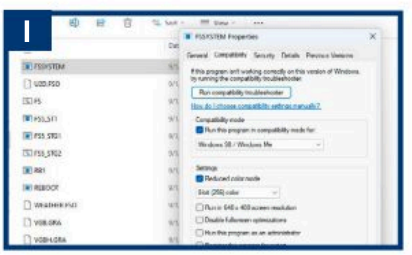

respectively. However, if you want to store your progress in different slots, use the F6 and F7 keys to move between them before saving or loading.

You can also press F1 during gameplay to return to pull up the RetroArch menu. From here, you can start recording or streaming games. You can also select 'Cheats' to explore memory addresses, but this requires programming knowledge.

If you need to fine-tune your game settings, exit the current game and return to the RetroArch main menu. First, choose 'Settings' from the left-hand pane, then 'Video'. Select 'Output' to change configurations like your video driver, resolution, and screen orientation.

The 'video' menu also allows you to configure both 'fullscreen' and 'window' settings, such as width, height, and opacity. To remove unwanted changes, choose 'Reset to Default Controls' [Image I].

After you've configured video output, return to the settings menu and select 'Input' > 'RetroPad Binds'. If you select the relevant controller, eg. 'Port 1', you can then see which particular keys are mapped to buttons. Click any of these, and press a key to map controls in a different way.

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

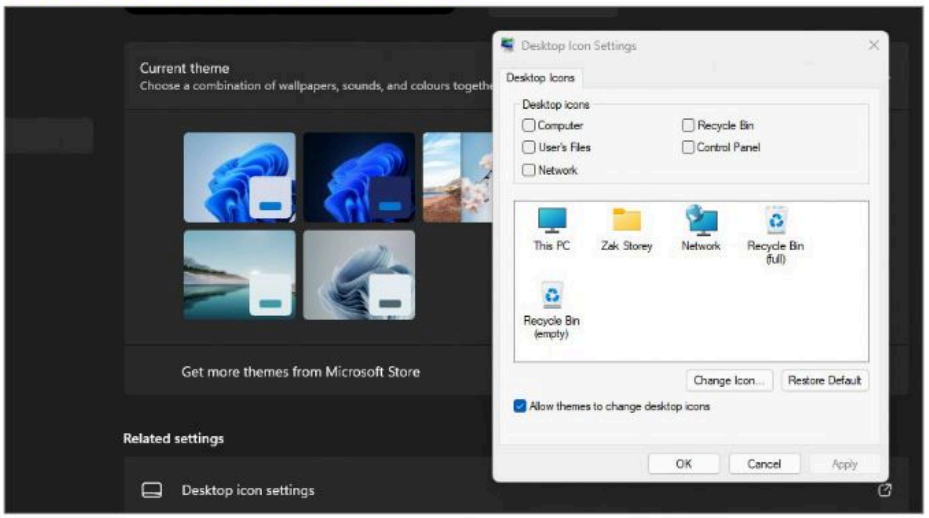

DESKTOP ICONS

Ever needed to get rid of your fixed desktop icons, or add them back in? Did you download Minibin because you just hate that Recycling Bin icon on your desktop? Click the Start menu, type 'Themes and Related Settings', then click 'Desktop icon settings'. From there, you can enable or disable whichever desktop icons you want.

MAKE – USE – CREATE

60
Optimize your PC for gaming

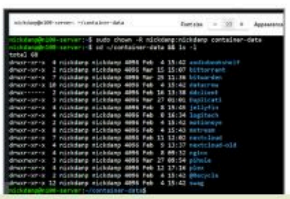

64
Migrate from Docker to Podman

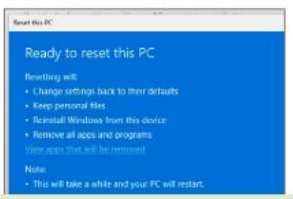

68
Debloat Windows 11

ZAK STOREY
CONTRIBUTOR

KEEPING PACE

Lately, I've been thinking about how we use our hardware, and how certain categories of components are becoming almost superfluous in their performance improvements.

Even if you 3D render, edit videos, or do anything related to data analysis, the majority of your time is still going to be spent not doing those heavy-load things. For the bulk of our operations, the current speed and performance we have with the best SSDs and DDR is fairly superfluous.

A SATA SSD or DDR4 will—at least day-to-day—perform fairly identically to the best PCIe 5.0 M.2 and 8,000 MHz DDR5. Even in gaming, we're starting to see this power creep—CPU, DDR, and SSDs don't affect much. I do benchmarks of SSDs that utilize game load times, and differentiate between a good or bad SSD via millisecond load times. How ridiculous is that?

We're getting to a point where technology is so fast, it's hard to see any benefit. Okay, you will see time saved and better results, but outside of fps with GPUs, there's really no single component that can alter our engagement with the PC. It's like owning a supercar; yeah, you can go fast, but most of the time you still have to deal with the speed limit, and the person in the Toyota can still get from A to B in the same amount of time.

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

Optimize your PC for Gaming

YOU'LL NEED THIS

Stable internet connection
(Optional) gaming mouse

THERE'S NO BETTER WAY to experience gaming than on a dedicated PC with a high-end graphics card, liquid cooling system and RAM to spare. Still, even the very best hardware can have issues running the latest games. You may find that textures don't live up to the promise of those 4K YouTube previews. Gameplay itself sometimes can be glitchy or the controls don't respond in the way you want them to. Such bugs can often be patched simply by updating the game itself. Even when a game seems to be operating at peak performance, though, you want to know that it and your operating system are making the very best use of the hardware available.

In this guide, you'll discover how to tweak Windows 11's settings to optimize your gaming experience. Some of these steps are straightforward, such as activating Windows' 'Game Mode' and enabling AutoHDR. You'll also learn how to check you're running the very latest drivers for all hardware components, as well as how to monitor game performance to measure the effectiveness of any changes via the Xbox Game Bar.

Some of these optimizations come at the expense of other system resources, such as background processes, so we always recommend using a dedicated gaming PC. –NATE DRAKE

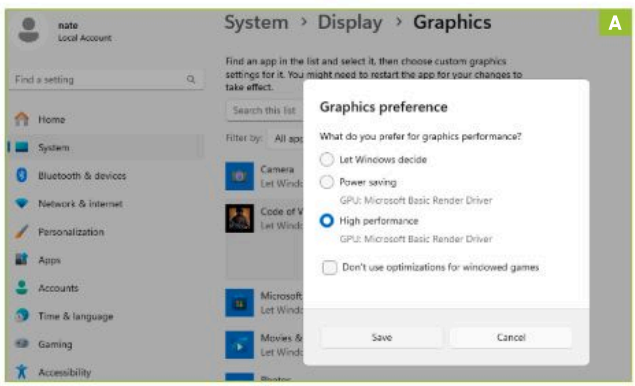

1 ACTIVATE AND TWEAK GAME MODE

According to the Xbox website, the Windows-dedicated 'Game Mode' is designed to 'prioritze your gaming experience'. This means that while activated, Windows won't perform driver installations or send any restart notifications.

» To get started, simply open Settings via the Windows search bar. From here, select 'Gaming' from the left-hand tab, then click 'Game Mode'.

» You can then enable Game Mode using the rocker switch. As the XBox website states, Game Mode is in fact enabled by default, so you may not have to do anything further.

» While Game Mode is enabled, any running games should be prioritized by your device's CPU and GPU. System resources for background processes will also be deprioritized.

» To fine-tune Game Mode further, hit 'Graphics'. From here, you can configure graphics performance for selected apps, such as your games.

» To get started, simply click the drop-down menu under 'Add an App' to choose either a Desktop or Microsoft Store program.

» Next, click 'Browse' to navigate to your chosen game, which will now appear in the list of apps at the bottom of the screen. Choose 'Options' to fine-tune settings. Unsurprisingly, [Image A] you're likely to achieve the best results by choosing 'High Performance', as well as leaving 'Don't use optimization for windowed games' unchecked. Click 'Save' to continue.

» As you browse through the list of other apps, you'll see the default is to 'Let Windows Decide' what resources to allocate. If you feel an app has been mistakenly flagged as 'High Performance', repeat the above steps to choose 'Power Saving' instead.

» You can also check current system performance by invoking the Xbox Game Bar (Windows + G). Click 'Performance' in the top pane to see a readout of current CPU/GPU usage.

2 DISABLE MEMORY INTEGRITY

In 2022, Windows 11 introduced new security features including Hypervisor-protected code integrity (HVCI), also known as Memory Integrity, and the Virtual Machine Platform (VMP).

» Memory Integrity is actually a powerful feature for keeping Windows safe—it double checks that drivers being installed onto the system are safe through verifying digital signatures and even checks for attackers injecting malicious code into system processes. VMP provides core virtual machine services.

» It wasn't long after these services were introduced and enabled by default that gamers noticed a hit to their FPS. Microsoft have acknowledged this, and provide instructions to disable Memory Integrity and VMP to remove any potential impact on game performance.

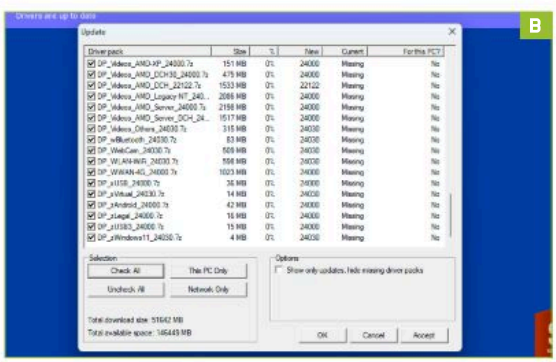

© MICROSOFT, SNAPPY DRIVER

» Needless to say, there are some very serious security implications to doing this. When researching this article, we also discovered that not every game experiences significant slowdown from HVCI/VMP, so proceed at your own risk.

» To get started, click the Start menu and enter 'Core Isolation'. Click on this to launch the Windows Settings window. Next, under 'Memory Integrity', hit the rocker switch to to 'Off'. You can also disable the 'Microsoft Vulnerable Driver Blocklist', though we didn't notice any significant improvement in gaming performance when doing so during our tests.

» To disable VMP, hit the Start menu and search for 'Windows Features'. Select 'Turn Windows Features on or off'. In the new window, find 'Virtual Machine Platform' and deselect.

» You may need to restart your device to save these configurations. Needless to say, you should enable Memory Integrity when you've finished gaming to keep your device safe.

3 UPDATE DRIVERS MANUALLY

When gaming, it stands to reason that you'll want the most bleeding edge drivers for full hardware compatibility, particularly when it comes to graphics cards.

» Though it may sound obvious, one of the best ways to make sure you're running optimal drivers is to enter 'Windows Update' in the search bar and 'Check for Updates'. Bear in mind if you've enabled 'Game Mode', you may not see all restart notifications.

» Once the updates are complete, click the Start menu once again to search for and launch 'Device Manager'.

» Expand the "Display Adapter' category and right-click your graphics card. Choose 'Properties', then hit the 'Driver' tab to note down your GPU model and version number.

» At this stage, you can click 'Update Driver' to 'Search Automatically for Drivers' but if you've just updated Windows, this is unlikely to have any effect. Instead, launch your browser and visit your GPU manufacturer's homepage, eg. www.nvidia.com. Here, you can visit the Support pages to download the correct drivers to your make and model. When selecting a driver, make sure it's designed for the version of Windows you're running.

» Once your drivers have finished downloading, return to Device Manager. This time, you can select 'Update Driver', then 'Browse my computer for drivers to select your recent download'.

» Some manufacturers will only make the drivers available as part of an automated installation program. If this happens, you can usually install the driver automatically just by running the downloaded setup file. The advantage of these type of installers is they can double check for compatible hardware before proceeding.

4 UPDATE DRIVERS AUTOMATICALLY

Manually updating each of your hardware drivers is one of the very best ways to ensure optimal gaming performance.

C

© XONTIC

» Still, this does involve detailing the exact make and model of every component, as well as visiting the manufacturer website to make sure that you have the best driver in question. You'll need to repeat this process regularly to make sure your drivers are still up to date.

» It's no surprise that manufacturers have developed control panels like NVIDIA's GeForce Experience or Snappy Driver Installer Origin to automate the process of checking for and updating drivers. [Image B]

» The specific steps you'll take to download and use such software will vary depending on the tool itself. There's also plenty of adware online thinly disguised as software designed to turbo charge your PC gaming experience.

» To this end, make sure you only download driver manager software directly from the manufacturer or from a trusted third party—for instance, the Snappy Driver Installer project is no longer in the hands of the original developers, but the 'Origin' fork remains free and open-source, so is safe to use.

» If you do use a more generic driver manager like this, then in the first instance we recommend only downloading the driver indexes. That way, you can check

USING AUTO HDR

If you're running Windows 11, depending on your hardware, you can enhance gaming performance through Auto HDR (Auto High Dynamic Range). In theory, this can improve the image and video quality of SDR (Standard Dynamic Range) in games to HDR.

To get started, type 'HDR' into the Windows search bar, and choose 'Use HDR settings'. Here, you can discover if your hardware supports HDR. If so, you can activate various options by selecting the corresponding rocker switch, such as 'HDR Video Streaming'.

You can play the sample clip in the top left of the window to view how video will look with your chosen settings.

It can also be accessed via the Gaming Bar. Hold Windows + G to launch, then click the settings icon in the top right. Choose 'Gaming Features' and enable both 'Use HDR' and 'Use Auto HDR with supported games'.

Be warned that for some titles, enabling HDR options will enhance game brightness considerably. This may have implications for battery life if you're playing on a portable device. You also won't observe any noticeable difference if you enable this feature for games that already have their own HDR implementation.

MANAGE VITALS USING XBOX GAME BAR

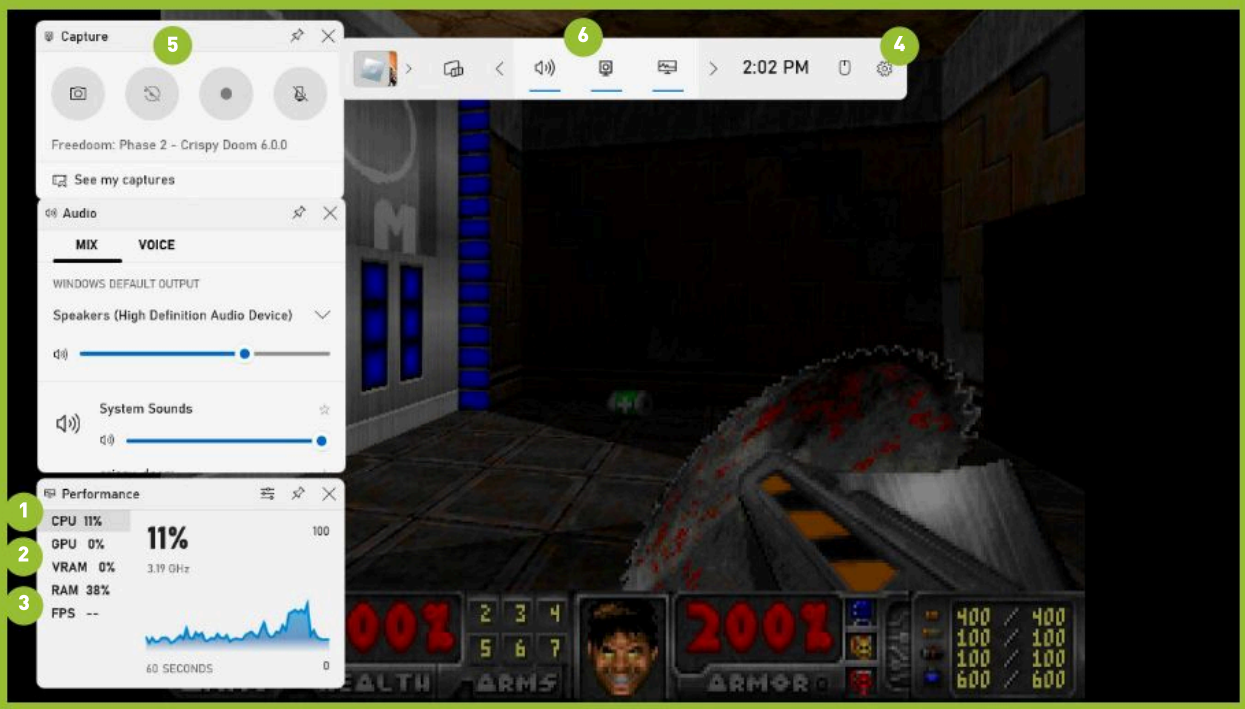

1. CPU
This element determines the CPU usage of the game in question. By default, the Xbox Game Bar also displays this in a line graph to the right.

2. GPU AND VRAM
The GPU metric naturally measures how much video processing power it is using (if any). VRAM simply represents how much video memory the graphics card itself is currently consuming.

3. FPS
The higher the FPS (fames per second), the smoother a game will look. The display graph can be configured to show average FPS over the past 60 seconds—helpful for checking optimizations.

4. SETTINGS
Use the XBox Game Bar settings to change configurations. Once opened, click into 'Gaming Features' to enable both 'Use HDR' and 'Use Auto HDR with supported games'.

5. CAPTURE
The Xbox Game bar can capture images and videos as you play, though this will place a greater burden on system resources. You can disable this feature via the settings in Windows.

6. WIDGETS
Click on any of these icons to display the relevant menus. Widgets can be moved, resized, and pinned elsewhere on screen. Use the arrow tabs either side of the default widgets to view others.

what components need to be updated without downloading lots of drivers needlessly.

5 TWEAK VISUAL EFFECTS AND VIRTUAL MEMORY

Windows has a number of fancy visual effects running in the background that when combined, can reduce system performance when you're running a game.

» To fix this, first type 'Performance' into the Windows search bar in order to choose 'Adjust the appearance and performance of Windows'.

» Next, select 'Adjust for Best Performance'. This will change a number of defaults, such as disabling the fade effects for ToolTips and Slide Menus. Click into the 'Advanced' tab. Under 'Processor Scheduling', ensure that resources are adjusted for best performance of 'Programs' rather than 'Background Services'. [**Image C**]

» Take a moment to read through the 'Virtual Memory' section. This is where you can specify the size of the paging file, which Windows can use as if it were RAM. By default, it should be around 25 percent of your physical RAM memory in size.

» We're drawing readers' attention to this, as some gamers think that having a huge paging file can improve game performance, given it superficially increases available RAM. This ignores the fact that read/write speeds to RAM chips are at least 10 times faster than that for hard disks, not to mention the increased disk usage caused by huge amounts of data being saved.

» From browsing gaming forums like Steam Community, we have seen examples of performance being improved for resource-hungry games like *Starfield* through disabling the paging file. This, of course, means Windows uses only physical RAM, which is much faster.

» If you choose to do this, click on 'Change' in the 'Virtual Memory', then uncheck 'Automatically manage paging file size for all drives.' Next, select 'No paging file', then 'Set' to confirm your changes. You may need to restart Windows for this to have an effect.

» Disabling the paging file can have unintended consequences. In the first instance, you'll need to be sure that your chosen game can work with whatever

physical RAM and GPU memory you have available. Certain features, like Hibernation, also won't perform correctly.

» While you're tweaking settings, this is also a good opportunity to double-check those within your game itself. These will differ from game to game, but we recommend first focusing on Anti-Aliasing. In simplest terms, this visual effects blurs lines between irregular textures or objects, making graphics appear smoother. Most modern games have different levels of anti-aliasing, so experiment with these to see what works best.

» Some titles may even have a specific 'Motion Blur' setting to make the game look more realistic. If your GPU is struggling, disabling this feature can improve game performance.

» Check if your chosen game supports Vertical Vsync (VSYNC). This feature can help prevent 'tearing', where a game's FPS exceeds the monitor refresh rate, displaying multiple frames at once. VSYNC can limit your game frame to stop this happening, but adds 'lag'. If this feature is available, we suggest trying the game both with and without it to test performance.

6 RECHECK YOUR REFRESH RATE

Every good gamer knows that a device's refresh rate is measured by how often it can refresh this screen within one second, and is measured in Hertz (Hz). In other words, a monitor with a 60Hz refresh rate can theoretically display new images at a rate of up to 60 per second.

» While this sounds great in theory, your device's refresh rate will depend on a number of factors, including what graphics card it has, your chosen screen resolution, and the quality of the monitor itself.

» Naturally, the higher your screen resolution and device refresh rate, the better the quality of your gaming experience will be when viewing high-quality images in quick succession.

» To double-check that your resolution and refresh rates are where they should be, click the search bar, and enter 'Display'. Click on 'Display Settings'. Here, you can adjust 'Display Resolution' via the drop-down menu in 'Scale & layout'.

» Next, scroll down to 'Related Settings' and select 'Advanced Display'. If you use multiple monitors for gaming, use the drop-down menu in the top right to select each one. Click into the drop-down menu next to 'Choose a refresh rate' to make changes.

7 CUSTOMIZE MOUSE SETTINGS

By default, Windows 11 will speed up or slow down mouse movements depending on how fast you move the cursor. This can be useful when navigating the desktop or productivity applications, but can be a nightmare for games that require pinpoint accuracy, like first-person shooters.

» To disable this feature, type 'mouse' into the Windows search bar, then select 'Mouse Settings'. Scroll down to the 'Related Settings' section, then click 'Additional mouse settings'. In the 'Mouse Properties' window, click the 'Pointer Options' tab.

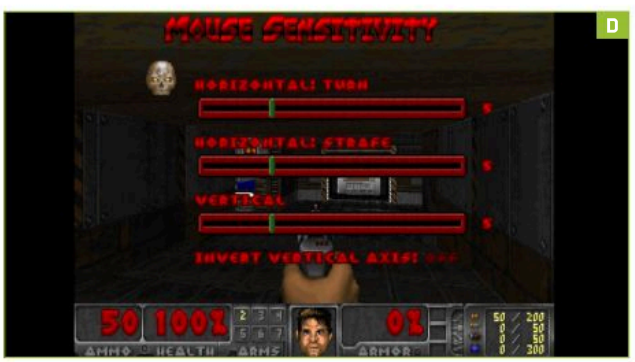

D

» You can now deselect 'Enhance pointer precision'. Although it's unlikely to affect gameplay, you can also disable 'Automatically move pointer to the default button in a dialog box.'

» From here, you can also vary the pointer speed, though don't fall for the assumption that a fast cursor will automatically lead to better gameplay, as higher speeds tend to lead to less accuracy. [Image D]

» Modern gaming mouses measure sensitivity in DPI (Dots per Inch). This indicates how many pixels the cursor will move for every inch you move the mouse itself. The higher the DPI, the further the cursor will move with minimal physical mouse movement.

» If you're still having sensitivity issues after adjusting the pointer speed in Windows settings, check if the mouse manufacturer has released their own control software to allow you fine-tune it with more precision. Certain games also allow you to adjust mouse sensitivity via the Settings menu.

SUPERSTAR BAR

The Xbox Game Bar can be a useful overlay for gamers, allowing you to view performance metrics and tweak settings like Auto HDR support. If you're as trigger happy as we are at *Maximum PC*, however, there's always a danger that you'll accidentally activate the Game Bar at an inconvenient time.

To remove the Xbox Game Bar altogether, type 'Terminal' into the Windows search bar and click on 'Windows Terminal'. From here, enter the command: Get-AppxPackage Microsoft.XboxGamingOverlay | Remove-AppxPackage

If you're looking for an alternative overlay, there are a number of options, including Razer Cortex (pictured), which contains a number of enhancements for game performance. These include disabling CPU 'Sleep Mode', clearing the RAM of non-essential programs besides the game itself, and pausing programs and services.

If you want to disable DVR to prevent the accidental recording of video while playing, access 'Settings' via the Windows search bar, and choose 'Gaming'. Here, deactivate the switch marked 'Record what happened'.

Though this can be less distracting, it also means there's no built-in way to record gaming via video. You can, however, search for other screen-recording applications, such as BandiCam or OBS Studio. These offer more advanced video recording features beyond the Game Bar, such as customizable hotkeys, so you don't accidentally start recording mid-game.

Migrate from Docker to Podman

YOU'LL NEED THIS

PODMAN

You can install this natively in Linux or through WSL in Windows. For a beginner's guide to Podman, see the February 2024 issue.

A COUPLE OF ISSUES AGO, we introduced you to a new, more secure way of running containerized apps on your PC: Podman. However, while Podman is designed to work with Docker containers, the way it functions means they're not 100 percent compatible—at least in their native form. That's down to two primary reasons: the first is that Podman runs under your own user account to provide containers with 'rootless' access to your system, and the second is that containers are set up to run independently instead of through an additional Docker daemon process.

Rootless access prevents containers from being given the keys to your system, but can cause problems for those that need elevated access in certain areas. The lack of a parent daemon also requires a different approach to configuring containers to auto-start with your system. In this tutorial, we'll examine these key differences and explain ways in which you can circumvent them to get a selection of popular containers. You'll also equip yourself with the know-how to get other containers working. We've focused on the Linux implementation of Podman, but you can easily tweak these instructions to work with Podman in Windows, too. –NICK PEERS

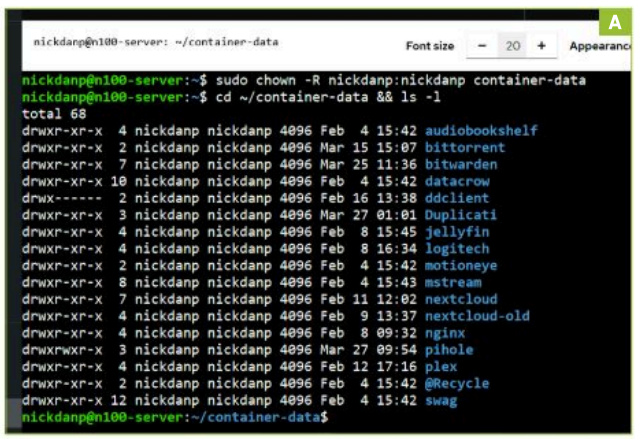

1 PREPARATORY STEPS

Whether installing Podman alongside an existing installation of Docker, or setting it up for the first time on a new machine, the first thing you need to do is migrate your existing container data across to Podman. We recommend storing everything inside a parent container data folder (container-data or podman-container-data, if Docker is already present) in an accessible location. This might be inside C:\Users\Username in Windows, or your home (~) folder in Linux. You then create individual subfolders for each container inside this.

» Because Podman runs rootless by default under your own username, you'll need to ensure that all your container folders are accessible. Permission issues don't exist for Windows users because of the way both Docker and Podman run within the Windows Subsystem for Linux, but if you're running a native Linux install like Debian or Ubuntu Server, you'll need to take ownership of the entire container-data folder. Input the following command, substituting username:username with your own Linux username, such as nick:nick):

```
sudo chown -R username:username ~/container-data
```

» Now, verify that the permissions have been updated:

```
cd ~/container-data && ls -l
```

» If all is as it should be, your username should be listed as the owner of all the sub-folders, as shown in [Image A].

2 YOUR FIRST MIGRATION

If you've not already done so, consult our beginners' Podman feature in the February 2024 issue. This introduced you to Podman, and revealed how to set it up in a Windows instance and install your first containers: Vaultwarden (a self-hosted instance of the Bitwarden password manager), and Nginx Proxy Manager (a reverse proxy). If you check the screenshot on page 45 of that issue, you'll see the commands required to create both instances in Windows. The steps are identical for Linux machines, except you don't use the character as an escape to indicate line breaks; instead, use the backslash (\) character, as shown in [Image B].

» If you compare this to the command required to run Vaultwarden in Docker (see https://github.com/dani-garcia/vaultwarden), there are two differences. First, we've referenced the bridge-for-podman network. We need to create this dedicated bridge network, which is required should you want to use Vaultwarden in conjunction with Nginx Proxy Manager (see step 4) to allow you to access it remotely. Thankfully, you can set this up with a single command:

```
podman network create bridge-for-podman
```

» The second difference can be found in the final line: we've added 'docker.io/'. This is because Podman needs to know which container repository you're pulling the container from.

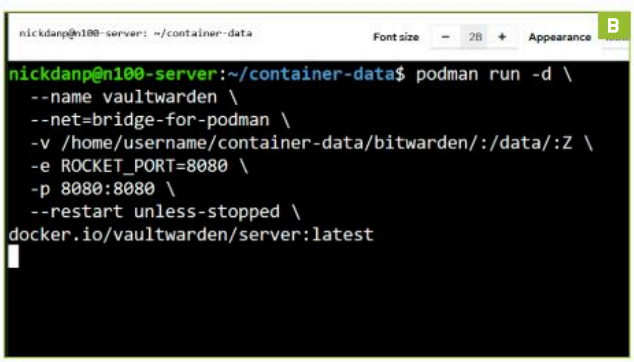

NAVIGATE PODMAN IN COCKPIT

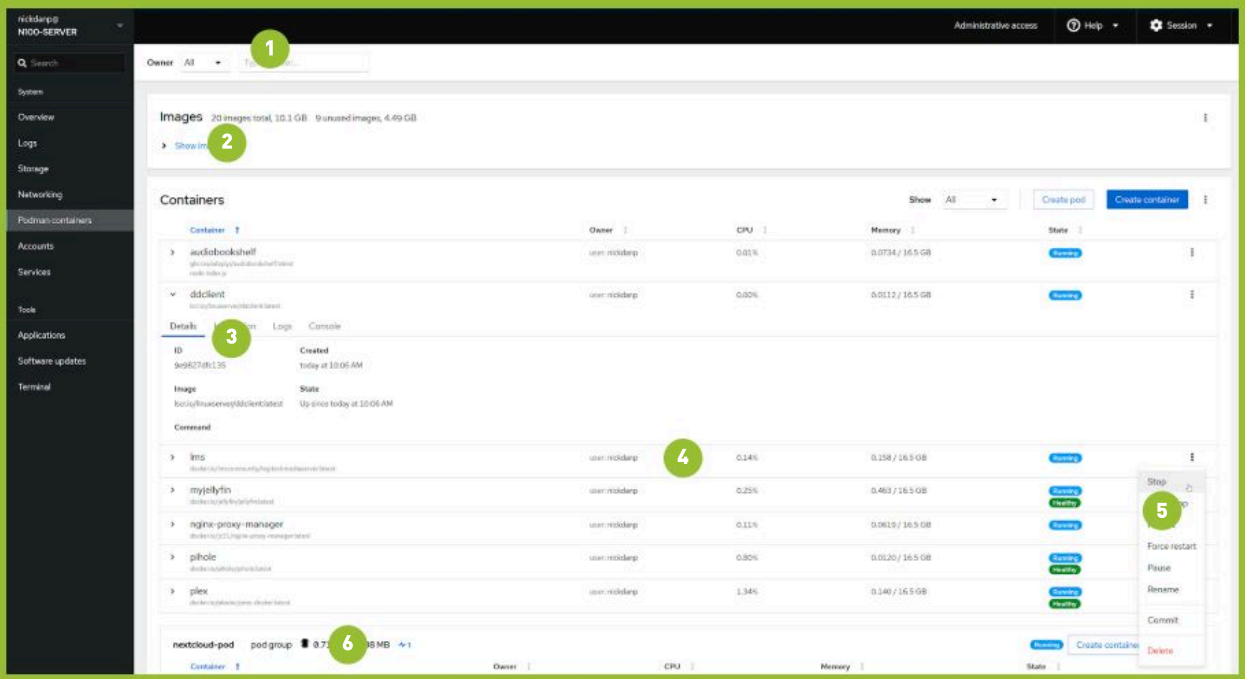

1. FILTER CONTAINERS
Cockpit's podman plugin displays all running containers regardless of which account launched them by default. Use the drop-down menu and keyword filter to create customized views.

2. IMAGES
Click 'Show images' to view all downloaded images—including those not currently being used. Click 'Create container' next to one to set up a container from it.

3. CONTAINER DETAILS
Click > next to a container to reveal more information, plus gain access to both system logs (for troubleshooting) and a console for directly interacting with the container.

4. SYSTEM RESOURCES
Each container reveals which user account is running it, as well as its CPU and RAM usage. This updates in real time to help you identify resource hogs.

5. CONTAINER ACTIONS
Click the tricolon to reveal a menu from which you can start, stop, restart, and pause the container. You can also rename it, commit a new image, perform a health check, and delete it.

6. POD DETAILS
You'll see a summary of the pod's current system resource usage—if the pod has set up port forwarding, click the blue link next to the memory usage to view a pop-up summary.

3 ACCESS PRIVILEGED PORTS

Our next migrations cover the reverse proxy Nginx Proxy Manager and ad-blocking service Pi-Hole. Both containers require access to so-called 'privileged' ports—ports 0 to 1023—that require elevated access. These ports cover common ports used by both services, such as 53/67 for Pi-hole's DNS filtering, and 80/443 for Nginx's proxy. Ordinarily, you'd make use of the -p flag to redirect ports to higher numbers (such as -p 8080:80), but this won't work with these two containers.

» There's a quick and dirty fix if you're in a hurry, which basically gives all unprivileged applications access to ports lower than 1024 with a simple system file edit:

```
sudo nano /etc/sysctrl.conf
```

» Add the following line:

```
net.ipv4.ip_unprivileged_port_start=53
```

» It's not an ideal workaround for security reasons, but it's the one recommended by Podman developers, so who are we to argue?

» With this in place, and assuming you've created the bridge-for-podman custom network in step 2, you can now deploy a fully working Nginx Proxy Manager container with the script, as shown in [Image C].

C

```
podman run -d \
  --name=nginx-proxy-manager \
  --net=bridge-for-podman \
  -p 443:443 \
  -p 80:80 \
  -p 81:81 \
  -v /home/username/container-data/nginx/data:/data \
  -v /home/username/container-data/nginx/letsencrypt:/etc/letsencrypt \
  docker.io/jc21/nginx-proxy-manager:latest
```

4 MORE NETWORK TWEAKS

As things stand, Pi-hole will now run too, were you to copy the Docker script, but it won't work properly. We can clear these hurdles by adding some

D

```
podman run -d \
  --name pihole \
  -e TZ="Europe/London" \
  -e INTERFACE=tap0 \
  -p 53:53/tcp \
  -p 53:53/udp \
  -p 8084:80 \
  -v "/home/username/container-data/pihole:/etc/pihole/" \
  --userns keep-id \
  --net=slirp4netns:port_handler=slirp4netns \
  --dns=127.0.0.1 \
  --dns=1.1.1.1 \
  docker.io/pihole/pihole:latest
```

E

```
podman run -d \
  --name duplicati \
  -e UID=1000 \
  -e GID=1000 \
  -v /home/username/container-data/:/data \
  -v /home/username/media/bay5/backups/:/backups \
  -v /home/username/:/n100-server \
  -p 8200:8200 \
  docker.io/duplicati/duplicati:canary
```

additional lines to the script. First, the following two lines adapt Pi-hole for Podman's networking implementation:

```
--net=slirp4netns:port_handler=slirp4netns \
-e INTERFACE="tap0" \
```

» This ensures Pi-hole can accurately see (and record) what devices are using it. Finally, if you plan to deploy Pi-hole as a DHCP server, you'll also need to add the following line to the Pi-hole script, as shown in [**Image D**]:

```
--cap-add=NET_ADMIN \
```

» The '--cap-add' option can be used to grant additional capabilities to unprivileged containers, and NET_ADMIN allows Pi-hole to create the network interface it needs to act as a DHCP server. You can now follow our setup guide in the January 2024 issue to block unwanted ads from your entire network.

5 PERMISSIONS ISSUES

Read the 'Rootless considerations' box for a primer on how user permissions work inside and outside of containers through user mapping. Most containers run internally as the 'root' user (UID 0), and Docker maps this to your own PC's root user (also 0). Podman's rootless nature means that it maps the internal root user to your own user (UID 1000).

» We've seen how you can work around certain demands for access to system resources normally restricted to root users. However, another problem area with rootless containers on Linux machines is with file permissions. One issue occurs when the container sets up internal users alongside root. These are mapped to different UIDs, and when those write data to your container-data folderm, they can take ownership of the folders.

» It's not something you need to necessarily worry about, but it can hamper efforts to back up those folders if their permissions are set in such a way that only the owner (or sudo) has access to the files. You can see if this has happened after setting up a container by navigating to the container-data directory in Terminal and typing the following:

```
ls -l
```

» This will list all your container folders along with the current owner—those marked with your username are fully accessible and can be safely ignored. However, you may find some are owned by numbers (typically a number between 100,000 and 165,535). Check the permissions on the left—if they're drwxr-xr-x then there's nothing to worry about—you can access those files from your user account, but if they're drwx------ (as they will with Linuxserver's ddclient container) then you won't be able to access them without sudo.

» When this happens, stop and delete the container using the controls in Cockpit or the command line—podman stop and podman remove. Next, take back ownership of the container folder using the same 'sudo chown' command, as outlined in step 1. Now, try inserting the following line into your Podman script:

```
--userns=keep-id
```

» This instructs Podman to map your user ID/GID to the same ID/GID inside the container (so 1000:1000). If the internal root user doesn't require outside access to your system, you may get lucky: the container will start, but you'll retain ownership of the folder. We've used this successfully with Nextcloud (see final step) and Jellyfin, among others. Sadly, it doesn't work with ddclient.

» You can configure --userns to behave in other ways, too. To avoid mapping your own user UID to the container's internal user for security reasons, try --userns=nomap instead, which will allocate the root user the first available UID within your userspace (so 100000 or upwards). Check out the RedHat documentation (www.

ROOTLESS PERMISSIONS

Containers are self-contained environments that utilize their own system resources. These resources include their own internal user accounts—such as a root user—that needs to be 'mapped' to a user on your system whenever the container needs access to outside resources. This feature—known as 'namespaces'—is designed to improve security by limiting the container's access to the rest of your PC.

By default, Podman maps the root user inside a container to your less-privileged user, restricting its access. This is part of Podman's 'rootless' nature that differentiates it from Docker, which runs in 'rootful' mode by default, mapping root users inside containers to the root user outside it. It's great for compatibility, but not for security. As an aside, you can run Podman in 'rootful' mode by prefixing your podman run command with 'sudo'.

How does this work in practice? Users and groups are allocated unique numeric IDs both inside and outside the container, called UIDs and GIDs respectively, and one of Podman's jobs is to map your user's UID/GID (typically 1000) to whatever user the container is running under, typically root (so 0).

You can auto update containers at https://rayagainstthemachine.net/linux%20administration/podman-auto-update/.

redhat.com/sysadmin/rootless-podman-user-namespace-modes) for a more detailed explanation.

» If userns doesn't help, all is not lost. Some containers, like Duplicati, respond to a different tweak—in this case, specifying the UID and GID as environment variables as shown in [Image E] fixed our permissions and access problems:

```
-e UID=1000 \
-e GID=1000 \
```

6 AUTO-START YOUR CONTAINERS

One key difference between Podman and Docker is that Podman ditches the parent daemon process that Docker uses to start and run your containers. Instead, each container is self-contained, so if another container crashes, it won't bring it down with it. One problem this creates is when you restart your PC. Docker's daemon is responsible for automatically restarting any containers with a suitable --restart policy (such as --restart=always or --restart=unless-stopped).

» Although Podman comes with a specialized service (podman-restart) whose job is to restart containers with restart policies of always or unless stopped, it only works on containers already running. To get containers starting with your PC, you'll need to generate 'systemd unit files', one per container. These allow containers to run without relying on a parent daemon.

» Windows users are covered by the Podman Desktop tool—see the February 2024 issue for details. Linux users, however, must generate these scripts themselves. Before doing so, enter the following command into a Terminal:

```
sudo loginctl enable-linger
```

» This allows containers running under your username to launch at startup without you having to log into your account first.

» You're now ready to generate the scripts. Podman can create a script from any running container with the 'podman generate systemd' command. The following generates a preview of the file that's created—just make sure the --name flag points to the name of a running container (such as Vaultwarden):

```
podman generate systemd --new --name vaultwarden
```

» If you're running Podman 4.6 or later, you'll see a warning about this being a depreciated command—this can be safely ignored. Instead, read through the script shown in [Image F]—you should see it looks vaguely familiar. Now, let's create that startup script for real:

```
podman generate systemd --new --name vaultwarden -f
```

» Now, you need to move the newly created file and enable it as a service:

```
mv -v container-vaultwarden.service ~/.config/systemd/user/
systemctl --user daemon-reload
systemctl --user enable container-vaultwarden.service
```

F

```
nickdanp@n100-server: ~                    Font size  –  17  +  Appeara
[Unit]
Description=Podman container-vaultwarden.service
Documentation=man:podman-generate-systemd(1)
Wants=network-online.target
After=network-online.target
RequiresMountsFor=%t/containers

[Service]
Environment=PODMAN_SYSTEMD_UNIT=%n
Restart=on-failure
TimeoutStopSec=70
ExecStart=/usr/bin/podman run \
        --cidfile=%t/%n.ctr-id \
        --cgroups=no-conmon \
        --rm \
        --sdnotify=conmon \
        --replace \
        -d \
        --pod=vaultwarden-pod \
        --name vaultwarden \
        --net-bridge-for-podman \
        -v /home/nickdanp/container-data/bitwarden/:/data/:Z \
        -e ROCKET_PORT=8080 \
        -p 8080:8080 docker.io/vaultwarden/server:latest
```

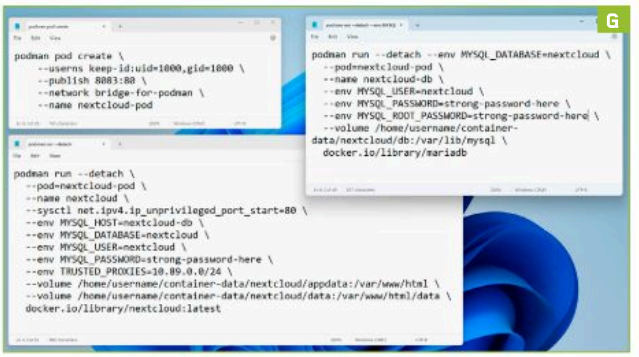

» If you remove and recreate the container with different settings, you'll need to generate a new systemd service file—just repeat the previous four commands.

7 WORKING WITH PODS

Some containers are quite complex affairs—Nextcloud is a case in point. Docker users can use the Nextcloud All-In-One container as outlined in the October 2023 issue, but one look at https://github.com/nextcloud/all-in-one/discussions/3487 reveals that it's not working in Podman (believe us, we've tried).

» In researching an alternative, we discovered one of Podman's standout features: pods. This lets you group related containers together under a single roof—perfect for security, but also to keep things organized. Our final migration sets up a Nextcloud pod. As you can see in Image G, certain options, including --userns, --publish, and --network are defined here, and will apply to all containers within the pod. Think of it as an extra wrapper.

» The --userns option is vital for running Nextcloud as a regular user while retaining access to all files and folders, including your data. You'll need to set up each element separately—start with the podman pod create command, then set up up the mariadb database. Finish with the Nextcloud instance at the bottom of [Image G].

» Navigate to 'Administration settings' in the Nextcloud web UI, and select Basic Settings. Set Background Jobs to 'Cron (Recommended)'. Now, return to the main Terminal and type the following command:

```
crontab -e
```

» If prompted, choose nano. Input the following line, substituting username with your Linux username:

```
*/5 * * * * podman  exec -t -u username nextcloud php -f /var/www/html/cron.php
```

» Save your changes and exit, then type:

```
sudo service cron reload
```

» Finally, you'll need to generate the auto-start files, which requires just a single command:

```
podman generate systemd --new --name nextcloud-pod -f
```

» This will generate three files: one pod, two containers. All three require moving:

```
mv -v pod-nextcloud-pod.service ~/.config/systemd/user/
mv -v container-nextcloud* ~/.config/systemd/user/
```

» Finally, enable everything to complete your fully functional Podman Nextcloud instance:

```
systemctl --user daemon-reload
systemctl --user enable pod-nextcloud-pod.service
systemctl --user enable container-nextcloud-db.service
systemctl --user enable container-nextcloud.service
```

Debloat Windows 11

YOU'LL NEED THIS

WINDOWS 11
An up to date installation.
(Optional) 8GB USB stick for clean install

THERE'S NOTHING QUITE as annoying as booting up your brand new Windows 11 PC for the first time, only to find that it's crawling to a snail's pace due to the presence of bloatware. Common culprits include 'registration' programs from your machine manufacturer, as well as trial versions of antivirus and backup programs.

Your OEM (Original Equipment Manufacturer) may want you to use their proprietary apps or ones that they've been paid to preinstall, but this isn't necessarily best for your machine. Bloatware consumes system resources, and can even pose a security risk due to undisclosed bugs.

In this guide, you'll learn how to go through your Windows 11 device and list apps, distinguish the unnecessary ones, then remove them. You'll also discover how to reset your device to factory settings, if necessary. Windows has safeguards to prevent you from deleting anything crucial but, we recommend making a full system backup before proceeding.–NATE DRAKE

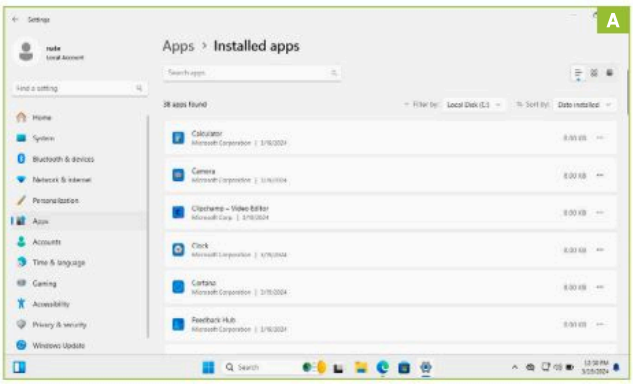

1 REMOVE BLOATWARE BY SETTINGS

The internet is full of apps that offer to automatically remove 'junk files' from your PC, but when it comes to removing bloatware, installing extra software can be counterproductive. It can also lead to security issues (as evidenced by the 2018 supply chain attack on CCleaner).

» Luckily, it's simple to remove installed applications using Windows' own built-in tools. To get started, open Settings, then choose 'Apps' in the left-hand pane. Select 'Installed Apps' to view a list of all applications.

» By default, this is a long A-Z list of all programs on all drives. Use 'Filter By' to list only applications for individual disks. Select the 'Sort By' menu to change from A-Z to something more helpful, eg. 'Date Installed'. [Image A]

» To remove an app, simply click the ... icon next to its name, and select 'Uninstall'. You'll see an alert informing you that the program and all related info will be removed. Hit 'Uninstall' to continue. If this is a third-party app, then the install wizard of the program in question will launch to continue the removal process. [Image B]

» You can't uninstall vital system apps, so it's usually safe to remove any apps here. Unlike malware, bloatware isn't designed to damage your, system so can encompass legitimate apps like Microsoft News or Nero BackItUp.

» If Windows comes bundled with your device, it's likely that your equipment manufacturer bundled a large amount of unnecessary apps. Some of these, like TOSHIBARegistration, are easy to detect.

» If you're uncertain about which programs to remove, open www.shouldiremoveit.com on a separate device. This site contains details of almost 85 million apps, as well as rankings for manufacturer bloatware (currently, Toshiba is in first place).

» You can also remove apps via the Start menu. Simply launch the menu, select 'All Apps', then Scroll down to your chosen program. Right-click to 'Uninstall' via the pop-up menu.

2 REMOVING STUBBORN BLOATWARE

Once you've gone through the Uninstall process, you may see that your chosen app is still listed in the 'Apps' menu.

» If a program persists, try manual removal via Powershell. Search for the utility via the Windows search bar, then right-click to 'Run as Administrator'.

Next, enter:

```
Get-AppxPackage -AllUsers | Select Name,
PackageFullName
```

» This will display a list of all currently installed Windows and third-party apps. Scroll down to find your chosen app in the list. [Image C]

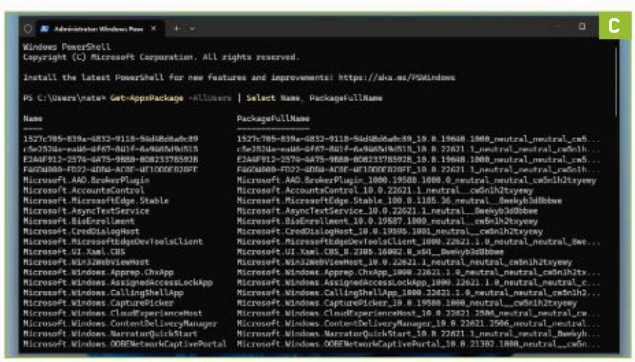

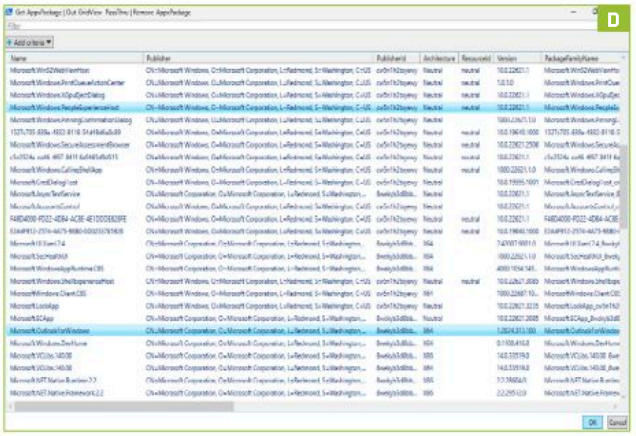

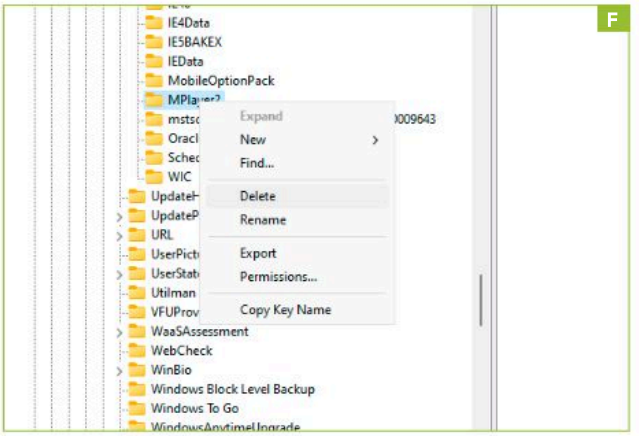

» You can use this to remove a program via Powershell. For instance, to rid your PC of 'Microsoft.BingNews' enter:

```
Get-AppxPackage -AllUsers Microsoft.BingNews | Remove-AppxPackage
```

» If you want to remove multiple apps, PowerShell offers a faster way to select and remove more than one program.

» Return to the command prompt and enter:

```
Get-AppxPackage | Out-GridView -PassThru | Remove-AppxPackage
```

» This performs the same function as the first command in that it lists installed apps, but uses 'Out-GridView' to pipe the output to a new window with a table of installed apps. [**Image D**]

» Here, you can use Ctrl + Click to select as many apps as you want, then choose 'OK' to proceed with removal. Some vital Windows apps are installed here, but in our tests, PowerShell threw up an error when we tried to remove anything vital.

» Even if your apps are removed, it may not remove all temporary files. These are in the 'Temp' file of your local user folder, eg. C:\Users\nate\AppData\Local\Temp. [**Image E**]

» From here, you can simply delete any files or folders relating to the bloatware you've just removed.

» If after rebooting bloatware is still listed in your installed apps, fire up the Registry Editor and navigate to HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Uninstall. [**Image F**]

» From here, you can right-click the key for individual programs and choose 'Delete'.

3 DOING A CLEAN INSTALL

If you don't enjoy tinkering with Windows settings, PowerShell, and the registry, you can remove all third-party bloatware in one go by doing a 'clean' install of Windows.

» You can use Windows settings to 'reset' your PC so it retains your own data and files, without third-party apps.

» To start, open Windows Settings. Select 'System' from the left-hand pane, then scroll down to 'Recovery'. Select this, then hit the button marked 'Reset PC'.

» In the new window, choose 'Keep my files' to keep your personal documents. Select a 'Local' or 'Cloud' reinstall as you see fit, then hit 'Next'.

» By default, this will remove all non-Windows apps and programs. Click 'View apps to be removed' to see a full list. [**Image G**]

» Click 'Reset' to continue. Your device will restart during the reset process.

» Once this is complete, if you've previously removed any Windows apps you don't need, such as Microsoft Edge, you may need to do so again via Settings by repeating the steps earlier in this guide.

» If you find that system performance hasn't improved significantly despite removing bloatware, you may need to do a clean install of Windows 11.

» This involves creating a USB install media, but this will erase everything on your device, including your personal files.

» If you're using an older PC, instead of stock Windows 11, consider installing Tiny11 instead (https://archive.org/details/tiny11-2311). The ISO is around 20 percent smaller than a regular Windows 11 one. Tiny11 is also based on Windows 11 2023 Update (23H2), but is much more lightweight, as most unnecessary software has been removed, including support for XBox and the now-defunct Cortana (CoPilot support is available).

» As speedy as Tiny11 is, as unofficial Windows software it could be prone to bugs and update issues. If you're unsure, use an official Windows 11 ISO, available from www.microsoft.com/en-ca/software-download/windows11.

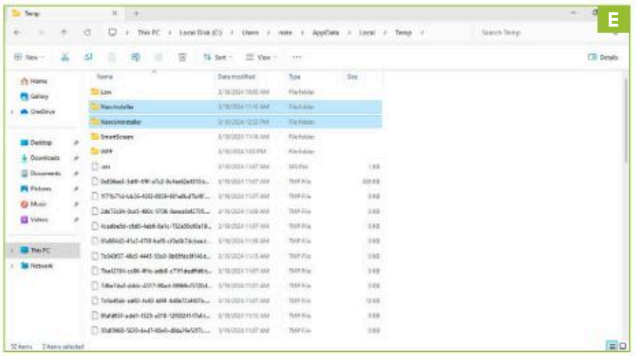

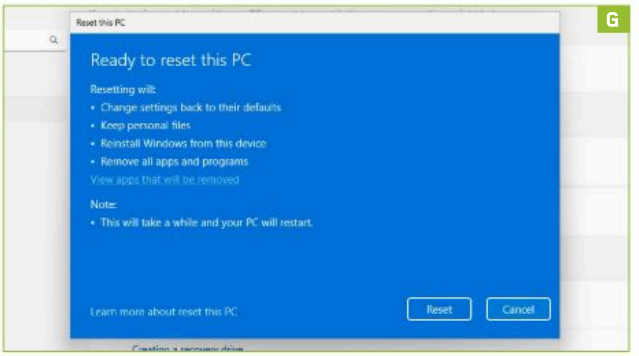

Create a mod with RTX Remix

YOU'LL NEED THIS

NVIDIA RTX 3060TI GRAPHICS CARD
or better
RTX Remix software
Retro game library
A lot of patience

WE'VE SEEN PLENTY of remastered or modded games recently that have updated the textures and even patched ray tracing into older titles. The RTX versions of *Quake II* and *Portal* are some of the best examples, as they were created by professional studios that know what they're doing.

If you don't know what you're doing, or just want to have a go because taking games apart to see how they work is interesting, Nvidia has released a beta version of RTX Remix that can do just that. It's time-consuming, fiddly, and there's no guarantee that you'll actually get results. However, in so many ways, it's the perfect pastime for PC gaming enthusiasts.

This being an Nvidia application, you'll need a graphics card from the green team. It'll need to be a fairly recent one, with RTX capabilities—the more powerful, the better. Apologies to AMD GPU users and those who bought Intel Arc cards—this will sadly not work for you.

It's also worth pointing out that this is largely going to be a story of failure. The first two games we tried didn't run, and the third had significant issues. Getting into modding with RTX Remix looks like being a long-term project with a steep learning curve. **–IAN EVENDEN**

1 INSTALL THE SOFTWARE

You can get the RTX Remix beta from nvidia.com with a simple internet search. Download the Omniverse app, log into it with your Nvidia account, then navigate to the Exchange section of the app, find the RTX Remix beta, and click install [**Image A**]. There are some minimum specs to consider—a four-core CPU and RTX 3060Ti, though it prefers eight CPU cores and an RTX 4070, and you'll need 16GB of RAM. While this is downloading (it's 6.75GB, so may take a while), there's a Tutorials button at the bottom of the interface that takes you to YouTube to get some pointers from the experts. There's also an FAQ with an installation guide and basic how-to, which does its best to put you off using the software by pointing out that it's for 'experienced modders'.

2 LAUNCH

There's a handy 'Launch' button in Omniverse's Library section—pressing it gets you into the RTX Remix app itself. This is beta software, and we found that this step could be a tiny bit crashy, especially on one of our PCs with a 16-core AMD Zen+ processor and an RTX 3080 GPU, but it ran first time on an i9-13900K and RTX 4090. There may be something about the older AMD chip, which only just scrapes the Windows 11 requirements, which Remix doesn't like. Whatever it was, the app wouldn't even run on that PC, despite it hitting all the other hardware requirements. Once it's open [**Image B**], you'll find a section called Exchange. In this, search for RTX Remix, click install, and wait while several gigabytes download.

3 HOW IT WORKS

RTX Remix comprises two parts: there's the runtime component, which exists as some .dll files, and the toolkit, which runs in the app you open from Omniverse. There's also a sample project, an .exe file that does nothing but make an R logo rotate within a rectangle with coloured walls. It's useful for learning how the system works before embarking on a game mod. It's not especially easy to find, as the file structure is obscure, but if you open the install location of RTX Remix from its Omniverse Settings page, it's in deps/remix_runtime/sample.

4 INTEGRATE THE RUNTIME

Sounds scary, huh? It's actually just a case of copying some files into the folder that contains the executable of the game you're trying to mod—in this case, sample.exe. The most recent versions of the runtime files are on GitHub, but RTX Remix comes with some too, and you can find them in /deps/remix_runtime/runtime. Copy everything in the /runtime folder to the /sample folder (or the location of the .exe for the game that you want to mod).

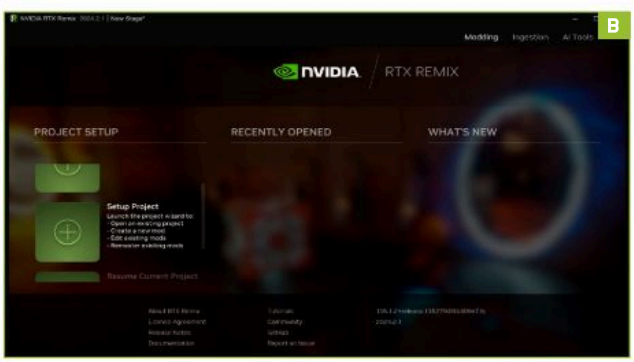

RTX HOME SCREEN

1. PREVIEW WINDOW
You can see what's going on in real time here.

2. DEVELOPER MENU
Tick the 'Always...' bot to get this every time you press Alt + X.

3. TABS
There are a lot of settings, arranged neatly in tabs.

4. SETTINGS
See what you can do by changing settings here.

5. BLOOM
Toggle bloom lighting for a quick demonstration.

6 SAVE BUTTON
It's important to press to save your changes.

5 RUN THE SAMPLE

If you now run the sample.exe file, you'll notice that it's a bit different to the way it was when you last saw it. Its lighting has been upgraded, and there are path-traced shadows present under the spinning letter. This is absolutely not what will happen when you try to mod a game—the sample file has been built with remixing in mind, while games from the early 2000s weren't, and you'll have to put quite a lot of work in before you see these kinds of upgrades. We tried it with some older titles, and they simply wouldn't start, or if they did launch, there was usually something badly wrong with them, such as NPCs' eyeballs displaying as blank white, or important parts of the in-game HUD being invisible.

6 ENTER THE MATRIX

With the sample app still open, press Alt+X, and you will get the RTX Remix User Graphics Settings window [Image C], from which you can hit the Developer Settings menu for a more fully featured options panel that slides in from the right. At this point, the beta nature of the software asserted itself again for us and it stopped responding, so remember to use that Save Settings button frequently.

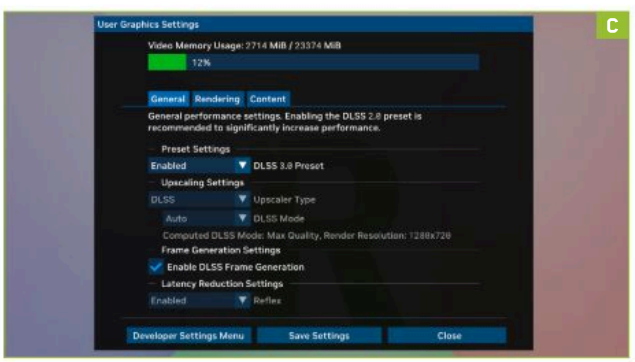

7 FIDDLE

One of the easiest things to do is turn bloom lighting on and off, which uses a single toggle under the Post-Processing section of the Rendering tab. You'll see the bloom disappear and reappear from the rotating image. To start upgrading a game properly, you'll need to perform an asset capture. Go to the Enhancements tab, and click Capture Scene after giving the capture a name. When it's at 100 percent, quit the sample game and open the /sample/rtx-remix/captures folder, where you'll find a .usd file with the name you gave it, plus folders with names like Lights and Meshes.

8 A NEW PROJECT

Go back to RTX Remix, and create a new project. You'll be asked for two folder locations. The top one can be wherever you like, perhaps user/documents/rtx-remix, but the second must be /rtx-remix. Choose the .usd file you made, and click 'create' to get the project started. This sets up a /mods folder in the /sample folder, and throws a couple of symlinks in—don't change the file or folder names. From this starting point, you can begin to create the remix mod for your chosen game, although if you have no dev experience, it's worth watching a lot of videos and reading forum threads first.

9 CHOOSE A GAME TO MOD

Now for the tricky bit. There are lists of compatible games on ModDB and Reddit, so it's worth perusing those before taking the plunge. A game that's fully compatible and already in your Steam (or GOG, or even Epic) library is a good place to start, though remixable titles tend to be of the older and cheaper variety. You'll need one that supports either DirectX 8 or 9, and there's a file you'll need to route DirectX 8 functions to DirectX 9, so check the small print before you attempt your Remix.

LAB NOTES

JEREMY LAIRD, CONTRIBUTOR

Arm is incoming

Qualcomm reckons games "just work" on its new Arm CPU

YOU'VE HEARD IT all before about Arm assimilating the PC at the expense of traditional x86 CPUs, but this time it might actually be happening. We're just months away from the first laptop PCs with the new Snapdragon X Elite chip, and Qualcomm is ramping up its marketing efforts.

Qualcomm reckons the chip has the measure of the best Intel and AMD CPUs for thin and light notebooks when it comes to both performance and efficiency. Even by Qualcomm's own numbers, Apple's latest M3 chip probably has the edge, but the performance and efficiency proposition looks promising.

The catch will be software support. Here, Qualcomm is again making big claims about the near-native performance of legacy x86 software running in emulation on Windows for Arm and the new Snapdragon X Elite processor.

The ability to run old x86 apps smoothly will be critical for Arm and the Snapdragon X Elite. Of all app types, it's games that are the most challenging to run with regards to emulation, but even there, Qualcomm reckons most games will "just run".

If it's true, then the final barrier to widespread Arm adoption on the PC will have fallen—almost. The Snapdragon X Elite's integrated GPU is claimed to be quicker than Intel's new Meteor Lake laptop chip, but it's no gaming powerhouse.

What will be required is pairing an Arm CPU with a proper graphics card. It's easy to imagine Nvidia supporting their GPUs on the Qualcomm Arm platform. On the other hand, will AMD want to encourage PC buyers to ditch x86 CPUs for Arm by providing full support for Radeon GPUs?

Whatever happens, we won't have long to find out whether Qualcomm really can deliver on its claims for the Snapdragon X Elite, as the first laptops go on sale in June. I'm pretty sure they'll blow away Intel and AMD laptops for efficiency and battery life. How they'll compare for performance and how well they'll cope with legacy x86 apps will be the really interesting bit.

Qualcomm has AMD and Intel in its sights with the new Snapdragon X Elite.

GUY COCKER
Editor-in-Chief

AMD has wheeled out new 'F' versions of its Ryzen 8000 Series desktop APU with the integrated GPU hacked out. Okay, the graphics are merely disabled rather than removed, but the result is the same: an APU with no graphics, which essentially amounts to a CPU.

The announcement came at a trade show in China, and it's not clear how the new chips will be sold. My guess is that they're aimed at system builders, and might not be sold separately to the public.

Either way, a warning. Two new 'F' chips were listed: the Ryzen 7 8700F and Ryzen 5 8400F. The presentation didn't reveal much, but the assumption is this is a Phoenix APU with the graphics turned off. The 8700F likely runs eight CPU cores, and is based on the Ryzen 7 8700G APU.

There's no 8400G equivalent in AMD's current product range, but the 8400F is probably a six-core model. At the right price, both could be appealing. However, the Phoenix APU isn't that great when paired with a discrete GPU, thanks to lower clockspeeds and half the L3 cache (16MB down from 32MB) versus a plain old eight-core Ryzen 7 7700 CPU.

In other words, these APUs will need quite the discount to the regular Ryzen 7000 desktop alternatives to make sense. If they're just $10 or $20 cheaper, you may as well go full-fat Ryzen. If you can get the 8700F for $250 or less, it might just make sense.

40 inches and 5K2K is just about perfect for productivity.

Editor's Pick: Philips 40B1U6903CH

There's a lot going for this ultrawide monitor

WHAT'S THE PERFECT premium productivity monitor form factor? 34-inch ultrawide? Too small. 32-inch 4K? Better, but arguably a bit constraining. What about one of those 49-inch 32:9 aspect monsters? Interesting, but they're a little too low-resolution for the panel size at 5,120 by 1,440 pixels. That's not the stuff of a high-DPI experience.

Then there's the Samsung Odyssey Neo G9, the 57-inch dual-UHD freak. That's a nice display, but apart from clocking in at $2.5K, it's almost too much of a good thing. The 7,680 by 2,160 native resolution can be problematic, especially at its 240Hz spec.

For now, then, the answer might be the Philips 40B1U6903CH. It's a 40-inch 21:9 model, so still pretty massive. It also sports a native 5K2K resolution or 5,120 by 2,160 pixels at the same density as a 32-inch 4K monitor. This is essentially a 32-inch 4K panel with some extra width.

Once you get used to being able to easily run three application windows side by side, even a 32-inch 4K monitor feels a bit constraining. That heightened pixel density over the various 49-inch 1440p models is a real boon, too.

As for image quality, this is a fairly middling IPS item. While the panel will process an HDR10 signal, there's no true HDR support or local dimming, and the brightness tops out at 300 nits. Likewise, the refresh rate tops out a 75Hz, and the response is rated at 4ms. So, this monitor isn't designed for gaming, even if it'll make a decent fist of that remit, provided you have a super powerful graphics card that can cope with the 5K2K resolution.

The lack of HDR support won't matter for most computing tasks, but the 300 nits brightness might. It's plenty for most workflows, but if you like a really bright and punchy panel, it could be marginal.

In all other respects, this is a productivity powerhouse. It has Thunderbolt connectivity with 90W of power delivery, plus full daisy-chaining support and plenty of USB ports. There's also a KVM switch, making it easy to share this display across multiple PCs, plus a 5MP webcam (roughly 3K in resolution terms). You can use this display as a glorified single-cable dock for your laptop at the same time as having your desktop PC hooked up. It really is a pleasant way to get computing done.

As for downsides, there are very few. The audio-out port is located on the side of the panel, which is fine for headphones, but less optimal for long-term usage connected to some desktop speakers. The 5K2K resolution can also be a bit problematic if you're sharing this screen across a PC and Mac.

It's fine for almost any modern PC. You're also good to go with a Mac running Pro, Max, or Extreme Apple silicon. The problem comes with entry-level Apple 'M' silicon. They top out at 6K resolution support. That's fine for running full native, but because of the way MacOS does its high-DPI UI scaling for non-native by essentially doubling the render resolution before scaling it down, if you want to run below native, but with a high-DPI experience, you have to step down to a virtual resolution of 3,072 by 1,296, which robs you of a load of desktop space.

That's a niche objection, but worth noting for MacBook Air and Mac Mini users. Otherwise, this is one heck of a productivity machine, and one that I'm sorely tempted to drop my own money on. That's about as strong a recommendation as any reviewer can possibly make. **–JL**

$1,499, www.philips.com

Reviewed...

Nvidia GeForce RTX 4080 Super

The $999 card that's really $1,200

OH, NVIDIA. This hasn't aged well. For those of you who are unaware, all of Nvidia's RTX 4000 Super series cards had a slightly different strategy towards their launch. The RTX 4070 Super touted a 25 percent internal componentry increase on average (CUDA cores, ray tracing units, ROPs, the lot). Pricing remained static. The RTX 4070 Ti Super received a 10 percent componentry increase, but more importantly, jumped from 12GB of VRAM up to 16GB, with pricing once again remaining the same. Lastly, there was the RTX 4080 Super. By far the least impressive of the three, it received only a meager five percent internal hardware bump. More importantly, it came with a $200 price drop on its recommended retail price. That would theoretically bring it all the way down to a humble $999—sub $1,000 at last, for an 80 series card. Huzzah!

Fast-forward two months, and we've hit a major roadblock: there's no stock. Cards that were debuting at $999 or slightly over (being overclocked variants with aftermarket PCBs) have skyrocketed in price by, you guessed it, $200-300 across the board, and all of the 4080 Supers that have been listed at $999 or thereabouts are now out of stock, on back-order, or just no longer available, even on Nvidia's own website. That's a hard pill to swallow, and makes that five percent hardware bump look pretty bad.

That problem extends even further when you look at the card from a performance perspective. The majority of journalists who received the RTX 4080 Super for launch were pretty underwhelmed with its overall performance. Take us, for example. We initially tested and compared it against an RTX 4080 Aero OC from Gigabyte. A slightly overclocked card, it did, however, have around a 6.5 percent overclock overall. On the other hand, the Founder's Edition we were given only had that five percent bump in internal hardware, and clock speeds were relatively the same. So in performance across the board, it remains identical, if not worse in a lot of our testing scenarios.

It was so poor that we paused our testing initially, and queried Nvidia about it, only for them to confirm that this was to be expected. Nonetheless, regardless of whether you're trying to pick up an RTX 4080, its Super-kin or an OC variant of either, you're going to be looking at $1,200 right now; that's the sad reality.

In isolation, it's an incredible piece of engineering. Complete with CUDA cores, 16GB of VRAM, and some chunky 2.76 GHz clock speeds, it's not to be trifled with. It's cool too, topping out at 77.3 C under-load. As for the frame rates, across our five testing titles at 4K, it averaged 102 fps. That includes some seriously aggressive benchmarks in the form of *Assassin's Creed: Valhalla*, *Cyberpunk 2077* (with ray tracing and upscaling), and *Total War: Three Kingdoms*, too. At 1080p, it nailed 206.6 fps. Those are some chunky numbers for top-tier AAA titles. If you're looking for stellar performance, this card certainly offers it.

It's an interesting GPU, and an incredible piece of design, packing some potent performance. If you can get it at $999, it's incredible. But right now, it's no different to the RTX 4080 at launch; overpriced and out of stock. **-ZAK STOREY**

BENCHMARKS

	Nvidia GeForce RTX 4080 Super	Nvidia GeForce RTX 4070 Ti Super	Nvidia GeForce RTX 4070 Super
3D Mark: Speedway (Index)	**7,309**	6,093	5,024
3D Mark: Port Royal (Index)	**17,948**	14,893	12,859
Max Power Draw (Watts)	582.8	531.2	**502.3**
Max Temperature (Celsius)	77.3	**72**	78.9
Total War: Three Kingdoms @ 1080p (avg fps)	**222**	184	173
Final Fantasy XIV @ 1080p (avg fps)	**260**	151	211
Far Cry 6 @ 1080p (avg fps)	**174**	151	136
Assassins Creed: Valhalla @ 1080p (avg fps)	**195**	163	152
Cyberpunk 2077 @ 1080p (avg fps)	**182**	145	132
Total War: Three Kingdoms @ 4K (avg fps)	**73**	71	53
Final Fantasy XIV @ 4K (avg fps)	**145**	128	110
Far Cry 6 @ 4K (avg fps)	**112**	102	83
Assassins Creed: Valhalla @ 4K (avg fps)	**103**	86	74
Cyberpunk 2077 @ 4K (avg fps)	**76**	61	55
Avg fps @ 1080p	**206.6**	173.6	160.8
Avg fps @ 4K	**101.8**	89.6	75
Avg fps per $ spent @ 4K (Index)	0.08	0.11	**0.13**

Best scores in bold. Our test bed consists of an Intel Core i9-14900K, 32GB of Corsair Dominator Titanium @ 7200, Corsair H150i AIO, and an Asus Z790 Dark Hero. All tests performed at 1080 & 4K, avg fps recorded, RTX & DLSS is enabled in Cyberpunk. Power Draw and Temperature recorded during Port Royal benchmarking. Avg fps per $ spent for RTX 4080 Super set to $1200.

VERDICT **7**

Nvidia GeForce RTX 4080 Super

+ SO GOOD Incredible 4K performance; Super cool; Great design.

- YOU CAN'T BUY IT Meager upgrade over RTX 4080; Not available at RRP; Out of stock.

$999, www.nvidia.com

SPECIFICATIONS

Architecture	Ada Lovelace
Manufacturing Process	TSMC 4N 5nm
CUDA Cores	10,240
ROPs	112
RT Cores	80
Tensor Cores	320
Memory	16GB GDDR6X
Memory Bus	256-bit
TDP	320W

By Grabthar's hammer, what a savings...

2TB Crucial T705 M.2 PCIe 5.0 SSD

The world's fastest PCIe 5.0 drive?

WELL, FOLKS, that's it. The PCIe 5.0 standard and its SSDs have been around for almost a year, and we've finally hit peak sequential throughput with the Crucial T705. That's a phenomenal achievement in and of itself.

The latest SSD from Crucial, the T705, is a slightly tweaked variant of the T700. It still features the same E26 Phison controller and the same 232-layer TLC NAND from Micron at its core, but all of it is ramped up to 11, with a mixture of binning and firmware adjustments necessary to hit those 14GB/s numbers. That's no small feat to achieve, and just goes to show how rapidly SSD tech is improving year on year.

The T705 is available in two primary variants: with or without a heatsink. You can also pick one up as a 'limited edition' white heatsink version, too. As for capacities, they range from 1TB to 4TB respectively, with the 2TB being the best performing drive, topping out sequentials at 14.5 GB/s read versus the 1TB's 13.6 GB/s and the 4TB's 14.1 GB/s. IOPS are also slightly improved for the 2TB variant, reaching 1,550K.

Endurance rating is fairly solid at 600 TBW per TB on the drive (our 2TB drive has that at 1,200 TBW respectively), combined with a five-year warranty. The heatsink is nothing if not chunky, though. Very similar in design to the T700, it still has that massively tall block of aluminum, complete with thick dense fins to encourage suitable airflow. One thing to note: unlike the T700, you can no longer remove the heatsink simply by taking off the small T5 torx screws on either side of the heatsink. If you do intend to take it off, we highly recommend buying the stock variant instead, and saving yourself some cash as you do so.

Speaking of price, it's not entirely unreasonable that you can pick one up for around $282 or so. It's had a huge drop in price since its debut only a month ago (at time of writing), dropping all the way down from $400 for the 2TB variant. Because of that, it's actually really reasonable, particularly with the performance you get for that outlay.

In our testing, Crystal Disk Mark reported QD32 Sequentials at 14,027 MB/s read and 12,280 MB/s write—not far off those advertised numbers. Access time was high, too, with a read speed of 0.015 ms. Under load, with the integrated heatsink, it topped out at around 74 C or so, making it fairly mid-range on the cooling front, all things considered, certainly among the 5.0 drives, with only the Aorus Gen5 12000 hotter at 78 C.

Random 4K performance, however, was somewhat disappointing, with Crystal Disk Mark reporting 86MB/s read and 303MB/s write. That's slower than Gigabyte's Aorus Gen5 12000 and Crucial's own T700 Pro, and more notably the T500 and the Kingston Fury Renegade we reviewed in this very issue.

We're seeing that happen a lot—sequentials do seem to be the driving force behind the majority of drive development. On the one hand, it makes sense, as performing large dataset file transfers takes the longest. On the other, game loading times are similar, often associated with Random 4K performance, and are now at a point where we're talking just a few seconds, even with read and write speeds at 86 and 303 MB/s. Still, holding back has a knock-on effect, both in terms of more advanced game development and the trickle-down tech effect limited to sequential performance.

So then, how do we wrap up a review like the T705? Well, it's seriously quick. Combine that with a decent heatsink, solid pricing (finally), and fairly decent all-round performance, and it might just be one of the best drives out there to date. Where do we go from here? Our guess: 300-layer TLC, a Phison E27 controller, PCIe 6.0 and 30 GB/s drives. Let's just hope those Random 4K figures go with them. **–ZAK STOREY**

BENCHMARKS

	2TB Crucial T705 PCIe 5.0 M.2 SSD	2TB Gigabyte Aorus Gen5 12000 PCIe 5.0 M.2 SSD	2TB Crucial T500 Pro PCIe 4.0 M.2 SSD	2TB Kingston Fury Renegade PCIe 4.0 M.2 SSD
AS SSD Sequential - Read / Write (MB/s)	**10,064** / 9,627	8,970 / **9,948**	5,631 / 4,455	5,598 / 3,970
AS SSD Random 4K - Read / Write (MB/s)	84.50 / 273.12	**86.65** / **289.12**	81.37 / 277.75	85.21 / 281.32
AS SSD Access Time (ms)	**0.015** / 0.045	0.017 / 0.037	0.018 / 0.017	0.062 / **0.016**
CrystalDiskMark Sequential QD32 - Read / Write (MB/s)	**14,027** / **12,280**	12,353 / 11,598	7,879 / 6,783	6,979 / 6874
CrystalDiskMark Random 4KQ1 - Read / Write (MB/s)	86 / 303	89 / 310	**92** / 331	88 / **357**
Max Temp Under Load (C)	74	78	72	**55**
Gigabyte per $ (GB)	7.09	7.41	**12.82**	10.71
Sequential Read MB/s per $ (MB/s)	49.74	45.75	**50.51**	37.37

Best scores in bold. Our test bed consists of an Intel Core i9-14900K, 32GB of Corsair Dominator Titanium @ 7200, an Nvidia GeForce RTX 4080, Corsair H150i AIO, and an Asus Z790 Dark Hero. Max Temp recorded via HWMonitor during benchmarking process.

VERDICT **9**

2TB Crucial T705 M.2 PCIe 5.0 SSD

CRUCIAL Top-tier sequential performance; Solid heatsink; Aggressive pricing; Good access time.

OPTIONAL Random 4K beaten by 4.0 SSDs; Heatsink non-removable.

$282, www.crucial.com

SPECIFICATIONS

Variants	Heatsink, Normal, Limited Edition
Form Factor	M.2 2280
Interface / Protocol	PCIe 5.0 / NVMe
Flash Memory	232-Layer TLC NAND Flash
Sequential Read	14,500 MB/s
Sequential Write	12,700 MB/s
Random Read	1550K IOPS
Random Write	1800K IOPS
Endurance (TBW)	1200
Warranty	5 Years Limited Warranty

So we've hit the PCIe 5.0 bandwidth limit already—now what?

Asus Zenbook Duo OLED 2024

Two screens are better than one with Asus's origami special

EVERY NOW AND THEN, something comes along that challenges your expectations about how a product will work. We're used to opening the lid of a laptop and using it like an L-shaped thing with a keyboard that's too close to the screen, but what if that keyboard came off and revealed another screen underneath?

There have been attempts at double-screened laptops before, but none have got it as right as this one. The pair of 14-inch OLEDs the Zenbook Duo is equipped with are capable of putting out almost 400 nits of brightness, are both touchscreens, and support a stylus. The stand that folds out behind the bottom screen is capable of holding them in landscape and portrait orientations, and despite the fact that unlike a folding phone, they are at all times two screens rather than one, the ability to have more app windows visible makes for a powerful productivity tool.

Inside, there's one of Intel's new Core Ultra 9 processors, with six performance cores (which can boost up to 5.1GHz), eight efficient ones, and two of the low-power efficiency cores that appeared with the Meteor Lake architecture. It pulls a base power of 45W, and has Intel Arc integrated graphics featuring eight Xe cores, ray tracing, and 8K output. It's not an obvious gaming laptop, though it manages over 100fps in the 3DMark Night Raid benchmark that's designed to test integrated graphics. It fared less well in Time Spy Extreme, a complicated DX12 scene that was probably unfair to ask it to complete—it managed 10fps across the tests, but did better in older games, especially at 1080p.

Our review model comes with 32GB of DDR5X RAM and a 2TB NVMe SSD. There's also a 65W USB-C charger in the box, and the Asus Pen 2.0 stylus. The keyboard and trackpad module sits on top of the second screen and is removable, connecting to the main part of the machine via Bluetooth once it's detached, but there's nothing stopping you using your wired or wireless devices. A pair of Thunderbolt 4 ports, plus a USB 3 Type-A, are found on the edge of the laptop, along with an HDMI 2.1. The keyboard module has a Type-C port for charging, and one of the Thunderbolt ports for charging the main body of the laptop. When attached, the keyboard charges via a dedicated connector. Battery life is hard to gauge, as it depends on how many screens you're using, but the 75Wh unit in the dual-screened machine acquits itself well.

This is an exciting step forward in laptop engineering. While it's not a gaming laptop, it allows remote workers to do more without needing a portable monitor, and the combination of the Core Ultra 9 and lots of RAM is potent if you want to use creative software and a web browser across two screens, or even on a third screen attached to the HDMI port.

The price is high, however, for something you could approximate with a portable monitor and a Bluetooth keyboard, but the Zenbook Duo is amazing to use. While it feels like a new generation of laptop, we can't wait to see what happens when smartphone-style folding screens become available in laptop sizes. **–IAN EVENDEN**

VERDICT **9**

Asus Zenbook Duo OLED 2024

DUO Twin-screened laptop marvel with top-quality internals.

BOO-O Expensive; Perhaps has only a niche appeal.

$1,499, www.asus.com

SPECIFICATIONS

CPU	Intel Core Ultra 9 185H
GPU	Intel Arc Graphics
RAM	32GB
Storage	2TB SSD
Screen	2x 14in 2880x1800 120Hz OLED
Connectivity	Wi-Fi 6E, Bluetooth 5.3, 1x USB 3.2 Gen 1 Type-A, 2x Thunderbolt 4, 1x HDMI 2.1, 3.5mm combo jack
Dimensions	12.3 x 8.5 x 0.8 inches
Weight	3.6lbs

Having two OLED screens means the battery can drain fast, but it's worth it.

HP's latest Omen laptop is big on portability, but the battery life is a bit borked.

HP Omen Transcend 14

Surprisingly affordable, but the battery life is a bummer

THE NEW HP Omen Transcend 14 gaming laptop is properly impressive. Usually, gaming portables come with a painful price tag. Not so with this beauty. It's still more expensive than some RTX 4060 gaming laptops, but you won't find many this compact, nor with such a lovely design and gorgeous, high-res OLED screen attached.

The standard sticker price for this model is $1,660 with an RTX 4060 GPU. It's a no-brainer at just $70 more than the RTX 4050 option. That's even more true at the $1,360 we've seen it offered for. That's a great price for this little machine.

Anyway, the Intel Core Ultra 7 155H chip at its heart is a decent 16-core, 22-thread device, which will turbo up to 4.8GHz and can mostly keep up with the likes of the AMD Ryzen 9 7940HS and 7840HS chips in either the Razer Blade 14 or Asus ROG Flow X13. It'll also keep the Nvidia RTX 4060 fed with enough data to keep the GPU running well.

The catch is that the RTX 4060 is just a 65W variant, which definitely limits performance. That's potentially problematic, given the 2,880 by 1,800 resolution of the lovely 120Hz OLED screen. That's a lot of pixels for the small Nvidia Ada GPU to cope with, especially with the low power limit.

The rest of the backup spec is pretty standard, including 16GB of dual-channel LPDDR5x-7500 memory and a 1TB NVMe SSD. In performance terms, it's pretty capable. You get great frame rates even in the latest games at 1080p, especially if you take advantage of Nvidia's DLSS and Frame Generation features.

While you can even get away with bumping up to native resolution in a few less demanding titles, that's not often the case. But the pixel density and speed of the OLED panel means that even if you are scaling down, you get a pretty sharp picture.

So, the HP Omen Transcend 14 is effective, even if it's not the fastest option out there. But that isn't really the aim. HP has kept the internals in check to the point where it doesn't get too toasty on the lap, and it also isn't tough on the ears. Even running the machine in its full power gaming mode doesn't offend in terms of fan noise. You could happily game on this system on public transport without feeling like a total heel—as long as you're plugged into the mains, that is. The battery itself is rated at 71Wh, but it's the off-mains performance of the Transcend 14 that really lets it down. Battery life is important for these kinds of compact laptops, but the PCMark Gaming battery life test result is just 58 minutes.

It's not just the gaming battery life that disappoints, either. Using it as an office machine, it only lasts until lunchtime, even just for writing and editing, and that's on the balanced power mode, without the screen set to full brightness. An all-day laptop this clearly is not.

The mostly lovely, svelte chassis is a bit of a worry, too. It feels beautifully machined and there's practically no give to the keyboard. But if you pick it up by the corner with the screen open, you'll hear a creaking noise from the chassis. There's also quite noticeable give on the lid behind that gorgeous OLED panel.

Of course, at this price that's probably tolerable. The likes of a 14-inch Razer is more solid. But then, it's at least $600 more expensive, too. If only it had a better battery life—this HP would be the sort of machine you'd definitely want in your life. It's otherwise a fantastic laptop with a decent level of gaming performance, a gorgeous OLED screen as standard, and impressive affordability. **–DAVE JAMES**

VERDICT 8

HP Omen Transcend 14

MOVIN' UP Fabulous OLED screen; Very portable; Decent gaming performance.

JOG ON Disappointing battery life; Slightly creaky chassis.

$1,360, www.hp.com

SSD BENCHMARKS

	HP Omen Transcend 14	Asus ROG Flow X13	Razer Blade 14
Cyberpunk 2077 1080p (fps)	20	23	**38**
F1 2022 1080p (fps)	45	44	**73**
Metro Exodus 1080p (fps)	47	45	**77**
Cinebench R23 (points)	15,225	**16,204**	15,115
X264 video encoding (fps)	44	**48**	**48**
PCMark 10 gaming battery life (minutes)	58	**134**	101

Best scores are in bold. Our test bench consists of an AMD Ryzen 9 3900X, a Gigabyte X570 Aorus Master, an Nvidia GeForce GTX 1080, 32GB of Corsair Vengeance RGB Pro DDR4, and a Corsair Neutron 240GB OS SSD.

SPECIFICATIONS

CPU	Intel Core Ultra 7 155H
Cores | Threads	16 | 22
Memory	16GB LPDDR5x-7500
GPU	Nvidia RTX 4060 (65W)
Screen size	14-inch
Screen type	OLED
Native resolution	2880 x 1800
Refresh rate	120 Hz
Storage	1TB WD SN810 SSD
Battery	71 Wh
Warranty	1 year

Iiyama GCB3480WQSU-B1 Red Eagle

The gaming goliath you're after?

WE'VE SEEN ultrawide monitors before, but they never fail to impress. This model from Iiyama offers a good-sized 34-inch IPS screen that's about as big as one and a half 28-inch widescreen displays, boasts nearly 5 million pixels, and brings with it some impressive brightness.

The curve in the screen means that not only is the box it's delivered in bigger than you'd expect, but it takes up more space. We're not entirely sold on the benefits of a gentle arc in your monitor—it means it can't sit flat against a wall, and that images in the corners are larger in your vision than those in the center. After decades of perfectly flat screens, having one that bends toward you can be a little disconcerting.

Pretty soon, however, it becomes second nature, and when playing games, the immersive nature of both the wider screen and the curve become apparent. This Red Eagle monitor has a few features that gamers will appreciate, with a 180Hz refresh rate and HDR compatibility, along with a maximum brightness of almost 450 nits. Games that can play to all its strengths are a bit scarce, however.

Creative app users will appreciate the color response, with 100 percent of sRGB, 88 percent of Adobe RGB, and 95 percent of P3 measured in our tests. While this means the colors you set in apps like Photoshop and InDesign are more likely to be accurately represented, for the rest of us it means a well-saturated picture that's pleasing to the eye.

There are four inputs at the back of the screen, and while there's no picture-in-picture mode so you can have a streaming stick playing a movie at the same time as you play a game, it does have automatic switching between inputs. There's no single-cable USB-C connection, so laptop users won't be able to charge while sending a video signal, but there is a small USB hub that converts a USB 3 Type-B connection from your PC to two Type-A ports, one of which is a quick-charger, allowing you to connect a keyboard dongle and maybe your cellphone.

Having more than one monitor has become popular, given the number of HDMI and DisplayPort outputs that hang out of our graphics cards, but an ultrawide monitor like this takes things in the other direction. A single ultrawide screen can simplify your setup, removing the need to have more than one screen and making good use of Windows' app snapping and tiling options. While it's not ideal for maximizing a single app—with the possible exception of video editing or graphic design software, where you want a large view of what you're doing—it still makes for a powerful productivity tool.

There's a VESA mount on the back of the monitor, onto which you can connect an arm or stand. This splayed V-shaped foot uses actual screws to connect, instead of the spring-loaded clips we've seen elsewhere, with a small screwdriver in the packaging to assist.

We're not sure that this will be for everyone, but if you do make the leap, this package has a lot to offer. The picture is sharp, bright, and colorful, there's plenty of connectivity so you can switch between devices, and the USB hub adds an extra layer of convenience. If you like what you see, and have the desk space, you won't be disappointed. **–IAN EVENDEN**

VERDICT **9**

Iiyama GCB3480WQSU-B1 Red Eagle

➕ **ULTRAWIDE** Good brightness and color; Sharp picture; Lots of inputs.

➖ **ULTRAHIDE** USB-C and picture-in-picture would have taken it to another level.

$539, iiyama.com

SPECIFICATIONS

Screen	34in 3440 x 1440 (21:9) IPS, 180Hz
Inputs	HDMI 2.0 x2, DisplayPort 1.4 x2, USB hub, headphones
HDR	Yes
Speakers	2x 3W
Height adjustment	130mm
Dimensions	31.8 x 21.7 x 10in
Weight	16.7lbs with stand

© UNSPLASH

Color reproduction is excellent, making it a good choice for games or design.

NZXT H6 Flow

Possibly the best budget chassis out there

OKAY, LET'S GET THIS OUT of the way. First and foremost, the H6 Flow is one of the best-value cases you can buy. If you're after a mid-tower, dual-chamber box for your next build, the Flow has you covered. For the money, it's just incredible what you can fit in this thing, and although there are a few foibles here and there, on the whole, there's very little to complain about. In fact, we'd go so far as to say that it even rivals the prestigious Corsair 4000D Airflow, which is one of our all-time favorite cases.

Why is it so good? Well, the real star of the show is the cooling. As standard, it will fit three 120mms in the front right, three 120mms in the roof (where your AIO will sit), two 140mms buried in the floor, and one 120mm in the rear. But it's that front-right segment that's the true champ. Positioned at a 45-degree angle, it's incredibly well-implemented. As standard, the case comes with three 120mm fans installed there, pre-routed through cable grommets, however, the big win is that angled design. Because most fishbowl cases are display pieces, typically sat on desks, these side vents (often populated with triple 120mms) are usually pushed up against the wall, with little airflow space, stifling your intake. Because the H6 Flow has those fans at an angle, no matter where you place it, it gives them far better access to air than you'd otherwise find in the competition, giving you a good healthy source of cool air for your rig.

On top of that, like most of the fishbowls out there, it also features two 140mm fan slots in the bottom of the case, giving you a direct cool intake feed to your GPU. In the H6 Flow's case, those are quite tight when it comes to fitting fans in there, but as long as your fan of choice conforms to the standard form factor, you'll be fine, and the super-tight fit looks incredible.

As standard, it supports ATX motherboards and power supplies, and has plenty of space around the back for cable management and secondary devices, along with support for two 2.5-inch drives and one 3.5-inch drive. It manages all this while keeping an incredibly tight footprint at just 17.13 x 11.3 x 16.34 inches, too.

It even supports cards up to 365mm or 14.37 inches in length as well, which should fit pretty much any of the modern-day GPUs, even including true monsters like Gigabyte's RTX 4080 Super Aero OC, which tips the scales just shy of 342mm.

It's not without fault, however. Some of the cable management, particularly surrounding those triple 120mm fans in the front-right and 140mm fans in the floor is incredibly tight, and once you've got your AIO installed, access to the top-most connectors on your motherboard can be a challenge. But those are minor issues that can be addressed by tweaking your build process.

It's the price, however, that seals the deal. Right now, you can pick up an H6 Flow in white for just $101, or the black variant for $110. NZXT will also sell you an RGB variant as well, which maxes out at $135, but swaps those three standard 120mm fans out for RGB models instead. At this price, there's no poor material choices or design corners cut; it's just pure, simple, elegant sophistication, and we just can't get enough. The only thing that would make it better (particularly given the number of fans you can install in it) is an integrated fan controller. Otherwise, this is a classic. **–ZAK STOREY**

VERDICT

9

KICK ASS!

NZXT H6 Flow

FLOW STATE Remarkable cooling; Clean aesthetic; Solid clearances; Great pricing; Good cable management.

DOESN'T FLOAT MY BOAT Needs a fan controller; Might need to think about your build order.

$110, www.nzxt.com

SPECIFICATIONS

Motherboard Support	ITX, Micro-ATX, ATX
2.5-inch / 3.5-inch Support	2x 2.5 / 1x 3.5-inch
Max Radiator Support	360mm roof, 120mm rear
Fan Support	3x 120mm front-right, 3x120mm / 2x140mm roof, 2x 140mm floor, 1x 120mm rear
Dimensions	17.13 x 11.3 x 16.34 inches
Graphics Card Clearance	14.37 inches
CPU Tower Clearance	6.42 inches
Warranty	Two Years

Go with the NZXT Flow.

SUBSCRIBE TODAY

Find your next great read

Subscribe today and get instant access on your iPad, iPhone or Android device.

www.magazinesdirect.com/B3XMPC

2TB Kingston Fury Renegade PCIe 4.0 M.2 SSD

A little variation goes a long way

SSDS ARE, on the surface, quite boring. They don't have the graphical fidelity of a GPU, they can't render 4K video files, or act as the beating heart of your machine like a CPU, and they aren't as critical to your system's operations as RAM. In fact, you can get away without using one of them at all. A simple HDD will allow you to run a machine and do all of those things you know and love, just more slowly.

So why have we've got so excited over this drive? Surely, at its heart it's just another boring SSD with slightly better sequentials, right? Well, yes, but also no.

Kingston has produced some fantastic drives, but in the world of cutting edge PCIe 5.0 storage and other marketing buzzwords, it doesn't have the same clout as some of its competition. Gigabyte, Corsair, and Crucial are all seemingly at the cutting edge.

The thing is, most of these drives are pretty much carbon copies of one another. Take the Gen5 SSDs, Gigabyte's Aorus Gen5 12000, the Seagate FireCuda 540, the Crucial T700, and the Corsair MP700. All feature the same Phison E26 controller, 232-layer TLC NAND from Micron, LPDDR4 cache, and in some cases, the same cooler. The differences come in slightly tweaked firmware, voltages, and chip binning.

That's led to very similar performance, if not on the sequential front (where everyone's trying to tell you that only sequential performance matters), then most certainly on the Random 4K side. It's that Achilles heel that our little Kingston Fury Renegade suitably sidesteps.

At its heart, the Fury Renegade has been built for the PS5, but what's unique is how Kingston has implemented the hardware. The Fury Renegade has a Phison PS5018-E18 controller, coupled alongside a 176-layer TLC NAND from Micron, and a full-sized DDR4 chip as its cache. Its closest competitor is the Crucial T500. It packs a Phison E25 controller, 232-layer TLC NAND, and LPDDR4 cache chip.

The difference comes in the form of endurance—the Renegade is rated at five years for 2000TBW. On top of that, this drive absolutely rips with random 4Ks, smashing the T500 in the Random 4K writes by a good 26 MB/s. Its reads are 3MB/s behind, but when you consider that it's using an older controller and lower-density NAND, it really does stand out. Sequentials are lower, and the price is slightly higher (clocking in at $187 for 2TB without a heatsink), but it's that full-fat DDR4 that really does sink the victory for Kingston here.

The real kicker is that those Random 4K figures beat every other drive we've tested, even the PCIe 5.0 ones and the latest Crucial T705. It manages all of that while staying exceedingly cool, too, topping out at a comfortable 55 celsius.

If gaming and system operation is your primary focus, and transfering large gobbets of photos and video isn't, the Renegade is phenomenal. PS5, PC, it doesn't matter; this rips. –ZAK STOREY

VERDICT **9** KICK ASS!

2TB Kingston Fury Renegade PCIe 4.0 M.2 SSD

ONE PUNCH MAN Phenomenal random 4K performance; Low temperatures; Incredibly well priced; Massive endurance rating.

SUPERMAN Slightly more expensive than the competition; Sequentials are slightly slower.

$187, www.kingston.com

BENCHMARKS

	2TB Kingston Fury Renegade PCIe 4.0 M.2 SSD	2TB Crucial T500 Pro PCIe 4.0 M.2 SSD	2TB Gigabyte Aorus Gen5 12000 PCIe 5.0 M.2 SSD
AS SSD Sequential - Read / Write (MB/s)	5,598 / 3,970	5,631 / 4,455	**8,970 / 9,948**
AS SSD Random 4K - Read / Write (MB/s)	85.21 / 281.32	81.37 / 277.75	**86.65 / 289.12**
AS SSD Access Time (ms)	0.062 / **0.016**	0.018 / 0.017	**0.017** / 0.037
CrystalDiskMark Sequential QD32 Read / Write (MB/s)	6,979 / 6874	7,879 / 6,783	**12,353 / 11,598**
CrystalDiskMark Random 4KQ1 Read / Write (MB/s)	88 / **357**	**92** / 331	89 / 310
Max Temp Under Load (C)	**55**	72	78
Gigabyte per $ (GB)	10.71	**12.82**	7.41
Sequential Read MB/s per $ (MB/s)	37.37	**50.51**	45.75

Best scores in bold. Our test bed consists of an Intel Core i9-14900K, 32GB of Corsair Dominator Titanium @ 7200, an Nvidia GeForce RTX 4080, Corsair H150i AIO, and an Asus Z790 Dark Hero. Max Temp recorded via HWMonitor during benchmarking process.

SPECIFICATIONS

Variants	Heatsink, Normal
Form Factor	M.2 2280
Interface / Protocol	PCIe 4.0 / NVMe
Flash Memory	176-Layer TLC NAND Flash
Sequential Read	7,300 MB/s
Sequential Write	7,000 MB/s
Random Read:	1000K IOPS
Random Write	1000K IOPS
Endurance (TBW)	2,000
Warranty	5 Years Limited Warranty

NZXT Function 2

Ultra-fast, fully configurable, and relatively affordable

A DO-EVERYTHING keyboard for a price you can afford—that's the idea behind the new NZXT Function 2. For starters, it features dual-actuation, and a choice between a light and responsive 1.0mm key actuation or a solid and full 1.5mm. This may not sound too extravagant, nor is it the first gaming keyboard to offer this, but it can entirely change the sound and feel of the board, which is handy when switching between work and play.

Next up, there's the 8,000Hz polling rate—something rarely seen on a gaming keyboard, especially one selling under $150. That's the kind of number you might expect from a high-end mouse, but not a keyboard, where 1,000Hz is more typical.

What's more, included in the box is not only a keycap removal tool, but also a keyswitch removal tool, as you're able to exchange both the cap and its switch. You get a handful of both 35g Yellow Linear and 45g Red Linear optical switches, so you can go for a heavier, louder-feeling gaming keyboard, or live on the quieter side. It's up to you.

The standard optical switches installed in the board have a 40g actuation force, so you can experiment with lighter or heavier switches under specific keys. It's also commendable that not only are the tools thrown in, but there's also a spare set of caps, which many companies would sell for an additional fee.

The experience is further bolstered by the NZXT CAM software. Through CAM, you can change the RGB lighting (which is per key), and remap your buttons as desired. More crucially, you can switch between the two actuation settings and change the polling rate from 125 Hz all the way up to 8,000 Hz. You can also set macros, as well as disable the Windows and FN keys, and save profiles.

Speaking of the RGB lighting, it's good, but not great. There's a dedicated button that cycles through four brightness settings, but its peak isn't as dazzling as you might hope for, nor as punchy as Razer's Chroma or Corsair's iCue lighting when properly configured. Further demerits involve some of the NZXT Function 2's unconventional features.

FEATURE SET

For example, the volume rocker and dedicated media bar is on the left side of the deck, which is something that's hard to adjust to. I kept reaching towards the top right to be greeted with nothing more than sheet aluminum. It has easily one of the worst wrist rests out there, too. Not only is it slippery, it's also rock hard. Okay, it's magnetic and stays in place, but it desperately needs some padding or texture to make it more comfortable to use.

All that said, gaming on the NZXT Function 2 is a stellar experience. Once you've spent a bit of time setting up, it really feels like a cut above other mechanical decks I've used. Straight away, you'll notice the response in games like *Far Cry Primal* and *The Finals*. Typing feels faster when working compared to my daily driver Razer Huntsman V3 Pro, with which the Function 2 shares very similar DNA, just for a lot less cash.

Then there's the sound dampening, which is now double layered when compared to the original Function. This board is considerably quieter than many other optical keyboards, including the aforementioned Huntsman V3 Pro and the Corsair K70 RGB OPX, which is also more expensive while offering a similar feature set.

Overall, the NZXT Function 2 gets a lot right with its 8,000 Hz polling rate, dual actuation, hot-swappable switches, and solid construction, especially for its sub-$150 price. If you're in the market for a custom keyboard, but don't want to spend too much or dabble in modding, this is a great choice. –ALEKSHA MCLOUGHLIN

VERDICT **8**

NZXT Function 2

PITCH PERFECT Very configurable; Super quiet; Ultra fast.

OFF KEY Weak RGB lighting; Uncomfortable wrist rest.

$140, www.nzxt.com

SPECIFICATIONS

Layout	Full-size (MiniTKL available)
Switch type	Optical
Switches	NZXT Swift
Backlighting	Yes
Anti-ghosting	Yes
N-key rollover	Yes
Discrete media keys	Yes
Connection	USB-C
Weight	910g | 2lb

NZXT Lift 2 Symm

Fast, affordable, and lightweight

DO YOU NEED to splash a ton of cash for a high-end lightweight mouse in the competitive gaming scene? Not so, according to the NZXT Lift 2 Symm. This pointer delivers a no-frills punch. It won't blow you away, but it does exactly what it sets out to.

NZXT has been smart to keep the pricing of the Lift 2 Symm similarly aggressive to its predecessor, while also shifting some of its weight in the process. The original Lift mouse clocked in at 67g, but this new iteration cuts its heft down to just 58g. That hasn't been without sacrifices—the RGB lighting is completely gone for both the Symm and Ergo variants, and the bottom has been hollowed out to its bare essentials.

Short of poking some holes in its back, NZXT has made the lightest mouse it could, and beefed it up with an 8,000Hz polling rate, plus a lightning-fast 26,000 DPI sensor. Those specs are enough to put the frighteners on the best gaming mouse, even without factoring in the price. If you include the very reasonable $50 sticker for both the Symm(etrical) and Ergo right-handed versions, the proposition only gets stronger.

Of course, when you build a gaming mouse down to a weight, you're going to end up with a pointer that feels insubstantial and cheap. The NZXT Lift 2 Symm is no exception. It isn't helped by a strangely textured scroll wheel that doesn't feel fluid when actively scrolling. It's precise enough, but not terribly smooth.

But there's still plenty to like, from the optical switches (rated with an implausible 100-million click lifetime) to the '100 percent virgin' PTFE mouse feet. The latter ensures that this gaming mouse glides over a mousepad with no resistance, which is particularly impressive given that it's a wired pointer. Cabled mice can be tiresome if you're used to wired, but the 2m braided cable gives no pull or drag. You shouldn't need an anchor here.

The Lift 2 Symm, as opposed to the Ergo variant specifically, has a symmetrical shape, which makes it ideal for both the claw and fingertip grips. I'm firmly in the latter camp, and had no issues with guiding the mouse, whether I was working or engaging in some FPS action. Within a couple of minutes of plugging in, you can be scoring headshots on beast and man alike in *Far Cry Primal*, and have little issue in holding your own after a couple close games of *The Finals* and *Overwatch 2* after hours.

One area where the low cost is more apparent involves the software, which is distinctly minimal. You have onboard memory with the Lift 2 Symm able to save custom profiles to the device via the NXZT Cam software. There's a total of five profiles you can copy over, each with five settings apiece and DPIs ranging from 100 all the way to 26,000. In other words, there's plenty of tunability, and you don't need to use the software on your PC to switch profiles. Likewise, you can unplug the mouse, plug it into another machine, and retain all your fine-tuning, which is nice. But all that aside, there are few additional fripperies.

Overall, the NZXT Lift 2 Symm is a great gaming mouse with a competitive price point, excellent polling rate, good sensor, and solid enough design. The brand has clearly gone for function over flashiness, and while this pointer is unlikely to turn many heads, it may just win you a couple of online games. There aren't all that many mice for $50 or less that genuinely deliver, so the Lift 2 Symm is an extremely welcome addition to the market. If you're a touch cash-strapped, but don't want to skimp, the Lift 2 Symm could be the right pick for you. –ALEKSHA MCLOUGHLIN

VERDICT **8**

NZXT Lift 2 Symm

LIGHT FANTASTIC Great price; Excellent performance; Very low weight.

OVERWEIGHT Limited software; Rough-feeling scroll wheel.

$50, ww.nzxt.com

SPECIFICATIONS

Sensor	PixArt PMW3395
Sensor type	Optical
DPI	26,000
Max acceleration	50g
Max speed	650 IPS
Polling Rate	8,000 Hz
Connectivity	USB 2.0
Buttons	5
Ergonomic	Ambidextrous
Dimensions	126.8 x 67.1 x 38.3mm
Weight	58g

When the dragon appears, a new Arisen also appears, their fates intertwined.

ACTION ROLE-PLAYING GAME

Dragon's Dogma 2

A cult classic in the making

I FINISHED *Dragon's Dogma 2* at 7am after an all-nighter. I should have gone to bed. I could barely keep my eyes open. But I did not go to bed. Instead, I immediately started a new game, resetting the story, but keeping most of my items and all of my levels. I then proceeded to play for five more hours before I literally passed out.

Dragon's Dogma 2 is glorious, thrilling, accidentally hilarious, frustrating, and maddening—literally all the adjectives. It's one of the very best RPGs, as well as being a huge pain in the behind, and then some.

You are the Arisen, a soldier killed by a dragon, returned to life despite the absence of a heart. Whenever the dragon appears, a new Arisen also shows up, their fates intertwined. Despite constantly flinging flat characters and stiff dialogue at you—everyone's lines are laden with faux medieval affectations and po-faced seriousness—the story itself is a creative yarn that's so much more elaborate than it needs to be.

Okay, the scripted NPCs are half-baked, but the pawns are a different matter. They're to die for, especially your main boy, Gorbo the goblin. Pawns are your loyal, chatty companions, putting their lives on the line for you every day. Your main pawn who will be with you always, leveling up alongside you.

The world of *Dragon's Dogma 2* is vast, with two main regions, one covered in forests and fields, the other scorching and arid, punctuated by the occasional oasis. Its vastness is not simply down to its literal size, what makes *Dragon's Dogma 2* feel gargantuan is the limitations placed on fast travel.

This is a game made up of thousands of random adventures, often in locations you're not explicitly told to visit, but only a few locations have fixed fast travel points, and you can only use them if you've found a ferrystone. By the end, you'll probably have quite a few, but the system discourages you from using them often.

Exploration also spits out some fun environmental puzzles, mostly relating to how to reach tantalizing chests or dungeon entrances. Mages and sorcerers can levitate, which often proves to be the solution. But you might also need to hitch a ride on a harpy, toss a pawn, or deploy some explosives.

As for the quests, they often forgo explicit instructions, which is where both the thrills and frustrations really kick in. *Dragon's Dogma 2* expects you to engage your brain and actually solve problems. A lot of the time that's great, but you'll also find yourself staring at your monitor at 2am screaming, "Just tell me where to go!"

We needed help from the developer, Capcom, to complete some quests. For most gamers, that's not an option. We still encountered gameplay and quest hitches where there's been no resolution. A second run through the game reveals just how much can be missed when quests don't work as intended. It's so frustrating.

Yet, it's impossible not to find yourself dragged back into the game. The combat system is so phenomenally tactile and kinetic, where each fight becomes a product of countless decisions and environmental factors. At its most basic, that might mean grabbing a loose boulder and chucking it at a dragon. But what if there are no boulders? Well, maybe you should summon a wall of ice, which you can then smash, and now you've got some massive ice cubes to throw around.

So yes, *Dragon's Dogma 2* is janky. Yes, it's going to need some patches to fix the performance issues and bugs, and even then it's going to be a pain in the ass because of some very annoying quests. But we promise that you'll still love this weird, ambitious RPG. **–FRASER BROWN**

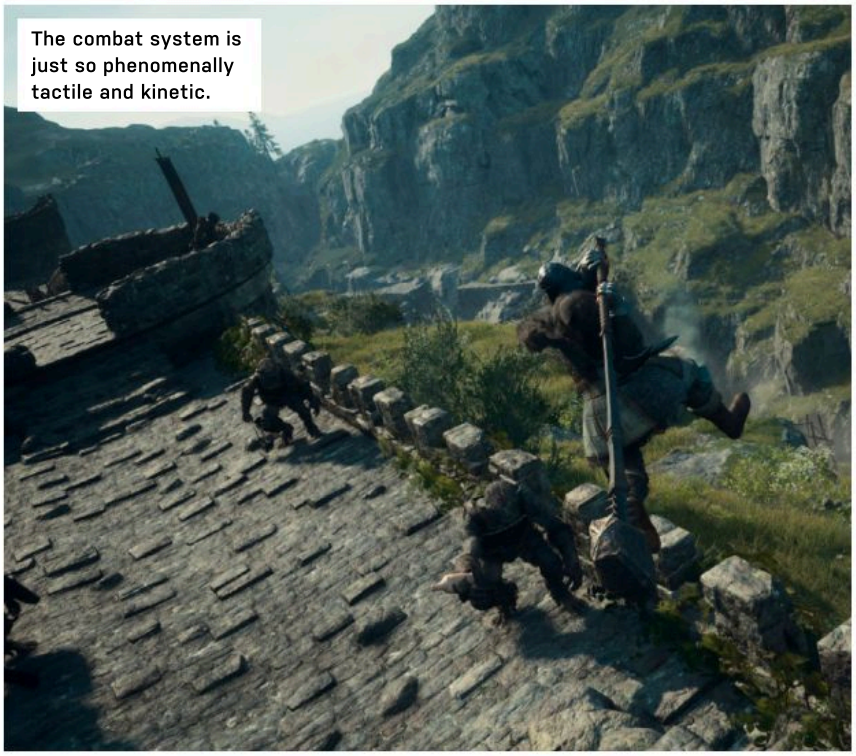

The combat system is just so phenomenally tactile and kinetic.

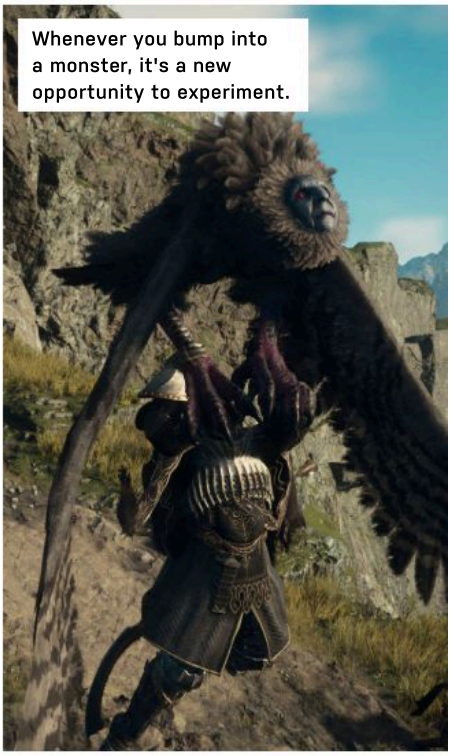

Whenever you bump into a monster, it's a new opportunity to experiment.

Monsters have their own lives, duking it out with NPCs, hunting other creatures, or taking a well-earned nap.

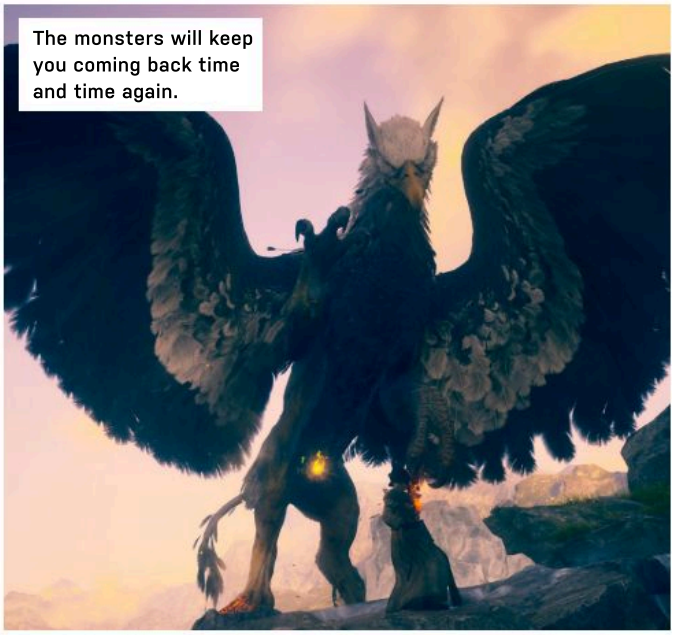

The monsters will keep you coming back time and time again.

Broken quests are a source of constant frustration.

VERDICT **8**

Dragon's Dogma 2

IT WAS THE BEST OF TIMES Great narrative; thrilling combat; vast game world.

IT WAS THE WORST OF TIMES Vague quests; Bugs can make it hard to progress.

RECOMMENDED SPECS CPU, Intel Core i7-10700 or AMD Ryzen 5 3600X. GPU, NVIDIA GeForce RTX 2080 or AMD Radeon RX 6700. RAM, 16GB..

$69.99, www.dragonsdogma.com, M-rated

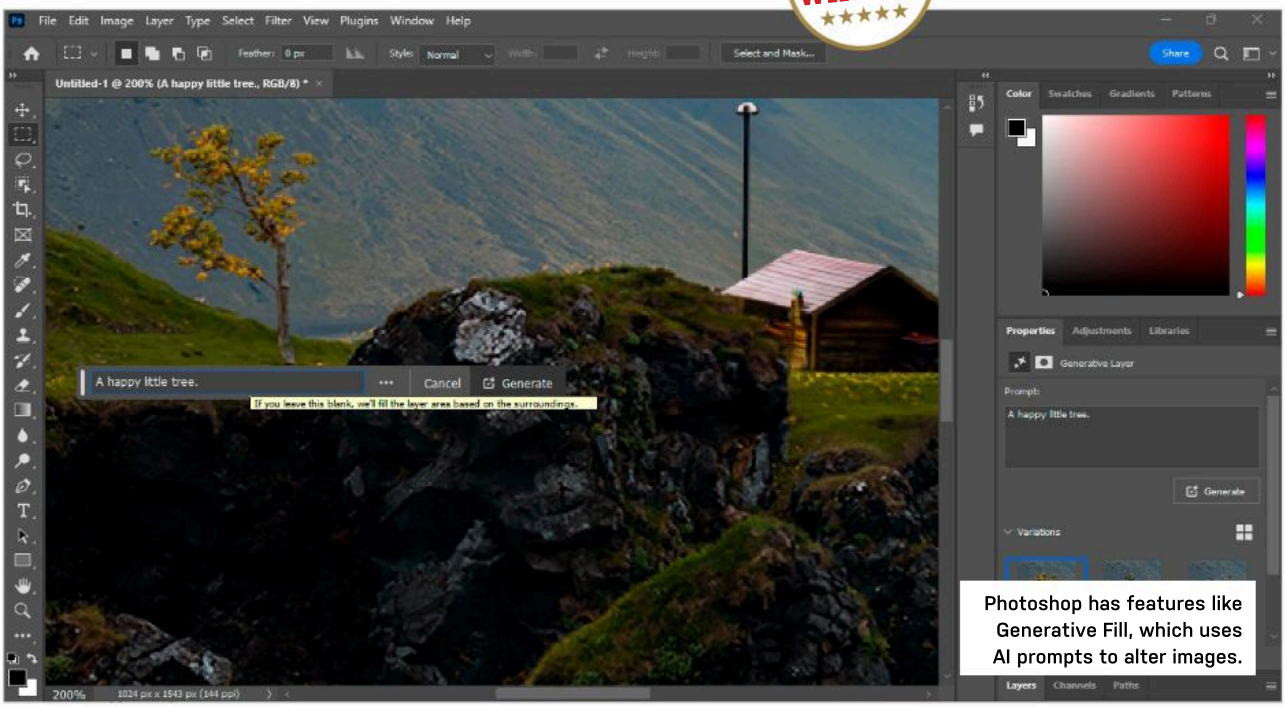

Photoshop has features like Generative Fill, which uses AI prompts to alter images.

Adobe Photoshop vs GIMP

Which picture manipulation app wins?

ADOBE Photoshop and GIMP have some very similar features: they're both raster image editors. Both programs also support virtually every image editing feature you're ever likely to need, including selection editing, layers, alpha channels, scripting, retouching, resizing, HDR, noise removal, and much more.

Photoshop is now a proprietary paid SaaS (software as a service) product. It's used by millions of professional graphic designers every day. In fact, the software has become so ubiquitous that the verb 'to photoshop' was even added to the Merriam-Webster dictionary in 2008, though the first use of the term was in a Usenet group in 1992, two years after the software was first released.

GIMP's initial public release was in 1996. Although the name was originally an acronym for 'General Image Manipulation Program', after meeting with Richard Stallman, the developers agreed to allow the program to form part of the GNU software collection, hence a slight name change from 'General' to 'GNU'.

The program is developed by volunteers and released as open-source. The upshot is that GIMP costs nothing to download and is cross-platform, available for Windows, macOS, and Linux.

The GIMP community also maintains a variety of plugins, such as darktable, which allows you to develop and enhance raw images. Traditionally, these were easy to search for and download via the official GIMP Registry (www.gimp.org/registry), but due to a lack of volunteer coders, the site died, and has yet to be resurrected.

Photoshop also supports a number of plugins, which you can browse via the software's own 'Plugin Marketplace'. These include popular choices like the powerful 'Nik Collection' DxO Labs, which offer enhanced image editing features like taming noise and correcting perspective issues—for a price.

PHOTOSHOP

Like most of its plugins, Photoshop is proprietary software. It has come a long way since the turn of the century, where users had to pay hundreds of dollars. Prices start at US$19.99 per month for Photoshop and Lightroom, though users can get better value for money by signing up for all Adobe Creative Cloud apps.

The software is available for both Windows and macOS. There's also a web version, with a slightly simplified user interface (Adobe states that the web version of Photoshop is optimized for Chrome). As of November 2019, there's even a version of Photoshop available for the iPad. This gives it an edge over GIMP, which has no mobile version, but naturally, not all features are supported.

Photoshop files use the default proprietary file extension PSD (Photoshop Document), which stores images with support for all the program's features, including masks, transparency, text, and alpha channels.

GIMP

One of the dangers of following a SaaS model and using proprietary formats is that if a user ends the subscription, they could be left with a ton of files they can't open. GIMP shines in this area, as it's capable of opening PSD files, though it may not support all features, such as Photoshop files, with smart objects or vector masks and layers.

GIMP's default XCF (eXperimental Computing Facility) image format is open-source, but follows an ad hoc standard. The upshot is that while it's relatively easy to open your Photoshop files in GIMP, there's no automatic way to do this with XCF files in Photoshop. Luckily, GIMP supports exporting files in PSD format.

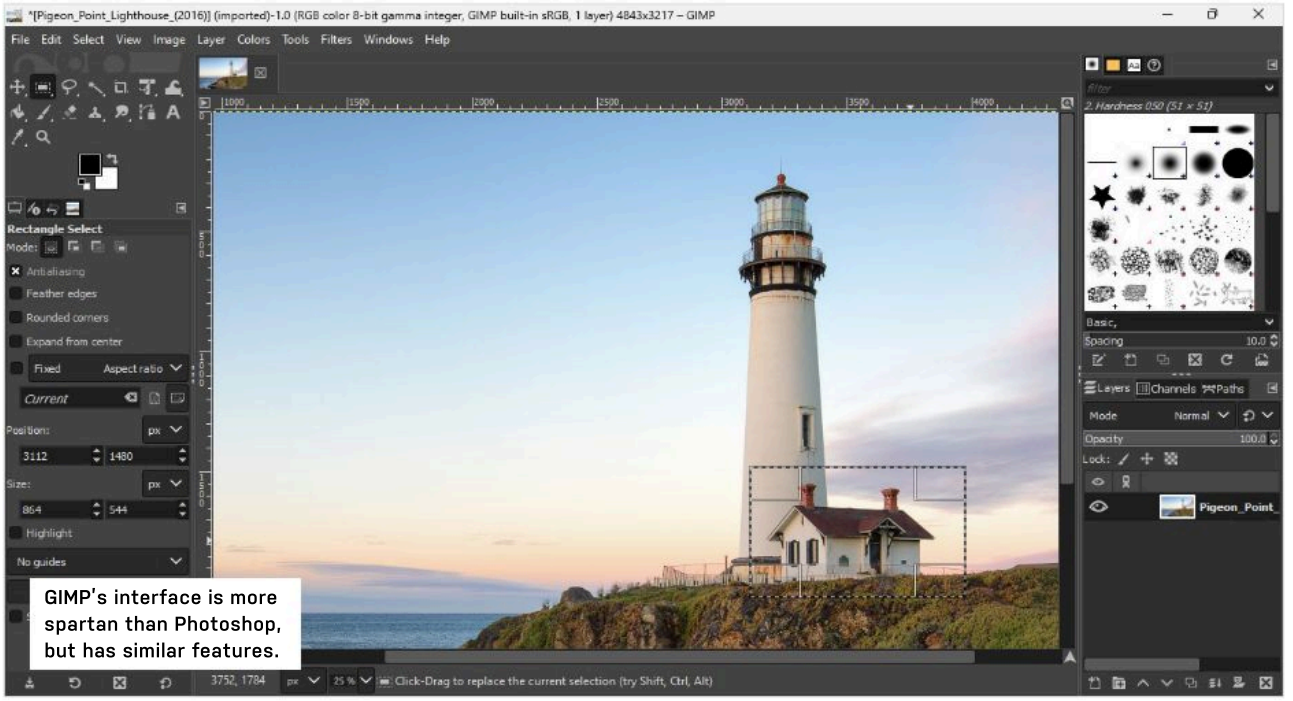

GIMP's interface is more spartan than Photoshop, but has similar features.

As free software, anyone wanting to use GIMP only needs to download via the website or Microsoft store.

Installing Photoshop is more involved due to the Creative Cloud Subscription. We downloaded Photoshop onto our test machine, and had to create an online account to receive a seven-day free trial. This made for a much longer setup time for Photoshop (15 minutes) versus GIMP (four minutes).

Newcomers will appreciate the way Photoshop takes you by the hand during setup. For instance, users are asked to gauge their level of experience. You can also specify how you plan to use Photoshop, e.g. to edit images or create graphics/visual effects. This is where we found the 'Discover Panel', which contains a mix of video tutorials, tools, and tips.

Photoshop's current hot feature is 'Generative Fill', showcased through an interactive tutorial based around an image of a house on a mountain. Users can select an area, then enter a prompt to alter the image. Generative Fill can also add new content to images, such as a tree.

The huge number of features and plugins available can be overwhelming, which is why we were impressed to see that Photoshop also has a 'Core Tools Workspace', with only the essentials.

Features like these make it easy to see why Photoshop is a popular choice with industry professionals. GIMP has no native generative AI features—the closest we could find was the 'Stable Boy' plugin, which makes use of Stable Diffusion's WebUI API. With well-crafted prompts this could provide similar functionality but is more difficult to set up.

Of course, Photoshop's range comes at a price. On our test machine, the combined install footprint of Adobe Photoshop 2024, Creative Cloud, and 'Genuine Service' (which apparently detects fake Adobe products) was over 9.3GB. GIMP's footprint was just under 1.2GB.

Photoshop also requires heavy system resources—the minimum amount of RAM is 8GB (with 16GB recommended), and it must use at least 1.5GB of GPU memory. There's no official support for graphics cards over seven years old.

GIMP, by contrast, doesn't specify any minimum system requirements, though it can make use of a 'tile cache', which operates as a type of paging file for the hard disk to process graphics.

Post-install, you'll need to be patient as GIMP loads for the first time. Introductory video tutorials are absent, but basic tools such as move, select, crop, and fill are in the top left-hand pane. On selecting one, all options are listed in the pane below.

As simplistic as this interface is, further help is available via tooltips and the GIMP Help Browser. Like everything with GIMP, this has been penned by volunteers, so may not be as comprehensive as Adobe documentation. Still, the GIMP project page links to dedicated forums and IRC channels for users with image-editing questions.

Overall, while Photoshop may have a wider range of features and plugins, we've tried to avoid saying it can do certain things GIMP can't, as native features can usually be recreated using scripting and plugins. For instance, Resynthesizer offers much the same functionality as Photoshop's 'Content Aware Fill'.

Still, features like Generative Fill, Smart Objects, and integration with Adobe Stock images gives Photoshop an edge when it comes to advanced image editing, provided you're able and willing to pay. If you don't edit images professionally, GIMP is likely to provide all the functionality you'll need for free. –NATE DRAKE

VERDICT 9

PHOTOSHOP

PHOTOSUP A huge range of advanced features and plugins; Excellent support.

PHOTOSTOP Ongoing expense of subscription and purchase of proprietary plugins.

From $22.99 per month (7-day free trial), www.adobe.com/products/photoshop.html

VERDICT 8

GIMP

GIMPRESSIVE A free editor with a simple interface; Advanced editing features.

GIMPOSSIBLE Native plugins don't support bleeding-edge features like generative AI.

Free, www.gimp.org

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> Ideal Pixel Density
> Where's My Case?
> Best of the Supes

Dinosaurs go digital

This is my first time reaching out, although I've wanted to write in on numerous occasions with questions or feedback. I've been in the semiconductor manufacturing industry since 1987, and have to admit, I'm a workaholic. I have also been subscribing to *Maximum PC* since the days that Alan Dexter was editing, but I've been reading it on and off since the 2000s.

Your deep dives into the history behind the hardware and software are really amazing. Your coverage of the manufacturing side of the industry is great, too, and I can relate to it, since that's been my career for many years. I think it's important for readers to understand how much goes into creating not just computer components, but all of the devices that make the world run.

I started out in 1987, learning to build ion beam etching and deposition equipment. I've assembled and repaired many systems, visited numerous places, and spent more time than I'd like in cleanrooms. 37 years later, I'm still in it and servicing the equipment. It's a great field to be in, and I really appreciate your coverage of the industry!

The motivation for this long-overdue email was the letter from B. Frost titled 'No love for digital' in the Letters section of your December 2023 issue. I could relate to some of what he was upset about in that I went through a similar phase when you guys switched over to a digital format. Like them, I call myself a dinosaur, maybe even a 'Neanderthal'. At the time, I was going to write in to say that I was disappointed, but your magazine was so great that I would 'adapt to change'—I think we all resist change at times.

Maybe all B. Frost needs is a good portable platform to read on. When it comes to portables, I've been an Apple guy for over 20 years—iPods, iPhones, and iPads. I've been using iPads for over ten years, and for the last three have been using a 4th gen iPad Pro 12.9" for work and play.

This is a great device to read magazines on, especially yours, where the layouts are done well and images are of high quality. Having the iPad is what made my transition so easy. Maybe your reader would be willing to invest in a large tablet, since he can afford to build a PC every year! He would also need to invest in a good case and/or stand combo, since the larger tablets can be heavy.

Finally, I want to say thank you to everyone at Future/*Maximum PC* for such a great publication. The writing, research, variety, photos are all top-notch—it's obvious to me how dedicated you all are and how much you enjoy technology.

Please pass this on to everyone there how much I/we appreciate what you all do for us.

–J. Karl McHenry

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: Thanks for this letter, we're very glad to have you as a long-term reader. And thank you for your kind words regarding the magazine, although it helps that we certainly enjoy putting it together for you.

On the subject of digital reading, we suspect that one of the reasons you enjoy using Apple products for digitally leafing through *Maximum PC* comes down to pixel density. Apple has made a point of targeting a minimum apparent pixel density for the screens on pretty much all its devices for some years now.

Apple calls these screens 'Retina' displays, the idea being that the pixels are small enough to be essentially undetectable by the human eye's retina. That's possibly hyperbole in many cases. However, Apple does ensure enough pixel density for a decent high-DPI experience across its devices. The result is super-sharp images and lovely, crisp fonts, and that makes a well-produced digital publication like *Maximum PC* look great.

Granted, Apple devices also tend to carry a premium next to those manufactured by Google, Amazon, or Samsung, and you don't need to buy Apple to get a high-DPI

↘ submit your questions to: editor@maximumpc.com

experience. Many laptops and tablets offer just that. Our advice would be to aim for something with at least 200 pixels per inch. If in doubt, you can use an online DPI calculator to work out whether a device can deliver on that remit.

Thanks again for your letter—we hope you continue to enjoy the magazine for many years to come.

Blast from the past

In the March 2024 issue of *Maximum PC*, there seems to be a phantom from the past. When I first saw it, I had to look twice, and sure enough, it was her. Like a girlfriend from the past you never forget, nor get over, her beautiful curves sang to my eyes and made my heart begin to race all over again. I thought, is that really her? I hadn't seen a photo of her in ten years.

I have her sister in a box in the attic, and I get her out and play with her from time to time, but to see her grace the pages of *Maximum PC* again took me right back to that torrid love affair.

In the March 2024 issue, on page 65 in the top-right corner of the page, you'll see her under the title, 'Turn Your Green PC'. It's a side shot of the best PC case I've ever owned. A fabulous feat of engineering from Antec, the 'P182', sitting there in all her majesty, opened to reveal her inner beauty.

My first thought was, "Yes, they're bringing it back!" But after a moment, I realized it was probably just a stock photograph, and wondered how your layout people managed to get a photo from the mid 2000s, and how it ended up in the March 2024 issue in an article that discusses making your PC green? There was nothing green about her; she was a demon of the evening, and looked as if she could destroy you with just a glance.

The enormous black monolith is something I've tried to find an equal to ever since they went out of production. I just can't find anything close. I've searched for any leftover P182s, and found one in Dallas, but it was the silver edition—not nearly as menacing as her midnight twin. They're quite rare, but they were solid as a rock. It's a beautiful PC case that had everything, built like a Spitfire.

I'll wait to hear the story of how an Antec P182 made the pages of the March 2024 *Maximum PC* issue. One of these days, with luck, I will find her equal, but until then...

Other than that, thank you for your hard work on the best PC magazine ever!

– Steven M. Sherk

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS:

Thank you for this letter, it certainly tickled me, hearing your response to seeing an image of your favorite classic chassis! There are just certain computer cases that will always be memorable to PC builders—my equivalent is the BitFenix Prodigy Arctic White that housed my first ever mini-ITX build. But I digress...

The story behind the Antec image is that it was provided by the writer, Nate Drake, as it's a nice side shot to highlight the key components that use energy, and therefore contribute to the carbon footprint of running your PC. However, I must admit that I didn't spot just how old the system was in this image when I edited it—that entire front section to house mechanical and SATA drives is pretty much redundant on any modern PC these days.

You make an interesting point on how case design has changed over the last 10-20 years, though—you liken the Antec to a Spitfire, and I presume you mean in terms of the fact that it was tough and would go for years. Today's cases are much more focused on delicate materials like glass, especially the newest fishbowl designs, which we've been using in builds and reviewing recently. Yes, they look great, and show off what's inside, but we have to treat them very delicately when transporting them home after a build to do the testing and benchmarking. I imagine that there's a market for people who still want rugged cases, especially those who maybe travel around for LAN party and eSports events.

My memorable case: the Bitfenix Prodigy Arctic White from 2012.

© BITFENIX

GPU showdown

Now that you've put all the RTX Super cards through their paces, which (if any) is the one to go for? Or should I be looking at AMD/Intel as well? I should mention that I game at 1440p and like to play story-driven games like *Cyberpunk 2077* and *Horizon Forbidden West*.

–G. Brown

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS:

I put this question to our builder, Zak Storey, in our 'Story of the Build' video this issue (embedded on page 16), and he said the RTX 4070 Super, hands down. It offers good value, with 20 percent more power for the same $599 price. It will allow you to game at your 1440p resolution at high or max settings, and it comes with all of the DLSS 3 goodness that will make your story-driven games look great and play at smooth frame rates. This month, I personally tried *Horizon* on a HP Omen Transcend 14 running an RTX 4060 (and a lowly 65W one, at that), and was impressed with how smoothly it ran with DLSS and frame generation turned on.

The AMD Radeon RX 7800 XT is worth a look at around $530 for its performance, but its AI and ray tracing features lag behind Nvidia's. Hope this helps!

THE BUILDS

THIS MONTH'S STREET PRICES...

IT FINALLY LOOKS like prices are stabilizing somewhat, certainly compared to last issue. Our blueprints this month have had some significant price decreases across the board, particularly for AMD processors. We've gone a bit off the cuff with these budgets builds, too, but more on that in a little bit.

AMD saw the Ryzen 5 7600 fall by $30, and the motherboard drop by another $6, along with a plethora of price increases across the SSD, RAM, PSU, and motherboard, leading to a $19 overall price increase, nullifying our decreasing CPU there.

Fortunately, we decided to go a bit renegade this month, re-testing some of Intel's latest GPUs, and have opted to swap the RX 7600 for an Intel Arc A750 8GB. Both cards feature 8GB of VRAM, and without a doubt the 7600 is the higher-performing card (averaging around 86fps in our 1080p testing suite versus 74 for the Arc A750), but where Intel has the edge is on price. ASRock's challenger D variant clocks in at $190, giving us 0.39fps per $ spent, versus the RX 7600's 0.33fps per $ spent. Both of those are way above what we'd consider good. Our favorite value GPU right now, the 4070 Super, only averages 0.26fps, for example.

For our Intel build, we've actually done the old switcheroo, and subbed in the 7600 from our AMD build. Still far cheaper than the 4060 we had from last issue, it's more than capable at 1080p, and even 1440p if we're honest, making it a surefire pick. We've also dropped down to a slightly lower-spec 14400F CPU. It saves us $40 in total, but does mean that we lose four cores.

Lastly, we've switched some of the memory out too, opting for a higher-spec 6,000 MHz kit for our AMD build (7th Gen Ryzen's sweet spot right now) and splurging on a super-budget $56 kit for Intel, as it's less dependent on memory speeds. That's allowed us to get our AMD build down to the sub-$1,000 mark once more, and our Intel build nearly down to three figures, but not quite.

AMD INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$68
Mobo	Asus Prime X670-P ATX AM5	$194
CPU	AMD Ryzen 5 7600	$199
GPU	ASRock Challenger D Arc A750 8GB NEW	$190
RAM	16GB (2x8GB) Kingston Fury Beast RGB @ 6000 C36 NEW	$75
SSD 1	1TB Solidigm P41 Plus PCIe 4.0 M.2	$70
SSD 2	1TB Adata Legend 800 PCIe 4.0 M.2	$64
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$987**

INTEL INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$68
Mobo	MSI Pro Z790-S WiFi ATX LGA1700	$170
CPU	Intel Core i5-14400 NEW	$210
GPU	ASRock Challenger OC RX 7600 8GB NEW	$260
RAM	16GB (2x8GB) Teamgroup Elite Plus DDR5 @ 4800 C40 NEW	$56
SSD 1	1TB Solidigm P41 Plus PCIe 4.0 M.2	$70
SSD 2	1TB Adata Legend 800 PCIe 4.0 M.2	$64
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$1,025**

SIMILARLY TO OUR BUDGET BUILDS, the mid-range has received a pretty healthy hammering on the pricing front. Even without our intervention, our AMD list has had its mobo slashed by $25, its GPU by $30. It isn't all buttons and charts, though, certainly not for AMD, as the 7700X we did have pegged for this issue had its price skyrocket by a phenomenal $50, quickly wiping away any savings we'd accrued from elsewhere.

To that end, we've switched out two key parts. That CPU, dropping to the Ryzen 7 7700 instead, the non 'overclocked' version (you can still tweak it in Ryzen Master, and it has the same core count), and we've swapped the memory for the OLOy Blade RGB kit. It's the former that made the biggest difference, though, as it's even a drop on last month's price. If we keep the CPU the same, it saves us a total of $80 on the overall price.

As for the memory, the Adata Lancer Blade is seemingly no longer in stock, so we've picked up this OLOy kit instead. You can actually go cheaper than this, and still retain a 32GB kit of 6000 with TeamGroup. However, the kit we've found comes with a staggering C48 latency, which you definitely don't want.

As for Intel, it saw an array of price drops and shuffles, too. The motherboard and CPU cooler in particular saw the biggest drops at $20 apiece, and even our trusty 14600KF had its price reduced by $5. To really cement the price cutting, though, we've again switched out our Nvidia unit for something blue in the form of Intel's A770. Not as potent as the 4060 Ti, what it lacks in firepower it makes up for with a stellar price, and average frame rates of 61.57 at 1440p (with hyper-aggro ray tracing enabled—non DLSS-enabled *Cyberpunk* is included in that figure, we should add) make it a solid pick for a mid-range gaming PC.

All in all, our Intel build has fallen by a phenomenal $114 this issue, with AMD being $75 cheaper, too.

AMD INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow	$109
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold	$95
Mobo	MSI PRO X670-P WiFi ATX - AM5	$200
CPU	AMD Ryzen 7 7700 NEW	$260
Cooler	Corsair A115 Air Tower	$100
GPU	ASRock Radeon RX 7700 XT Challenger OC	$400
RAM	32GB (2x16GB) OLOy Blade RGB DDR5 @ 6000 C36 NEW	$92
SSD 1	1TB Corsair MP600 PRO LPX M.2 PCIe 4.0 SSD	$90
SSD 2	1TB Adata Legend 800 M.2 PCIe 4.0 SSD	$64
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$1,442**

INTEL INGREDIENTS

PART		PRICE
Case	NZXT H7 Flow	$109
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold	$95
Mobo	MSI Z790 Gaming Pro WiFi ATX	$189
CPU	Intel Core i5-14600KF	$280
Cooler	EK AIO Basic 360 - 360mm AIO	$109
GPU	Sparkle Titan OC Arc A770 8GB NEW	$300
RAM	32GB (2x16GB) Corsair Vengeance DDR5 @ 6000 36 NEW	$93
SSD 1	1TB Corsair MP600 PRO LPX M.2 PCIe 4.0	$90
SSD 2	1TB Adata Legend 800 M.2 PCIe 4.0 SSD	$64
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$1,361**

YES, we managed to get both of these builds to exactly $3,119, completely by accident. Perhaps unsurprisingly, these saw the biggest shift in pricing this month, with the CPU on AMD and the GPU on Intel each receiving a $50 price bump. The RTX 4080 Super stock situation continues to be a complete pain to deal with. There are GPUs out there, but as we mentioned in our review this issue, they all seem to be around the $1,200 mark. There are a few GPUs listed for less than $1,100, but they're typically on pre-order or out of stock.

We've gone with PNY's GeForce RTX 4080 Super for our Intel build to counter that, but we're still sat at $1,050 even doing that. Other shocks included a fairly significant bump to our 48GB kit of Trident Z memory, going up by $15 and putting it outside of our tolerance range. We've gone for a slightly lower-spec model in the Ripjaws instead, dropping from 6,800 MHz to 6,400 MHz, but you're not likely to see the difference in the real world, and that price drop down by another $35 is tasty.

Speaking of RAM, our AMD build also did the DDR shimmy, ditching its old kit for a shiny set of 64GB 6,400 MHz C30 kit from Mushkin. It's not quite as fancy as some of the competition, but that top-line spec sheet is to die for. One of the last changes that permeates across both systems is our SSD 2 choice. We've ditched the P3 Plus and its shoddy overall performance, and gone for one of our favorite latest and greatest PCIe 4.0 drives, the Crucial T500. Not only does it monster the sequential transfers for a PCIe 4.0 SSD, but its random 4K performance is also stellar. The only downside being that we are giving up 2TB of storage to facilitate that. If you do need the extra capacity then the P3 Plus is still the drive to go with, but for us, speed is king.

On the surface, it has been a rough issue for our Turbo builds, particularly with AMD and the 7950X skyrocketing by $52, but we still managed to keep the price down with some tweaks, cutting the overall cost by $65. Intel saw an even bigger price cut, with $110 lopped off the total cost, even with Nvidia's RTX 4080 Super being rarer than, well, an RTX 40 series card.

AMD INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$140
PSU	Super Flower Leadex Platinum SE 1200W - 80+ Platinum	$160
Mobo	Asus Prime X670E Pro WiFi - AM5	$300
CPU	AMD Ryzen 9 7950X	$600
Cooler	NZXT Kraken Elite 360 RGB - 360mm AIO	$277
GPU	Sapphire Nitro+ RX 7900 XTX 24GB	$1,030
RAM	64GB (2x32GB) Mushkin Redline ST @ 6400 C30 NEW	$175
SSD 1	2TB Corsair MP700 PCIe 5.0 M.2	$249
SSD 2	2TB Crucial T500 PCIe 4.0 M.2 NEW	$156
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$3,119**

INTEL INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$140
PSU	Super Flower Leadex Platinum SE 1200W - 80+ Platinum	$160
Mobo	Gigabyte Z790 Aorus Elite AX-W ATX	$370
CPU	Intel Core i9-14900KF	$535
Cooler	NZXT Kraken Elite 360 RGB - 360mm AIO	$277
GPU	PNY Verto Overclocked RTX 4080 Super 16GB NEW	$1,050
RAM	48GB (2x 24GB) G.Skill Ripjaws DDR5 @ 6400 CL36 NEW	$150
SSD 1	2TB Corsair MP700 PCIe 5.0 M.2	$249
SSD 2	2TB Crucial T500 PCIe 4.0 M.2 NEW	$156
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$3,119**

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US LLC, 130 West 42nd Street, 7th Floor, New York, NY 10036. USA. Website: www.futureus.com. Future US LLC also publishes *MacLife*, and *PC Gamer*.

UNEARTH THE FULL STORY BEHIND THE LEGEND OF ZELDA

Explore every game in the series and what makes each one so great. If you're a fan, you'll enjoy the trip through time. If you're a newcomer, we offer you this sage advice: "It's dangerous to go alone, take this!"

ON SALE NOW

FUTURE

Shop easily online at:

magazinesdirect.com

bit.ly/3xbOGXq

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